Hybrid Particle-Continuum Simulations of Low Knudsen Number Hypersonic Flows

Thomas E. Schwartzentruber[*], Leonardo C. Scalabrin[†] and Iain D. Boyd[‡]

Department of Aerospace Engineering, University of Michigan, Ann Arbor, MI 48109

A modular particle-continuum (MPC) numerical method is used to simulate steady-state hypersonic flows which exhibit local regions of non-equilibrium embedded within mainly continuum flow fields. The MPC method loosely couples direct simulation Monte Carlo (DSMC) and Navier-Stokes (NS) methods which operate in different regions, use different mesh densities, and are updated using different sized timesteps. The MPC method is applied to both a hollow cylinder flare and planetary probe geometry and results are compared with full NS and DSMC simulations as well as with experimental data. MPC simulations are demonstrated to reproduce experimental and full DSMC simulation results for surface and flow field properties including velocity slip, temperature jump, thermal non-equilibrium, heating rates, and pressure distributions with high accuracy. The hollow cylinder flare problem provides an insightful test case for the MPC method, however, it is found un-suitable for practical hybrid simulation. Orders of magnitude variation in mean-free-path for the planetary probe problem make it an excellent candidate for hybrid simulation. For this case, MPC results are obtained approximately 12.5 times faster than a full DSMC simulation while requiring 20% of the memory.

I.   Introduction

HYPERSONIC vehicles generate multi-scale gas flows as a result of both large variations in mean-free-path and short characteristic length scales. A particular scale of interest that captures the degree of collisional-equilibrium within a gas is the Knudsen number (Kn), which is defined as the ratio of the mean-free-path (λ) to a characteristic length scale of interest (l_c). At most altitudes λ is small and there is an enormous number of gas-particle collisions occurring over most practical length scales ($Kn < 0.01$). As a result, for the majority of aerospace flows, the gas behaves as a continuous fluid for which the continuum Navier-Stokes (NS) equations provide an accurate physical model that can be solved efficiently using techniques from Computational Fluid Dynamics (CFD). However at high altitudes where λ may approach the characteristic dimension of a vehicle ($Kn > 0.1$), the continuum assumption breaks down and a particle or kinetic description of the gas is required. The direct simulation Monte Carlo (DSMC) method[1] is currently the most popular method for high-speed non-equilibrium flows. DSMC directly simulates a gas flow by tracking a representative number of simulated particles 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 λ and the mean-free-time (τ_c) respectively.

Generally, the global Knudsen number (Kn_g) which uses the free-stream value of λ and an overall dimension of the vehicle for l_c, provides a good indication of whether the flow is continuum or non-equilibrium. However, difficulty arises for low Knudsen number flows ($Kn_g \leq 0.001$) where hypersonic velocities induce localized regions of non-equilibrium imbedded within a mostly continuum flow. For example, DSMC may be required to accurately model the flow over sharp leading edges, inside shock waves, boundary layers, and wake regions, while remaining computationally expensive for the majority of the flow. For blunt body flows, in order to obtain accurate results in the wake region without resorting to a full DSMC simulation, "zonally decoupled" DSMC-NS simulations have been performed.[2–4] Here, the dense fore-body flow is computed separately using a NS solver and the exit-plane solution is specified as inflow conditions for a decoupled DSMC simulation of the entire wake region. Although accurate and more efficient than full DSMC, this approach requires a-priori knowledge of a suitable particle-continuum interface and

---

[*]Graduate Student, Student Member AIAA. Email: schwartt@umich.edu.
[†]Graduate Student, Student Member AIAA. Email: lscalabr@umich.edu.
[‡]Professor, Associate Fellow AIAA. Email: iainboyd@umich.edu.

therefore lacks generality. Various researchers have proposed hybrid numerical methods which adaptively re-position the particle-continuum interfaces during a single hybrid simulation and couple particle and continuum regions by transferring information across this interface.[5–10] Such hybrid DSMC-CFD algorithms have been under development for over a decade, and they have yet to clearly show their promise of significant speedup over full DSMC simulation. A more detailed overview of these various methods and comparison with the hybrid numerical algorithm used in this article can be found in Ref. 11. In this study, the modular particle-continuum (MPC) algorithm[11] is used to simulate steady-state hypersonic flows. The MPC method was initially developed and tested for 1D normal shock waves[12] as well as for hypersonic flow over a 2D cylinder geometry.[13] The MPC algorithm loosely couples DSMC and NS regions which have different mesh densities and are updated using different timesteps. Such an approach enables both spatial and time-scale decoupling while lending itself to a modular implementation[14] which utilizes existing, state-of-the-art, DSMC and NS codes (unmodified) within the hybrid code. For flow over a 2D cylinder, the MPC method has been shown to reproduce the flow fields, surface properties, and local velocity distributions predicted by full DSMC simulations 1.4 to 3.3 times faster while requiring less memory. Recently, a similar loosely-coupled approach was used by Wu *et al.*[15] to simulate hypersonic flow over a wedge as well as expansion of nitrogen gas from a 3D nozzle into a near-vacuum. The authors use a shell script to couple existing DSMC and NS codes. The resulting numerical cycle differs from the MPC algorithm in a subtle but important manner regarding *when* information is transferred between DSMC and NS regions; an issue which is discussed in section III of this article. Finally, while their hybrid solution for the wedge flow agreed well with full DSMC flow field results, the simulation was reported to take longer than a pure DSMC simulation.

In this article, the MPC method is applied to two popular benchmark problems; a hollow cylinder flare geometry and a planetary probe geometry. These axi-symmetric low Kn_g flows provide a significant challenge for the MPC algorithm compared with the previous MPC research cited above. High-quality experimental data is available for both problems and has been used extensively to validate both DSMC and CFD methods in the literature. A detailed description of both problems as well as a comparison of experimental results with pure NS and pure DSMC simulations is presented in section II. A brief overview of the MPC numerical cycle is presented by showing the progression of a hybrid simulation in section III. The accuracy of the final MPC solutions with respect to flow field features, velocity slip, temperature jump, thermal non-equilibrium, and surface properties is investigated in section IV. The computational efficiency gained over pure DSMC simulation is also discussed in section IV and conclusions are presented in section V.

II.   Navier-Stokes and DSMC Simulation

A.   Numerical Methods

Particle regions are simulated using MONACO,[16] a general, cell-based implementation of the DSMC method.[1] The variable hard sphere (VHS) collision model is employed which results in the following macroscopic viscosity model:[1]

$$\mu = \mu_{ref}\left(\frac{T}{T_{ref}}\right)^{\omega}, \qquad \mu_{ref} = \frac{15\sqrt{\pi m k T_{ref}}}{2\pi d_{ref}^2(5-2\omega)(7-2\omega)} \; . \tag{1}$$

All numerical results presented in this article 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[17] 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.

Continuum regions are simulated using the LeMANS code.[18] For the results of this article, 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, 2D laminar, compressible, Navier-Stokes (NS) equations. The viscosity in the NS solver is modelled using Eq. 1 in order to match exactly the viscosity model used in DSMC. LeMANS solves this set of equations using a finite-volume formulation. The inviscid fluxes between the mesh volumes are discretized using a modified form of the Steger-Warming Flux Vector Splitting[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. 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, no-slip conditions are applied to both velocity and temperature on all surfaces for the results of this article. Specific details of the numerical method are contained in Ref. 18.

American Institute of Aeronautics and Astronautics
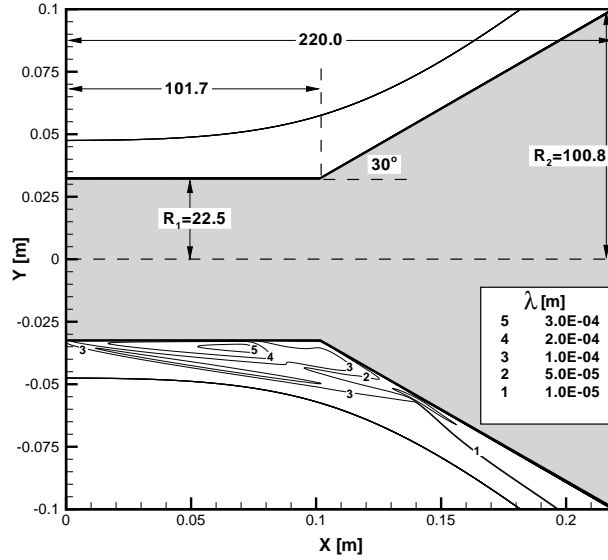
B. Hollow Cylinder Flare

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

Table 1. Experimental flow conditions for hollow cylinder flare and planetary probe problems.

Case	M_∞	$U_\infty \ [m/s]$	$\rho_\infty \ [kg/m^3]$	$T_\infty \ [K]$	$Tv_\infty \ [K]$	$Tw \ [K]$	Kn_g^*
Hollow Cylinder Flare - CUBRC Run 11 -	12.4	2484.1	5.566×10^{-4}	95.6	2486.7 (not used)	297.2	0.0008
Planetary Probe - SR3 Case 3 -	20.5	1633.0	4.660×10^{-4}	15.3	15.3 (not used)	300.0	0.001

$^*Kn_g$ is based on the length of the cylinder, the diameter of the planetary probe, and the VHS value for λ

The hollow cylinder flare geometry is detailed in 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 the stronger shock generated by boundary layer separation at 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 separation region. This strongly coupled and highly nonlinear flow provides a challenging case for the validation of both CFD 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 organized by Holden at the 2001 AIAA Aerospace Sciences Meeting in Reno, NV.[20,21] For the Run 11 conditions, Navier-Stokes simulations[22] predicted a larger separation 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 CFD code. Further investigation by Candler *et al.*[23] determined that the free-stream conditions had a significant amount of thermal non-equilibrium due to frozen vibrational relaxation during the rapid nozzle expansion in the LENS facility. As a result, new revised free-stream conditions were defined for this case. The current study uses these revised flow conditions which are listed in Table 1. 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 article. The LeMANS code as described above is used to solve the NS equations for the hollow cylinder flare geometry detailed in Fig. 1 and for the flow conditions listed in Table 1. Solutions are obtained on three mesh densities of 550x125, 1100x250, and 2200x500 cells for mesh1, mesh2, and mesh3, respectively. Results for the coefficient of pressure [$C_P = (p - p_\infty)/(0.5)\rho_\infty u_\infty^2$] are displayed in Fig. 2(a) in order to verify grid-independence

American Institute of Aeronautics and Astronautics

and compare the NS solution with experimental results. As seen in Fig. 2(a) the solutions on mesh2 and mesh3 are virtually identical and the NS equations predict a separation region significantly larger than measured experimentally. As a result, the separation shock forms earlier which causes the strong interaction region, and thus the peak surface pressure, to occur further downstream than observed experimentally. Although not shown, the NS equations (using no-slip boundary conditions) significantly over-predict surface heating and shear stress right at the leading edge of the hollow cylinder. The surface pressure distribution predicted by LeMANS in Fig. 2(a) is in excellent agreement with simulations performed by Candler et al.[23] with respect to the precise location of flow separation, the peak surface pressure, and the over-prediction of surface pressure along the flare.
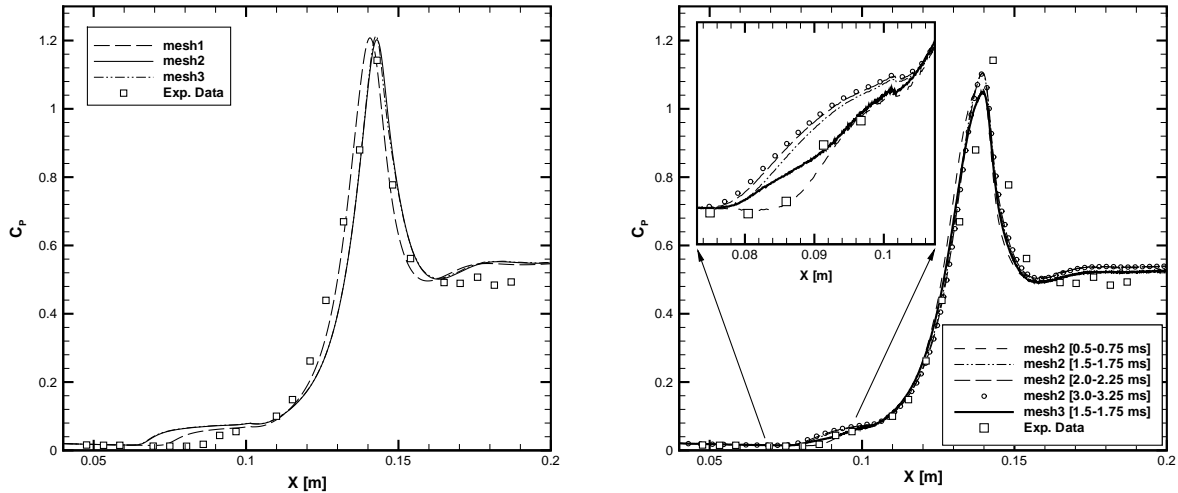
(a) Navier-Stokes results (LeMANS).

(b) DSMC results (MONACO).

Figure 2. Navier-Stokes and DSMC grid convergence studies for the hollow cylinder flare geometry.

Table 2. DSMC prediction for the location of flow separation and approach to steady-state.

sampling interval [ms]	x_{sep} [m]	% change / 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%

In the blind code validation exercise, the DSMC method predicted leading edge surface properties and the size of the separation region well.[21] The largest discrepancy with experimental results was found downstream of the flare junction where λ becomes small. This can be seen in the bottom of Fig. 1 where 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[24] have improved the resolution and produced excellent results compared with the experimental data. The MONACO code as described above is used to calculate the flow field and surface properties for the Run 11 flow conditions listed in Table 1 using both mesh2 and mesh3. The least resolved region in mesh2 and mesh3 is at the surface along the flare where cell dimensions are 4λ and 2λ respectively, in the direction normal to the wall. The mesh3 simulation uses twice as many simulated particles as the mesh2 simulation and both contain at least 15 particles per cell. The flow is axi-symmetric where cell-volumes in DSMC become larger as distance from the axis increases and therefore contain more particles. However, since 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 timesteps 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. In order to demonstrate this, Fig. 2(b)

American Institute of Aeronautics and Astronautics

shows the MONACO solution on mesh2 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 2. 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 this same case performed by Moss and Bird[24] draw the same conclusion and predict a similar level of agreement with experimental data as the MONACO results in Fig. 2(b). Compared with the NS results in Fig. 2(a), DSMC is seen to better predict the extent of flow separation for the Run 11 conditions. Finally, comparison of the solutions obtained on mesh2 and mesh3 (for the same sampling interval) demonstrates that mesh2 provides adequate resolution for a DSMC simulation. Although the mesh3 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 mesh2 solution. It is interesting to note that for the Run 11 conditions, the mesh density required to obtain a grid-independent NS solution (mesh2) is the same as that required by a corresponding DSMC simulation.

C. Planetary Probe
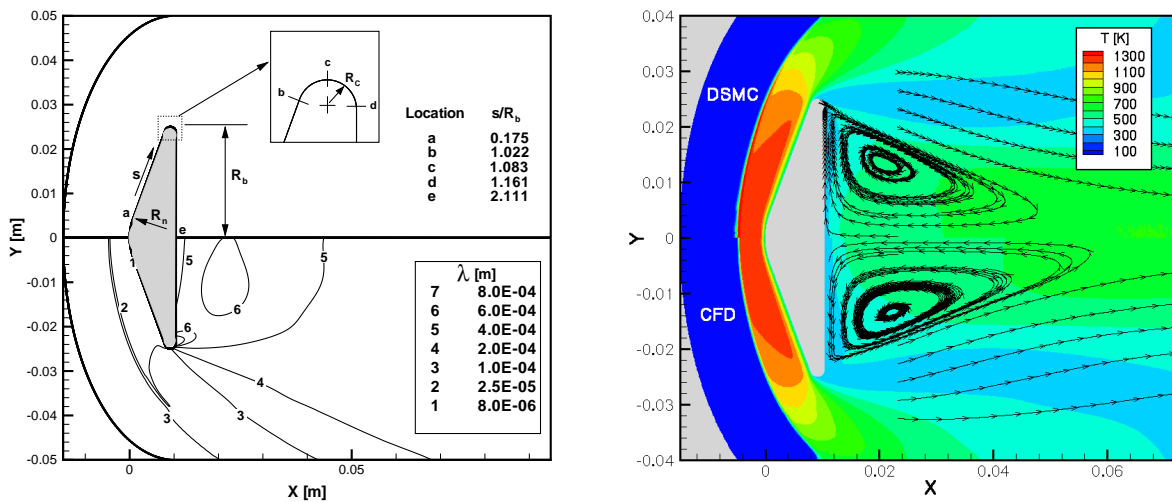
(a) Planetary probe geometry (dimensions in mm) and variation of λ.

(b) NS and DSMC translational temperature.

Figure 3. Particle and continuum simulation of the planetary probe configuration.

The planetary probe geometry is detailed in the top of Fig. 3(a) where the diameter of the probe is 5 cm ($R_b = 0.025\ m$), the nose radius $R_n = 0.5 \times R_b$, and s is the distance around the surface. Experimental data was obtained in the SR3 wind tunnel in Meudon, France. Three test conditions were considered where the free-stream was nitrogen at a nominal Mach number of 20 with free-stream densities of 1.7×10^{-5}, 5.2×10^{-5}, and $46.6 \times 10^{-5}\ kg/m^3$, corresponding to Kn_g values of 0.03, 0.01, and 0.001 respectively. Experimental data and numerical simulation results for these cases are summarized in Ref. 25. A "zonally decoupled" approach[2] has been applied to this same case where the dense fore-body flow is computed separately using a NS solver and the exit-plane solution is specified as inflow conditions for a decoupled DSMC simulation of the entire wake region. It should be noted that although the experiment supported the model with a sting, the zonally decoupled simulations did not include the sting. In order to investigate the size of the wake vortex and compare with the zonally decoupled results, the present study also does not include the sting. The low density (high Kn_g) cases are found to lie entirely in the rarefied (non-equilibrium) regime and can be solved efficiently using full DSMC simulation. This article will focus on the highest density case whose initial conditions are listed in Table 1. As depicted in the bottom of Fig. 3(a), the hard-sphere value of λ varies by 2 orders of magnitude between the stagnation and near-wake regions. In a full DSMC simulation, since the local cell size (in each dimension) must scale with λ, an enormous number of cells (and therefore simulation particles) are

American Institute of Aeronautics and Astronautics

required to simulate the fore-body flow. In addition, the flow is axi-symmetric where regions of the computational domain include the axis of symmetry. Cells far away from the axis have orders-of-magnitude larger volumes than cells near the axis and as a result would require a prohibitive number of particles to model the same density. For this reason, radial cell-weighting must be used for DSMC simulation. The effective weight in each cell (defined as the number of real particles represented by each simulated particle) is increased by a factor of 4 for each factor of 4 increase in distance from the axis of symmetry. The translational temperature fields calculated using both MONACO (DSMC) and LeMANS (CFD) are shown in Fig. 3(b). Here both DSMC and CFD methods predict the same fore-body flow and only differ near the shoulder and in the near-wake region. The temperature shock predicted by DSMC is seen to be thicker and to begin slightly upstream of the location predicted by CFD simulation. This is a well known prediction of DSMC, whereas the density shock location (shock standoff distance) is verified to be the same for both simulations. Streamlines originating from identical points are plotted in Fig. 3(b) for both DSMC and CFD solutions which show that the NS equations predict a larger vortex in the wake than DSMC. In addition, streamlines are seen to lie closer to the surface in the DSMC solution, signifying velocity slip in the shoulder and base regions of the planetary probe.

III.   Modular Particle-Continuum (MPC) Numerical Method

A.   Problem setup

The modular particle-continuum (MPC) numerical method used in this article is completely detailed in Ref. 11 and is now briefly outlined. The MPC method begins with a CFD 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[26]

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

where λ is calculated consistent with the variable-hard-sphere (VHS) model and the maximum is taken over the
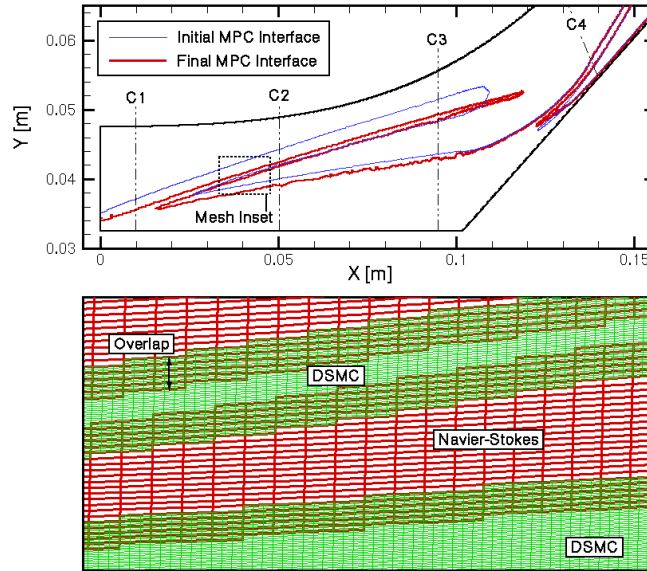

Figure 4.   Interface locations (top) and mesh refinement (bottom) for the hollow cylinder flare.

flow quantities (Q) of interest: density (ρ), temperature (T), and velocity magnitude ($|V|$). Cells in which $Kn_{GL} > Br_{cutoff}$ are labelled as DSMC cells and the remainder as NS cells, thus defining an interface between the two regions. Previous studies have recommended[26] and validated[11–13] a cutoff value of $Br_{cutoff} = 0.05$ which is the value used for the planetary probe simulation. Whereas, a lower value of $Br_{cutoff} = 0.03$ is used for the hollow cylinder flare simulation. In order to create an overlap region where both methods are used, the DSMC regions are then extended further into the continuum region by a number ($n_{overlap}$) of cells. Both the hollow cylinder flare and planetary probe simulations use a value of $n_{overlap} = 4$. Extended DSMC regions are further surrounded by a layer of DSMC boundary

cells and NS regions are surrounded by a layer of NS boundary cells. 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 Ref. 14 which also uses flow gradient information from the NS solution to set more stringent refinement in the direction of flow gradients and relax the refinement in other directions. The initial particle region (including the overlap) for the hollow cylinder flare is shown in the top of Fig. 4 and a portion of the mesh is depicted in the bottom of Fig. 4. Likewise, the initial particle region for the planetary probe problem is shown in Fig. 5(a) and a portion of the mesh near the capsule shoulder is shown in Fig. 5(b). Initially particles are generated in all DSMC cells (using a Chapman-Enskog velocity distribution[27]) such that they are consistent with the 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 towards the correct non-equilibrium solution. As this progression and relaxation occurs, DSMC and NS regions are loosely coupled and interface locations are periodically updated throughout the simulation. As an example, the final steady-state interface locations for the hollow cylinder flare and the planetary probe are shown in the top of Fig. 4 and in Fig. 5(a) respectively. Finally, since 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.
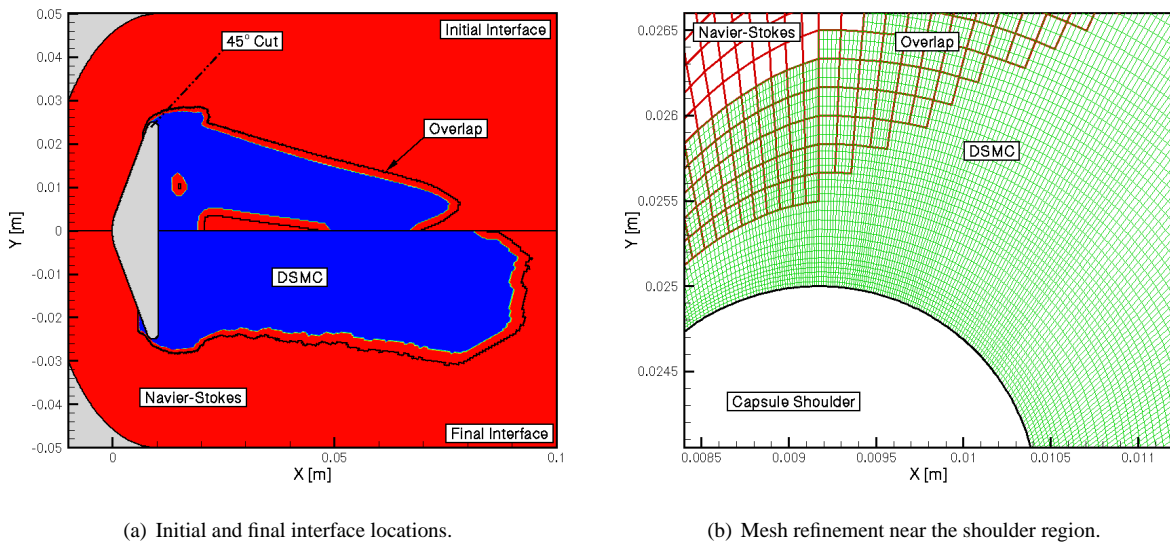

(a) Initial and final interface locations.

(b) Mesh refinement near the shoulder region.

Figure 5. Interface locations and mesh refinement for the planetary probe.

B. Numerical cycle and progression of hybrid solution

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

1. Using Kn_{GL}, setup 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 (utilizing 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 the state in any DSMC boundary cell $\longrightarrow$ return to 2.

   ELSE , if the new NS solution does not change the state in any DSMC boundary cell $\longrightarrow$ continue to 4.

4. Since 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 solvers (coupling occasionally) until the DSMC scatter and NS residual fall below threshold values.

American Institute of Aeronautics and Astronautics

The hybrid numerical cycle used by Wu *et al.*[15] is similar to the above MPC cycle and adds further evidence that a loosely-coupled approach is well suited for high-speed, steady-state flows. One difference is that while Wu *et al.* use a cumulative (sampled) average inside time-varying DSMC regions, the MPC method uses a sub-relaxation average[28] to track macroscopic variations with low scatter in DSMC regions. Tracking variations in DSMC regions with the sub-relaxation 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 towards 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]
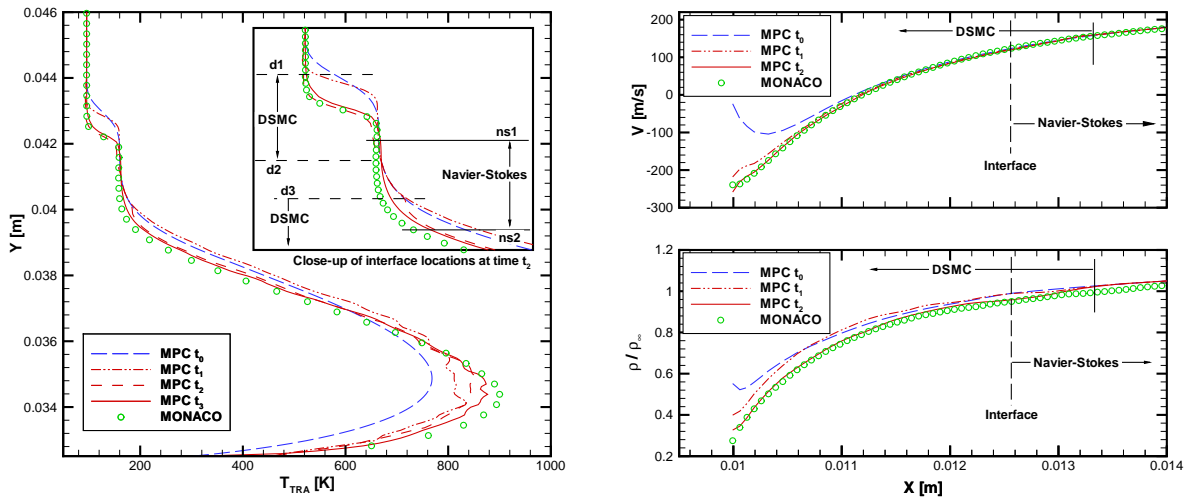
(a) Hollow cylinder flare MPC simulation.

(b) Planetary probe MPC simulation.

Figure 6. Progression and relaxation of DSMC regions within the MPC cycle.

Figures 6(a) and 6(b) show MPC solutions at various stages during step 2 of the above MPC cycle for both the hollow cylinder flare and planetary probe cases respectively. In Fig. 6(a) 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. Since the initial boundary layer development (at the leading edge) is enclosed entirely by a DSMC region (see 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 towards 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 towards 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. 6(a) 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 timesteps has progressed to the profile labelled as t_1, which at this point has not clearly progressed in the proper direction. 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 timesteps, the solution reaches the profile labelled 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 is fixed by incorrect boundary conditions at line $d3$, within this overlap region it 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 to the NS regions, the new NS boundary condition is taken on

line *ns*2. After convergence of the NS regions, the profile labelled t_3 is obtained which shows significant progression towards 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 *d*1, *d*2, and *d*3.

For the planetary probe simulation, the same progression and relaxation process is seen along the 45^o cut through the shoulder of the capsule (shown previously in Fig. 5(a)), and is plotted in Fig. 6(b). Figure 6(b) demonstrates again how the overlap region provides a buffer which allows DSMC and NS solutions to differ. The more accurate DSMC information is then used to update the NS boundary conditions. This is evident in the bottom of Fig. 6(b) where although the DSMC density at the edge of the overlap region is set by the NS solution, inside the overlap region the DSMC density differs from the NS result. When the DSMC solution is used to set new NS boundary conditions, the NS density profile will be shifted into better agreement with the full DSMC simulation (a lower density). The shifted NS solution will then provide more accurate boundary conditions for the DSMC region and the above MPC cycle will repeat itself until step 4 is reached.

IV. Hybrid Results

Previous research focusing on 1D normal shock waves[12] and hypersonic flow over a 2D cylinder,[11, 13, 14] for which $Kn_g = 0.01$, demonstrated the accuracy of the MPC method for computing flow field data, velocity distributions, and surface properties. The low Kn_g hollow cylinder flare and planetary probe problems studied in this article provide a significant challenge for the MPC algorithm compared with this 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 which significantly alters the solution in hybrid-NS regions, substantial variation in mesh density between continuum and particle regions for more complex geometries, and finally, the ability to handle axi-symmetric flows which may necessitate the use of cell-weighting factors in DSMC regions.

A. Hollow Cylinder Flare

The current implementation of the MPC method is restricted to serial (single CPU) processes. Due to computational time and memory restrictions, the initial NS solution used is that obtained on mesh1. DSMC regions are further refined to the local value of λ which results in a mesh density slightly higher than that of mesh2. A portion of this hybrid mesh was shown earlier in Fig. 4. As seen in Fig. 2(a) the NS solution on mesh1 is very close to the grid-converged result on mesh2 except that the separation region is slightly smaller. The size of the separation region on mesh1 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 mesh1 also allows testing of the mesh refinement procedures for the hollow cylinder flare.

In order 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 *C*1, *C*2, *C*3, and *C*4 (from Fig. 4) are shown in Figs. 7(a), 7(b), 7(c), and 7(d), respectively. Near the leading edge, the breakdown parameter sets up a DSMC region which completely envelopes the initial boundary layer growth and leading edge shock. As seen in Fig. 7(a), the shock/boundary layer height predicted at location *C*1 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 non-equilibrium 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 *C*2, the shock and boundary layer become more distinct and Fig. 7(b) 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 non-equilibrium and temperature jump predicted by full DSMC very accurately. The MPC profile in Fig. 7(b) also demonstrates how the changing DSMC regions have successfully set new NS boundary conditions which significantly shift the continuum solution between the shock and boundary layer. The progression of the MPC solution at cut *C*2 was detailed in section III. It is important to note in the continuum region between the shock and boundary layer that the flow is in thermal equilibrium and is adequately modelled by the NS equations. This adds further credibility to the use of the gradient-length Knudsen number[26] to detect continuum breakdown. Location *C*3 cuts through the separation region which is evident in Fig. 7(c) where the x-velocity (U) is negative. Here, DSMC is seen to predict a smaller separation region than the initial NS solution. Similar to a full DSMC simulation, an MPC simulation also requires significant time for the separation region to transition from the initial NS solution and reach a steady-state. However, since the MPC

interface envelopes the entire separation 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 since little or no interface movement signifies steady-state at which point the MPC cycle wants to lock the interfaces and begin sampling. In order to allow the MPC cycle to handle slow recirculating flows a user-defined parameter, Δt_{ss}, is introduced which specifies the minimum number of iterations (or minimum physical time) before sampling begins. This parameter is analogous to the user-defined parameter specifying the sampling interval employed in full DSMC simulations. The MPC solution along cut
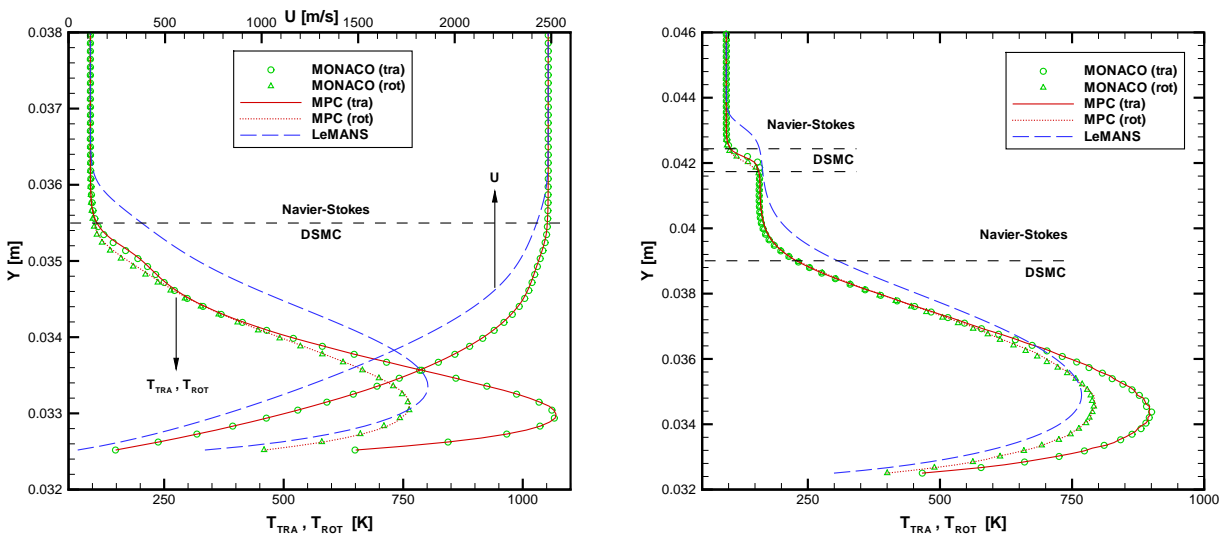
(a) Temperature and velocity profiles along C1.

(b) Temperature profiles along C2.
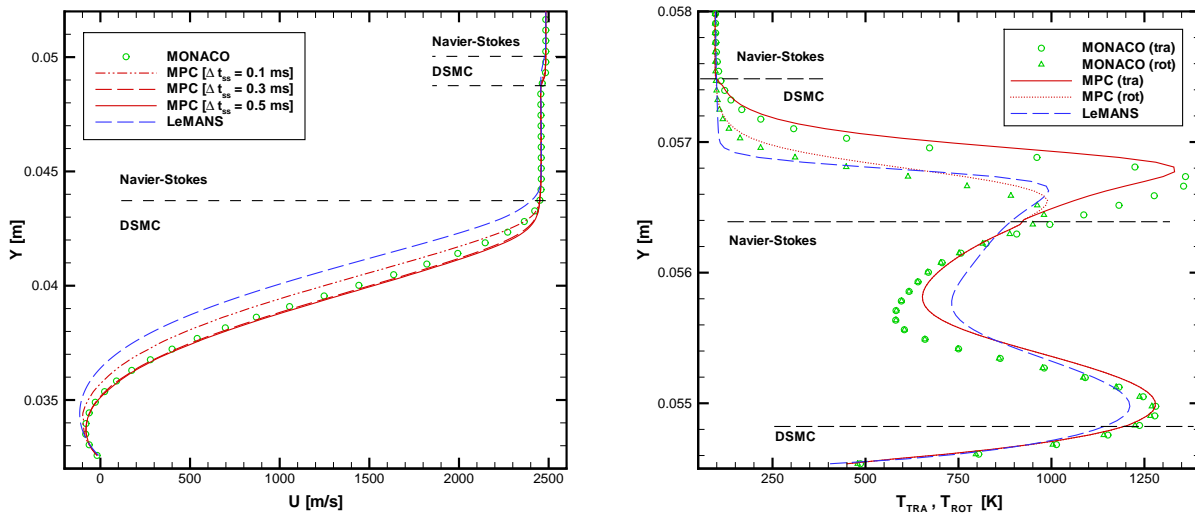
(c) Velocity profile along C3.

(d) Temperature profiles along C4.

Figure 7. Full DSMC, NS, and MPC flow field results at various locations along the hollow cylinder flare.

$C3$ is plotted for various values of Δt_{ss} and Fig. 7(c) 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. 2(b)) and is entirely attributed to an MPC simulation being initialized with particles corresponding to a full NS solution. Finally, Fig. 7(d) depicts the temperature profiles at location $C4$ which lies after re-attachment and 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 non-equilibrium inside the shock region. Both of these effects are captured well by the MPC method. In addition to using DSMC for the strong shock, the continuum breakdown parameter sets up a very

American Institute of Aeronautics and Astronautics

thin DSMC region inside the hypersonic boundary layer next to surface (only 6-10 DSMC cells thick). This causes problems with the heat transfer predicted by the MPC method and will be discussed in the next paragraph. Other than the shock, it is important to note that the remainder of the flow at location $C4$ is seen to be in thermal equilibrium and the MPC method is seen to move the initial NS solution into better agreement with the full DSMC solution. However, it should also be noted that at this location, full DSMC and full NS solutions agree quite well, especially close to the surface, indicating that at this point on the flare, the flow may lie in the continuum regime.

The coefficient of pressure (C_P) predicted by full DSMC, NS, and MPC simulations is plotted in Fig. 8(a). The size of the separation region has been successfully reduced by the MPC method and agrees well with that predicted by full DSMC. The coefficient of heat transfer, also referred to as the Stanton number [$S_t = q/(0.5)\rho_\infty u_\infty^3$] where q represents

(a) Coefficient of pressure.

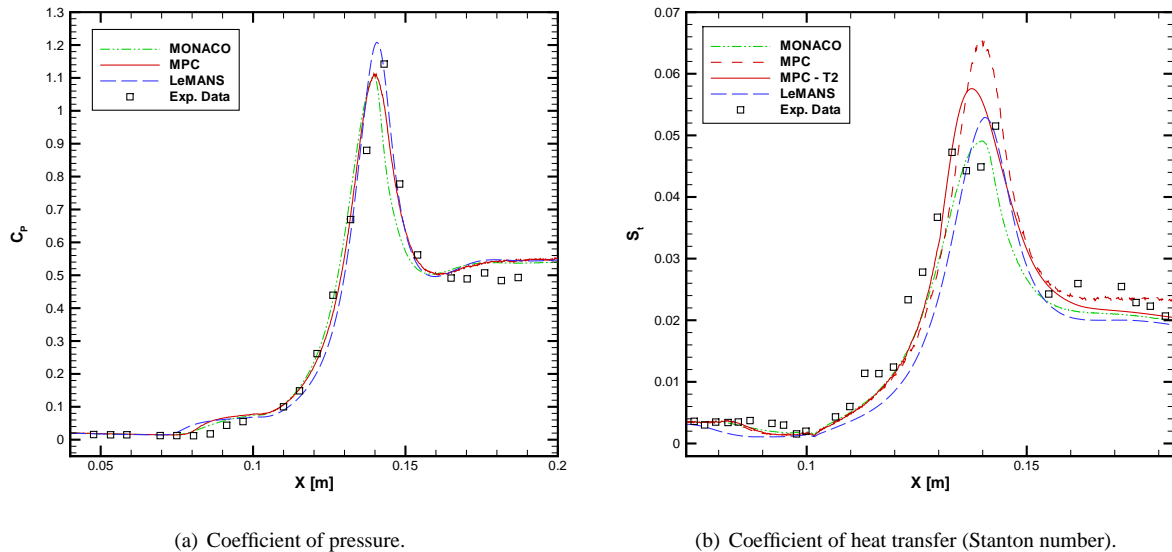(b) Coefficient of heat transfer (Stanton number).

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

the amount of energy transferred to the surface per unit time per unit area, is plotted in Fig. 8(b) for each simulation method. Although not shown, the MPC simulation reproduces exactly the results for heat transfer predicted by full DSMC near the leading edge. This is expected, since the MPC method was shown in Fig. 7(a) to accurately capture velocity slip, temperature jump, and thermal non-equilibrium near the leading edge. In Fig. 8(b), the MPC method is seen to improve the prediction of heat transfer in the separation region over the initial NS prediction. However, downstream of the flare junction the MPC results for heat transfer begin to differ from full DSMC results and over-predict the heat transfer by 20-30% for $X > 0.135\,m$. As mentioned above, and shown previously in Figs. 4 and 7(d) the DSMC region created by the MPC method next to the flare surface is very thin. A similar, thin DSMC region is found near the wedge surface for the hybrid simulations of Wu *et al.*[15] Although surface property results are not analyzed in the article by Wu *et al.*, close analysis of the thin DSMC region next to the wedge surface[15] shows similarities with Fig. 7(d) of this article. Specifically, small differences in the steep temperature gradient compared with the gradients predicted by either pure DSMC or NS are visible which may alter the calculated heat transfer significantly. In order to alleviate this problem, the Br_{cutoff} parameter could be lowered in order to enlarge the DSMC regions. However, such a change is found to adversely affect other regions of the flow. Instead, DSMC regions could be enlarged near the surface only. However this approach may be highly dependent on the geometry and flow conditions and would be difficult to implement in a general manner. A better approach comes from realizing that aside from the strong shock wave, the re-attached flow near the flare surface is well within the continuum regime and DSMC may not be necessary at all. Comparing the initial NS solution to the full DSMC solution in Fig. 7(d) reveals no significant difference between DSMC and NS solutions (except inside the shock wave). Additionally, in Fig. 2, NS simulation is seen to accurately predict the experimental data downstream of the flare junction. Attempting to use DSMC in a thin, continuum region, involving steep flow gradients next to the surface is very difficult and unnecessary. In order to address this, a second MPC simulation (MPC - T2) is run where DSMC regions are only allowed to develop prior to $X = 0.13\,m$ and the NS equations are solved for the remainder of the flare. As seen in Fig. 8(b), this improves the heat transfer result over that predicted by the original MPC simulation.

American Institute of Aeronautics and Astronautics

B. Planetary Probe

The final steady-state particle region for the planetary probe MPC simulation was shown previously in Fig. 5(a) and includes the shoulder region and a large portion of the wake. The heat transfer results for full DSMC, NS, and MPC simulations are displayed in Fig. 9(a) and compared with experimental measurements. Since the experiment was sting supported and the numerical simulations were performed with no sting, experimental results are only plotted for $s/R_n < 3.0$ where the presence of a sting has little impact on the heating rate. Comparing full DSMC and NS results in Fig. 9(a), it can be seen that both predict the same peak heating rate at the stagnation point. Along most of the fore-body (locations a-b), DSMC predicts a slightly lower heating rate than the NS equations, however both simulations predict heating rates ranging from 2-3 times larger than measured experimentally. While the reason for this remains unclear, both DSMC and NS results agree very well with simulations performed by other researchers.[25] Around the capsule shoulder and along the capsule base (locations b-e), DSMC is seen to predict a much lower heating rate than the NS equations, and DSMC is in better agreement with experimental results. In Fig. 9(a), the MPC simulation is shown to reproduce DSMC results very accurately. In the highly compressed fore-body region where DSMC and NS simulations produce similar results, the MPC method successfully uses the NS equations and therefore reproduces full NS results in this region. Just prior to the shoulder ($s/R_n = 1.6$), where the MPC method switches to DSMC, we see the heating rate transition from the NS result to the DSMC result. For the entire shoulder and base region, the MPC method is seen to reproduce full DSMC results with a high degree of accuracy.

The velocity magnitude around the surface of the planetary probe (just 10 μm off the surface) is extracted and plotted in Fig. 9(b) after being normalized by the free-stream velocity. The extraction is performed using Tecplot[C] where a curve is created parallel to the probe surface at a distance 10 μm normal to the surface. Data from the flow field is then extracted along this curve which involves Tecplot[C] linearly interpolating data from the mesh nodes to the curve. In the fore-body region both full DSMC and NS simulations predict virtually no slip velocity at the probe surface. However, DSMC shows very large velocity slip near the shoulder region which persists throughout the entire base region ($s/R_b > 0.8$), whereas the NS simulation assumes no slip conditions around the entire surface. It should be noted that since the velocity magnitude is extracted slightly off the surface that even the NS result will show a velocity magnitude greater than zero in some regions. Furthermore, the results in Fig. 9(b) agree both qualitatively and quantitatively with those of other researchers,[25] accounting for the fact that the exact distance from the surface directly influences the precise magnitude of the velocity. Finally, the MPC method is seen to accurately reproduce full DSMC results along the entire planetary probe surface while successfully solving the NS equations in the fore-body region where both DSMC and NS predict no velocity slip. Both MPC results plotted in Figs. 9(a) and 9(b) lend further confidence in the use of the gradient-length Knudsen number (Kn_{GL}) and a value of $Br_{\text{cutoff}} = 0.05$ for predicting continuum breakdown in blunt body flows.
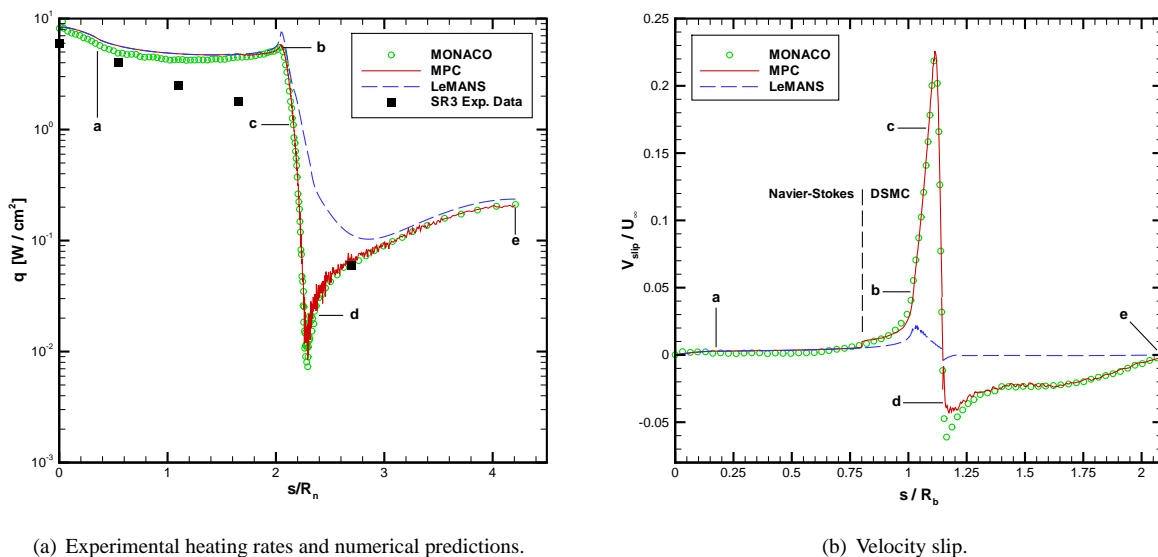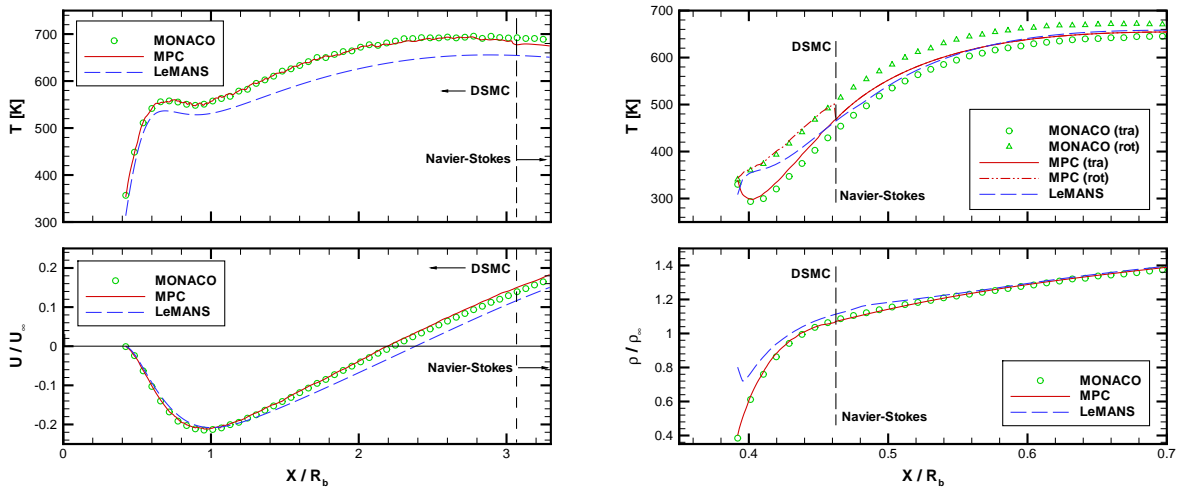

(a) Experimental heating rates and numerical predictions.

(b) Velocity slip.

Figure 9. Heating rates and velocity slip along the planetary probe surface.

The translational temperature and normalized x-velocity (U/U_∞) directly in the wake, along the axis of symmetry,

(a) Flow properties in the planetary probe wake.

(b) Flow properties near the planetary probe shoulder.

Figure 10. Flow field results for the planetary probe.

are plotted for each method in the top and bottom of Fig. 10(a) respectively. DSMC is seen to predict a 17% temperature jump at the base of the probe and $10-13\%$ higher temperatures in the wake compared with that predicted by the NS equations. In the bottom of Fig. 10(a), DSMC is seen to predict the vortex to close 10% earlier than the NS simulation. The MPC method captures these flow field features very accurately. In the process, the variations in DSMC regions provide more accurate boundary conditions and succeed in shifting the NS portion of the solution into better agreement with full DSMC results. The flow field properties extracted from the 45^o cut, normal to the shoulder surface shown previously in Fig. 5(a), are displayed in Fig. 10(b). Both translational and rotational temperatures are plotted in the top of Fig. 10(b). DSMC predicts the gas to be in thermal equilibrium at the wall at a slightly higher temperature (a temperature jump of 7%) compared with the NS solution. It should be noted that a full thermal accommodation boundary condition is used in DSMC which assumes that particles colliding with the surface leave with both translational and rotational energy modes thermalized to the wall temperature. As seen in the top of Fig. 10(b), close to the surface the MPC method uses DSMC which individually models translational and rotational energy modes, however on the other side of the interface the MPC method solves the NS equations which assume both modes are in equilibrium and are described by a single temperature. It is quite remarkable that despite particles being introduced to the DSMC region at the interface with incorrect translational and rotational energies, the hybrid-DSMC solution quickly recovers and reproduces full DSMC results accurately near the wall. This behavior was first observed while applying the MPC method to 1D shock waves[12] and provides a fundamental mechanism which drives an MPC solution towards the correct non-equilibrium result. Of course the MPC method is unable to predict the thermal non-equilibrium which persists in the expansion region behind the probe where the hybrid-NS region assumes a single temperature. From a practical perspective however, accurate prediction of surface properties as well as the flow near the surface and within the vortex are more important than predicting this small degree of thermal non-equilibrium away from the probe surface. The density along the 45^o cut is plotted in the bottom of Fig. 10(b). At the surface, the density predicted by DSMC is 50% of that predicted by the NS equations. Again, the MPC method is seen to accurately reproduce full DSMC results, provide more accurate boundary conditions for the hybrid-NS region, and shift the NS portion of the MPC solution into better agreement with full DSMC results.

C. Computational Efficiency

Precise determination of the speed-up factor achieved by the MPC method over full DSMC simulation becomes somewhat difficult for challenging problems such as the hollow cylinder flare and planetary probe. One factor that is difficult to quantify is the time required by a user to generate a mesh for a DSMC simulation versus a mesh for a NS simulation. In general, NS simulations are much faster than DSMC simulations for near-continuum flows. However significant time and effort is sometimes required by the user to generate NS meshes for complex geometries, whereas

American Institute of Aeronautics and Astronautics

with DSMC, mesh generation is typically performed automatically within the simulation. That being said, mesh generation for all DSMC and NS simulations presented in this article requires little time (less than two hours of user-time per mesh). 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 article are run in parallel on a cluster of CPUs 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 which is included approximately into the timing comparisons which 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 section III B (MPC-steady-state), plus the time required to sample and converge DSMC and NS regions (step 4). In order to remain consistent, each MPC solution is sampled for the same physical time and number of timesteps as the corresponding full DSMC solution. The performance parameters of most interest are listed in Table 3 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** 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 compared to 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 3. In order to show the relative cost of hybrid-NS operations, the ratio of time required for **initial NS** simulation compared with the full DSMC simulation as well as the ratio of the **timestep** used in NS regions compared with that used in DSMC regions are included in Table 3. 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. Since 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 3 for each case.

Table 3. Computational efficiency of the MPC method.

Hollow Cyl. Flare	speedup	memory	particles	initial NS	timestep	steady-state	sampling
DSMC	100%	100%	100%	N/A	1	100%	100%
MPC	71% (1.4x)	80%	53% (1.9x)	4%	28	20%	100%
MPC-T2	10% (10x)	15%	7.3% (13.7x)	4%	28	20%	100%
Planetary Probe	speedup	memory	particles	initial NS	timestep	steady-state	sampling
DSMC	100%	100%	100%	N/A	1	100%	100%
MPC	8% (12.5x)	20%	14% (7.1x)	.75%	150	30%	100%

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 only 1.4). This signifies that the time required for the initial NS solution and NS operations within the MPC cycle are not negligible compared with the time spent simulating DSMC regions. Since the flare region has a high density and is located further away from the axis (recall no cell-weighting is used for the hollow cylinder flare), a large portion of the particles in a full DSMC simulation are located in this flare region. As seen in Table 3, when the flare region is simulated entirely using the NS equations (case MPC - T2), a significant number of particles are eliminated within the MPC simulation which increases the speedup factor proportionally. Specifically a full DSMC simulation requires 13.7 times more particles than the MPC - T2 simulation, which obtains an accurate solution approximately 10 times faster. However, as mentioned previously in section II, NS simulation of the hollow cylinder flare problem (Run 11 conditions) requires a similar mesh resolution as DSMC. As seen in Table 3, the initial NS simulation takes 4% of the time required by a full DSMC simulation. If recent developments in the DSMC method are used in the simulation of the hollow cylinder

flare, it may be possible to reduce the number of simulated particles and increase the speed of full DSMC simulation by an order of magnitude.[24] The time required to obtain a full DSMC solution then becomes comparable to the time required to solve the NS equations and hybrid simulation is no longer practical or beneficial. As demonstrated in this article, the hollow cylinder flare configuration is certainly a challenging and valuable test case for a hybrid method. However, a more practical hybrid simulation may involve a full DSMC simulation of the shock-interaction (the entire hollow cylinder flare problem) which is itself embedded as part of a much larger continuum flow over an entire vehicle.

The planetary probe problem, on the other hand, is very well suited for a hybrid method. Due to the large variation in λ, mesh refinement within the MPC method is able to eliminate many computational cells and many particles. The MPC simulation uses 7.1 times fewer particles and implicit NS timesteps are 150 times larger than the time-step used by DSMC in the fore-body region. As a result, hybrid and NS operations within the MPC cycle are indeed negligible compared with the time spent simulating DSMC regions. The MPC solution is obtained approximately 12.5 times faster and uses only 20% of the memory compared with full DSMC simulation. For this case, full DSMC simulation requires a long time for particles to fill the wake region and reach steady-state. Since the MPC method begins with a NS solution, particles generated in the wake region are much closer to steady state immediately. As a result, the time required for the MPC method to reach steady-state and begin sampling is much less than the corresponding time required by full DSMC simulation. This results in a larger actual speedup (12.5x) than would be achieved by the elimination of particles alone (7.1x).

V.   Conclusions

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

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, temperature jump, thermal non-equilibrium, and surface properties predicted by full DSMC near the leading edge. MPC simulation also successfully reduces the over-sized separation region predicted by the NS equations to the size predicted by DSMC and experimental results. With the exception of the flare region, MPC simulation accurately reproduces DSMC results approximately 1.4 times faster than full DSMC simulation while using 80% of the memory. When the NS equations are used to simulate the entire flare region downstream of re-attachment, the resulting MPC simulation achieves the same level of accuracy approximately 10 times faster using only 15% of the memory.

3) The gradient-length Knudsen number predicts continuum breakdown very well for both blunt body and shock-interaction flows with the exception of the dense flare region of the hollow cylinder flare problem. The parameter has difficulty in very dense regions with large gradients where it predicts continuum breakdown near the surface even though the NS equations are shown to be valid in this region. Furthermore, the transfer of information across such thin DSMC regions is is shown to lead to a 20-30% over-prediction in heat transfer along the flare. It is recommended that such thin DSMC regions be avoided in a loosely-coupled hybrid approach which may require further investigation into the failure of the continuum breakdown parameter in this region.

4) The fine resolution required to solve the NS equations for the hollow cylinder flare problem combined with the possibility of efficient full DSMC simulation (utilizing recent advances in the DSMC method) lead to the conclusion that the hollow cylinder flare problem is not practically suited for a hybrid simulation. Rather, the entire shock-interaction flow is better simulated using pure DSMC and embedded as a single particle region (perhaps around a control surface) in a hybrid simulation of an entire vehicle.

5) The planetary probe problem involves a variation in mean-free-path of 2 orders of magnitude between stagnation and wake regions. The fine spatial and temporal resolution required by a full DSMC simulation in the fore-body makes the planetary probe problem a very practical candidate for a hybrid simulation. The MPC method is shown to reproduce full DSMC flow field and surface property results with a high level of accuracy using 7.1 times fewer particles and obtaining the solution approximately 12.5 times faster than full DSMC while requiring only 20% of the memory.

Application of the modular particle-continuum (MPC) method to the hollow cylinder flare and planetary probe benchmark problems clearly show that a loosely coupled hybrid DSMC-NS approach is very promising for the solution of hypersonic steady-state flows where large variations in local Knudsen number are seen. The MPC method is shown

to be particulary promising for blunt-body and leading edge flows. Complex shock interaction flows are recommended for hybrid simulation only if the flow is part of a much larger, full vehicle simulation.

Acknowledgments

References

[1]Bird, G. A., *Molecular Gas Dynamics and the Direct Simulation of Gas Flows*, Oxford University Press, New York, 1994.

[2]Wilmoth, R. G., Mitcheltree, R. A., Moss, J. N., and Dogra, V. K., "Zonally Decoupled Direct Simulation Monte Carlo Solutions of Hypersonic Blunt-Body Wake Flows," *Journal of Spacecraft and Rockets*, Vol. 31, No. 6, 1994, pp. 971–979.

[3]Glass, C. E. and Gnoffo, P. A., "A 3-D Coupled CFD-DSMC Solution Method With Application to the Mars Sample Return Orbiter," *NASA TM-2000-210322*, July 2000.

[4]Glass, C. E. and Horvath, T. J., "Comparison of a 3-D CFD-DSMC Solution Methodology With a Wind Tunnel Experiment," *NASA TM-2002-211777*, Aug. 2002.

[5]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.

[6]Hash, D. B. and Hassan, H. A., "A Decoupled DSMC/Navier-Stokes Analysis of a Transitional Flow Experiment," *AIAA Paper 96-0353*, Jan. 1996, presented at the 34th AIAA Aerospace Sciences Meeting and Exhibit, Reno, NV.

[7]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.

[8]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.

[9]Wang, W. L. and Boyd, I. D., "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.

[10]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, 2004, pp. 768–777.

[11]Schwartzentruber, T. E., Scalabrin, L., and Boyd, I. D., "A modular particle-continuum numerical method for hypersonic non-equilibrium gas flows," *Journal of Computational Physics*, 2007, In Press, Available online Feb. 03, 2007.

[12]Schwartzentruber, T. E. and Boyd, I. D., "A hybrid particle-continuum method applied to shock waves," *Journal of Computational Physics*, Vol. 215, 2006, pp. 402–416.

[13]Schwartzentruber, T. E., Scalabrin, L. C., and Boyd, I. D., "Hybrid Particle-Continuum Simulations of Non-Equilibrium Hypersonic Blunt Body Flow Fields," *Journal of Thermophysics and Heat Transfer*, In Press, May 2007.

[14]Schwartzentruber, T. E., Scalabrin, L. C., and Boyd, I. D., "Modular Implementation of a Hybrid DSMC-NS Solver for Hypersonic Non-Equilibrium Flows," *AIAA Paper 07-613*, Jan. 2007, presented at the 45th AIAA Aerospace Sciences Meeting and Exhibit, Reno, NV.

[15]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, 2006, pp. 579–607.

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

[17]Boyd, I. D., "Analysis of Rotational Nonequilibrium in Standing Shock Waves of Nitrogen," *AIAA Journal*, Vol. 28, No. 11, 1990, pp. 1997–1999.

[18]Scalabrin, L. C. and Boyd, I. D., "Development of an Unstructured Navier-Stokes Solver for Hypersonic Nonequilibrium Aerothermodynamics," *AIAA Paper 05-5203*, 2005, presented at the 38th AIAA Thermophysics Conference, Toronto, Ontario, Canada.

[19]MacCormack, R. W. and Candler, G. V., "The Solution of the Navier-Stokes Equations Using Gauss-Seidel Line Relaxation," *Computers and Fluids*, Vol. 17, 1989, pp. 135–150.

[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 Paper 01-1031*, Jan. 2001, presented at the AIAA Aerospace Sciences Meeting and Exhibit, Reno, NV.

[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 Paper 01-1031*, Jan. 2001, presented at the AIAA Aerospace Sciences Meeting and Exhibit, Reno, NV.

[22]Candler, G. V., Nompelis, I., and Druguet, M. C., "Navier-Stokes Predictions of Hypersonic Double-Cone and Cylinder-Flare Flow Fields," *AIAA Paper 01-1024*, Jan. 2001, presented at the 39th AIAA Aerospace Sciences Meeting and Exhibit, Reno, NV.

[23]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," *AIAA Paper 02-0581*, Jan. 2002, presented at the 40th AIAA Aerospace Sciences Meeting and Exhibit, Reno, NV.

[24]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.

[25]Moss, J. N. and Price, J. M., "Survey of Blunt Body Flows Including Wakes at Hypersonic Low-Density Conditions," *Journal of Thermophysics and Heat Transfer*, Vol. 11, No. 3, 1997, pp. 321–329.

[26]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.

[27]Garcia, A. L. and Alder, B. J., "Generation of the Chapman-Enskog Distribution," *Journal of Computational Physics*, Vol. 140, 1998, pp. 66–70.

[28]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.

American Institute of Aeronautics and Astronautics