Multiscale Particle-Continuum Simulations of Hypersonic Flow over a Planetary Probe

Thomas E. Schwartzentruber[*]
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

A modular particle-continuum numerical method is used to simulate the flow over a 70 deg blunted cone planetary probe geometry under various steady-state hypersonic conditions. The conditions studied correspond to low global Knudsen number flow where hypersonic velocities induce a large range of temporal and spatial scales that must be modeled. The modular particle-continuum algorithm loosely couples direct simulation Monte Carlo and Navier–Stokes methods, which operate in nonequilibrium and continuum regions, respectively. In addition, particle and continuum regions have different mesh densities and are updated using different sized time steps. Modular particle-continuum simulations are shown to reproduce the heating rates, velocity slip, and thermal nonequilibrium on the probe surface, as well as the flowfield properties predicted by the direct simulation Monte Carlo method with high accuracy. For the first time, such a hybrid particle-continuum method is shown to achieve high accuracy while achieving an order of magnitude decrease in required computational time and memory compared with pure particle simulation. The modular particle-continuum simulations agree well with Navier–Stokes simulations and experimental measurements in the dense forebody flow. In the rarefied wake of the planetary probe, the modular particle-continuum results are in better agreement with experimental measurements than Navier–Stokes simulations.

Nomenclature

Cp	=	coefficient of pressure, $(p - p_\infty)/(0.5\rho_\infty U_\infty^2)$
D	=	probe diameter, m
d	=	diameter of molecule, m
F	=	mesh scaling factor
Kn	=	Knudsen number, λ/l_c
Kn_{GL}	=	gradient-length Knudsen number, where l_c refers to gradient length, $l_c = (Q/|\nabla Q|)$
Kn_g	=	global Knudsen number, where l_c refers to a global vehicle dimension
k	=	Boltzmann constant, 1.38×10^{-23}, J/K
l_c	=	characteristic length, m
l_f	=	length of cell face f, m
M	=	Mach number
m	=	mass of molecule, kg
$\hat{n}_f$	=	unit normal vector of cell face f
p	=	pressure, N/m^2
Q	=	macroscopic property
q	=	heat transfer per unit area, W/cm^2
R	=	probe geometry radius, m
r_f	=	mesh refinement factor for cell face f
s	=	distance along surface beginning at stagnation point, m
T	=	continuum combined translational-rotational temperature, K
T_{ROT}	=	rotational temperature, K
T_{TRA}	=	translational temperature, K
Tw	=	wall temperature, K
U	=	bulk X velocity of gas, m/s
V	=	bulk Y velocity of gas, m/s
$|V|$	=	velocity magnitude, $\sqrt{U^2 + V^2}$, m/s
V_{slip}	=	velocity magnitude $|V|$ at the surface, m/s
X, Y	=	Cartesian coordinates, m
λ	=	mean-free path, m
μ	=	viscosity, (Ns)/m^2
ρ	=	density, kg/m^3
τ_c	=	mean-free time, s
ω	=	power law temperature dependence of μ

Subscripts

b	=	refers to base of planetary probe
c	=	refers to shoulder of planetary probe
n	=	refers to blunted nose of planetary probe
ref	=	reference value of the parameter
s	=	refers to sting of planetary probe
∞, o	=	freestream value of the parameter

[*]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 vehicles entering a planetary atmosphere generate multiscale 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 molecular 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 nonequilibrium flows. The DSMC method 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 collision time τ_c, respectively.

Generally, the global Knudsen number Kn_g, which uses the freestream value of λ and an overall dimension of the vehicle for l_c, provides a good indication of whether the flow is continuum or nonequilibrium. However, difficulty arises for low Knudsen number flows ($Kn_g \leq 0.001$) where hypersonic velocities induce localized regions of nonequilibrium imbedded within a mostly continuum flow. For example, the DSMC method may be required to accurately model the flow inside shock waves, boundary layers, and wake regions, while remaining computationally expensive in the dense forebody flow. For blunt-body flows, 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 forebody 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 a full DSMC simulation, this approach requires a priori knowledge of a suitable particle-continuum interface and therefore lacks generality. Various researchers have proposed hybrid numerical methods which adaptively reposition 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 [11].

In this article, a modular particle-continuum (MPC) numerical algorithm [11] is used to simulate low global Knudsen number hypersonic flow over a planetary probe geometry. Such a flow involves large variations in both λ and τ_c, which enables a hybrid particle-continuum method to be far more computationally efficient than full particle simulation. The MPC algorithm has been shown to reproduce full DSMC simulation results with a high degree of accuracy for 1-D normal shock waves [12], as well as for 2-D flow over a cylinder [13]. The MPC method loosely couples DSMC and NS solution techniques, which operate on different mesh densities with different time steps in nonequilibrium and continuum regions of the flow, respectively. Such an approach enables both spatial and time-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. The purpose of this article is to outline how existing DSMC and NS codes can be combined using the MPC algorithm and to demonstrate that such a loosely coupled approach is able to produce highly accurate results using significantly less computational resources than full DSMC simulation. Researchers studying hypersonic flows involving variations in λ and τ_c comparable to, or greater than, the variations studied in this article can expect similar computational savings by using a modular particle-continuum algorithm. In Sec. II, the setup of a hybrid simulation will be outlined, including the mesh refinement process between particle and continuum regions. The MPC numerical cycle will also be outlined in Sec. II by tracking the hybrid solution at various times during an MPC simulation. Numerical simulation results are presented in Sec. III for three experimental

configurations involving hypersonic flow over a planetary probe geometry. MPC results are compared with full NS and full DSMC simulation results, as well as with experimental data. Section IV details the computational efficiency of the MPC method where, for the first time, a hybrid particle-continuum algorithm is shown to reduce the computational time and memory by an order of magnitude compared with full DSMC simulation with a high level of accuracy. Finally, the conclusions drawn from these simulations are presented in Sec. V.

II.   Hybrid Simulation Setup

A.   Geometry and Flow Conditions

The planetary probe is a 70 deg blunted cone geometry detailed in Fig. 1, where s is the distance around the surface of the probe beginning at the stagnation point. Experimental data were obtained in the SR3 wind tunnel in Meudon, France for a 5-cm-diam probe. Three test conditions were considered, where the freestream was nitrogen at a nominal Mach number of 20 with freestream densities of 1.7×10^{-5}, 5.19×10^{-5}, and 46.6×10^{-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 [15]. Only the medium-density (case 2) and high-density (case 3) experimental conditions are analyzed in this study. The low-density ($Kn_g = 0.03$) SR3 experiment is found to lie entirely in the rarefied (nonequilibrium) regime and is not a candidate for hybrid particle-continuum simulation. The precise experimental conditions are summarized in Table 1. A zonally decoupled approach [2] has been used to simulate the same SR3 cases, where the dense forebody 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 experiments supported the model with a sting, the zonally decoupled simulations did not include the sting. To investigate the size of the wake vortex and allow for direct comparison with the zonally decoupled results, the present study does not include the sting for the high-density run (case 3). However, the present study does include the sting for the medium-density run (case 2) to compare with experimental data. Experiments using the same geometry as detailed in Fig. 1, but with a larger base diameter ($D_b = 15.24$ cm) and different flow conditions, were performed in the Large Energy National Shock (LENS) facility operated by the Calspan–University of Buffalo Research Center. The LENS experiment was also conducted with nitrogen, and conditions for this experiment are summarized in Table 1. Modular particle-continuum simulations of the three experiments listed in Table 1 will be analyzed in this article.

B.   Initialization of a Hybrid Simulation

The MPC method begins with a pure NS solution obtained on a quadrilateral mesh designed for the continuum equations. Although research into triangular (or 3-D tetrahedral) meshes for hypersonic flows is ongoing, quadrilateral (or 3-D hexahedral) meshes are currently required to properly resolve the large gradients found in hypersonic shock waves and boundary layers [16,17]. The NS equations are solved using the LeMANS code [18]. For the results of this article, it is assumed that rotational and translational energy
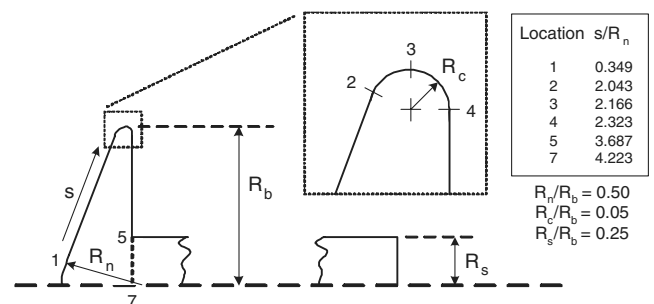

Fig. 1   Planetary probe geometry.

Table 1 Experimental flow conditions for the planetary probe configurations

Case	M_∞	U_∞, m/s	ρ_∞, kg/m³	T_∞, K	λ_∞, mm	Tw, K	D_b, cm	Kn_g[a]
SR3 facility case 2	20.0	1502.0	5.19×10^{-5}	13.6	0.54	300.0	5.0	0.01
SR3 facility case 3	20.5	1633.0	46.60×10^{-5}	15.3	0.06	300.0	5.0	0.001
LENS facility	15.6	3246.0	13.06×10^{-5}	103.7	0.35	294.0	15.24	0.002

[a]Kn_g is based on the base diameter of the planetary probe D_b and λ_∞

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, Navier–Stokes equations. The viscosity in the NS solver is modeled by the following relation, which is consistent with the variable-hard-sphere (VHS) viscosity model used in the DSMC method:

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

All numerical results presented in this article are for diatomic nitrogen with a reference diameter of $d_{\text{ref}} = 4.17 \times 10^{-10}$ m at a reference temperature of $T_{\text{ref}} = 273$ 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. LeMANS solves the NS 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 the Steger–Warming method 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 LeMANS numerical method are contained in [18].

Using this initial NS solution, a measure of continuum breakdown is calculated in each cell. The parameter used by the MPC method is the gradient-length Knudsen number [20] given by

$$Kn_{\text{GL}-Q} = \frac{\lambda}{Q}|\nabla Q| \quad (2)$$

where the maximum is taken over the flow quantities Q of interest: density ρ, temperature T, and velocity magnitude $|V|$. In addition, because this parameter does not account for thermal nonequilibrium, a condition of near thermal equilibrium is added to the definition of the continuum breakdown parameter, which then becomes

$$Kn_{\text{GL}} = \max\left(Kn_{\text{GL}-\rho}, Kn_{\text{GL}-T}, Kn_{\text{GL}-|V|}, 5 \times \frac{T_{\text{TRA}} - T_{\text{ROT}}}{T_{\text{ROT}}}\right) \quad (3)$$

Thus, the final value of the breakdown parameter accounts for compression thermal nonequilibrium ($T_{\text{ROT}} < T_{\text{TRA}}$) but ignores expansion thermal nonequilibrium ($T_{\text{ROT}} > T_{\text{TRA}}$), because the difference in Eq. (3) is then negative. The local mean-free path is calculated in each cell of the continuum solution as

$$\lambda = \frac{2(5-2\omega)(7-2\omega)}{15}\sqrt{\frac{m}{2\pi kT}}\left(\frac{\mu}{\rho}\right) \quad (4)$$

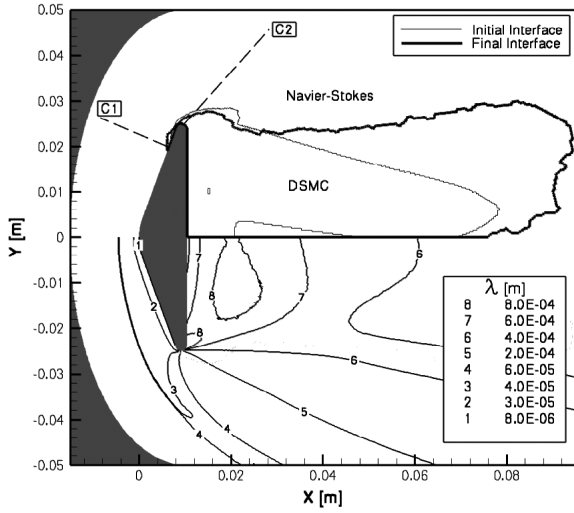where ρ is the local density and μ is the local viscosity. This expression for λ is consistent with the VHS viscosity model shown in Eq. (1), and all variables have the same values as previously outlined. In addition to its use in determining continuum breakdown, λ is also a measure of how much refinement is necessary in the DSMC regions. This cell-centered value of λ is then compared with the length of each cell face to determine the refinement required in each NS cell. For an accurate DSMC simulation, the refined cell size should be less than

the local value of λ. This restriction is most important in the direction of the flow gradient, however, it is less important in the direction normal to this gradient. The mesh refinement procedure used in this article refines the cell size to λ in the direction of the gradient and relaxes the cell size in the direction normal to the flow gradient to 4λ. This is accomplished by setting the refinement for each cell face as
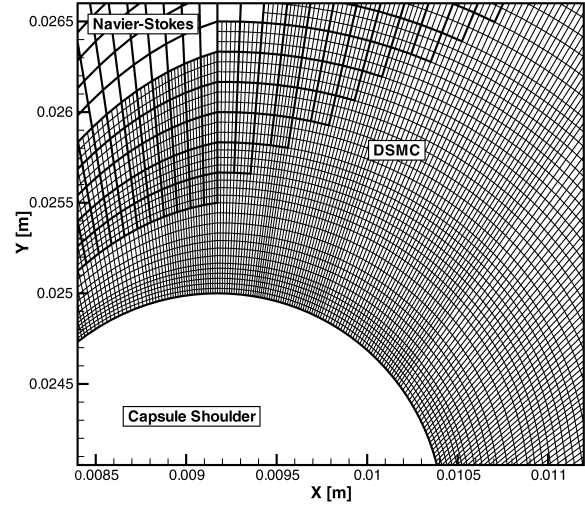
$$r_f = l_f \times \left[F\lambda\left(3\left|\hat{n}_f \cdot \frac{\nabla Q}{|\nabla Q|}\right| + 1\right)\right]^{-1} \quad (5)$$

Here, l_f is the length of the cell face which is compared to a multiple of the mean-free path λ. The dot product between the face unit normal vector $\hat{n}_f$ and the unit gradient vector determines the alignment of the cell face with the direction of the gradient. If the cell face is aligned parallel to the flow gradient (which means $\hat{n}_f$ is *normal* to the flow gradient), the dot product in Eq. (5) is zero, and the refinement is scaled to $F\lambda$. Whereas, if the cell face is perpendicular to the flow gradient (meaning $\hat{n}_f$ is *parallel* to the gradient vector), the refinement is relaxed to $4F\lambda$. Because the refinement is required because of continuum breakdown, for consistency, the gradient in Eq. (5) is calculated using the same flow quantity (ρ, T, or $|V|$) which resulted in the maximum value of Kn_{GL} in that cell [Eq. (2)]. Finally, the factor F in Eq. (5) allows scaling of the entire mesh by any fraction (or multiple) of the local mean-free path. The entire expression is rounded to the nearest integer to obtain r_f. Because opposite faces in a quadrilateral cell must be refined equally, each NS cell will contain two integers (r_1 and r_2) specifying the minimum required refinement in orthogonal directions. The current DSMC module also requires that each cell face border only one neighboring cell. As a result, the refinement cannot be set arbitrarily in each cell but must be consistent with neighboring cells. A simple algorithm is employed which ensures all cells are refined consistently while still meeting the minimum required refinement. This restriction may produce more cells than required in certain regions. More advanced mesh refinement procedures would result in fewer computational cells and increased efficiency of the MPC method. All cells in which $Kn_{\text{GL}} > 0.05$ are labeled as DSMC cells. This cutoff value of 0.05 for continuum breakdown has been recommended [20,21] and previously validated [12,13] within the hybrid code for similar flow conditions. These DSMC regions are then extended a few cells further into the continuum regions to create an overlap region. As an example, the variation in λ and the initial DSMC–NS interface are shown for the SR3 case 3 simulation in Fig. 2a. The initial nonequilibrium region captured by the Kn_{GL} parameter includes the rapid expansion around the shoulder of the capsule and a portion of the rarefied wake. It should be noted that, although the interior of the bow shock wave is also a highly nonequilibrium region, previous research [13] has shown that particle simulation of the bow shock interior is not required for accurate prediction of surface properties.

Once all cell types including the overlap regions have been defined using the continuum breakdown parameter [Eq. (3)], the refinement is carried out. For each cell labeled as nonequilibrium, $r_1 \times r_2$ refined cells are created and added to the DSMC mesh data structure. This creates a separate mesh refined to the mean-free path to be used by the DSMC module. A portion of the initial mesh refinement, including the overlap region in the vicinity of the capsule shoulder, is shown in Fig. 2b for the SR3 case 3 simulation. To start the hybrid simulation, particles are generated inside all DSMC cells with Chapman–Enskog velocity distributions [22]. The Chapman–Enskog velocity distribution is a first-order perturbed Maxwellian velocity

a) Interface locations and variation of λ

b) Mesh refinement near the capsule shoulder

Fig. 2   Interface locations and mesh refinement for the SR3 case 3 simulation.

distribution, which is an inherent assumption made by the continuum NS equations [1]. Thus, particles are initially generated to be consistent with the NS solution and are given near-equilibrium velocity distributions within each DSMC cell. During an MPC simulation, particle regions are simulated using MONACO [23], a general, cell-based implementation of the DSMC method. For the simulations presented in this article, MONACO employs the VHS collision model and the variable rotational energy exchange probability model of Boyd [24], where the reference temperature for rotational energy exchange is specified as 91.5 K and the maximum rotational collision number as 18.1. Energy transfer to vibrational modes is not considered. In addition, the flow is axisymmetric 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. In addition, up to 16 collision subcells are used as needed inside DSMC cells to ensure that collision cells are less than λ in dimension. Finally, diffuse reflection and full thermal accommodation is assumed for particles colliding with the probe surface. During DSMC simulation, the solution in nonequilibrium regions transitions from the near-equilibrium solution (consistent with the initial NS solution) toward the correct nonequilibrium solution. As this transition occurs, information is periodically transferred to the continuum NS regions, and DSMC and NS regions are loosely coupled for the remainder of the simulation. In Fig. 2a, the mean-free path λ is seen to vary by over 2 orders of magnitude from the stagnation point to the wake region. Thus, by modeling the dense forebody region using the NS equations, a large number of cells and DSMC particles are eliminated. Particle-continuum interfaces adapt as the hybrid simulation proceeds. For future reference, the *final* interface location for the SR3 case 3 simulation is also shown in Fig. 2a.

C.   Information Transfer and Coupling Procedure

DSMC and NS regions are loosely coupled using state-based information transfer depicted in Fig. 3. The overlap region allows the solution in DSMC regions to differ from the continuum NS solution. Information is then transferred to continuum regions by averaging over particle information [25] and updating the macroscopic boundary conditions in NS boundary cells. On the other side of the interface, particles are generated [22] in DSMC boundary cells using the macroscopic state and gradients of the corresponding NS cell. The loosely coupled numerical cycle employed in the MPC method is completely detailed in [11] and 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) Iterate DSMC regions and allow the solution to progress with the boundary conditions provided by the current NS solution, while adaptively repositioning the interfaces (using the overlap region) without transferring any information to NS regions.

3) After the DSMC solution and interfaces stop changing, use the current DSMC solution to set new NS boundary conditions. Significantly converge the NS region.
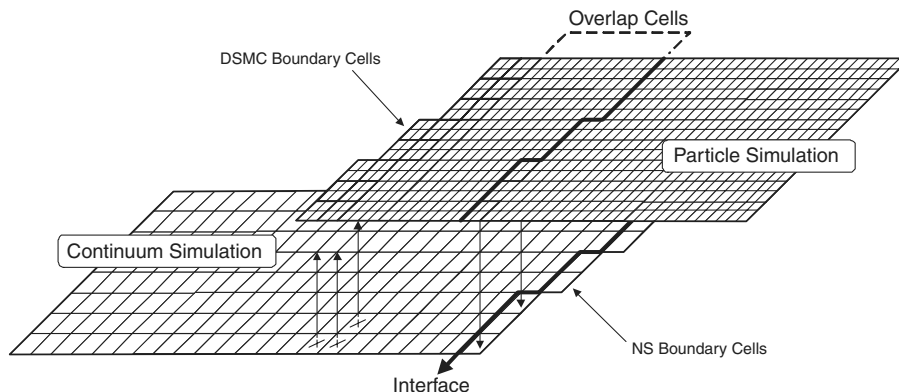

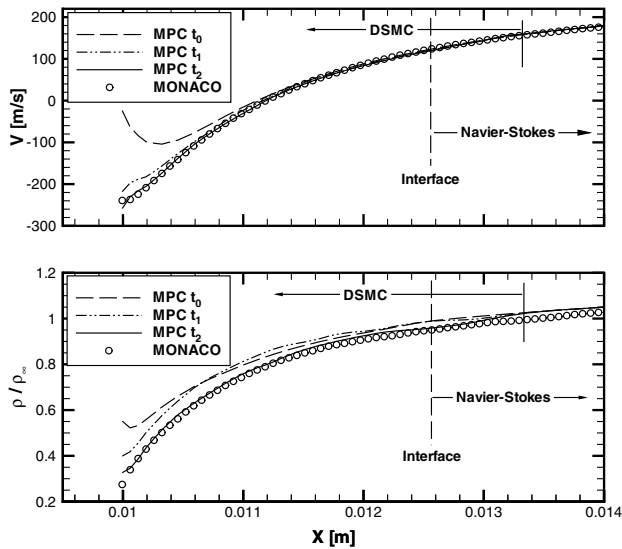Fig. 3   Information transfer and mesh refinement.

Fig. 4 Progression of the solution in DSMC regions within the MPC cycle.

Fig. 5 Interface locations and variation of λ for the SR3 case 2 simulation.

IF the new NS solution changes the state in any DSMC boundary cell $\rightarrow$ **return** to 2.

ELSE, if the new NS solution does not change the state in any DSMC boundary **cell** $\rightarrow$ **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.

Figure 4 shows MPC solutions at various stages (time t_0, t_1, and t_2) during step 2 of the preceding MPC cycle for the SR3 case 3 simulation. The progression of the solution in DSMC regions is seen along the 45 deg cut through the shoulder of the capsule (line C2 shown previously in Fig. 2a) and is plotted in Fig. 4. Figure 4 demonstrates how the overlap region provides a buffer that 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. 4 where, although the DSMC density at the edge of the overlap region is set by the NS solution (MPC t_0), 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 preceding MPC cycle will repeat itself until step 4 is reached.

III. Hybrid Simulation Results

A. SR3 Case 2 Experiment

The freestream conditions for the SR3 case 2 simulation result in a global Knudsen number $Kn_g = 0.01$. As a result, the flowfield is significantly rarefied and a large portion the hybrid simulation uses the DSMC method. To increase the size of the continuum regions, the cutoff value of the continuum breakdown parameter is raised from the recommended value of 0.05 to a value of 0.10. Figure 5 shows the variation in λ for this simulation, as well as the initial and final DSMC–NS interface locations. Initially, the continuum breakdown parameter (with a cutoff value of $Kn_{GL} = 0.10$) sets up DSMC regions inside the shock and for the majority of the wake flow, since λ is very large in the wake. The final interface location for this simulation remains similar in the wake, however, the particle region in the forebody expands significantly upstream to enclose the entire diffuse shock wave. Despite the rarefied conditions for this simulation, a region of continuum flow exists downstream of the
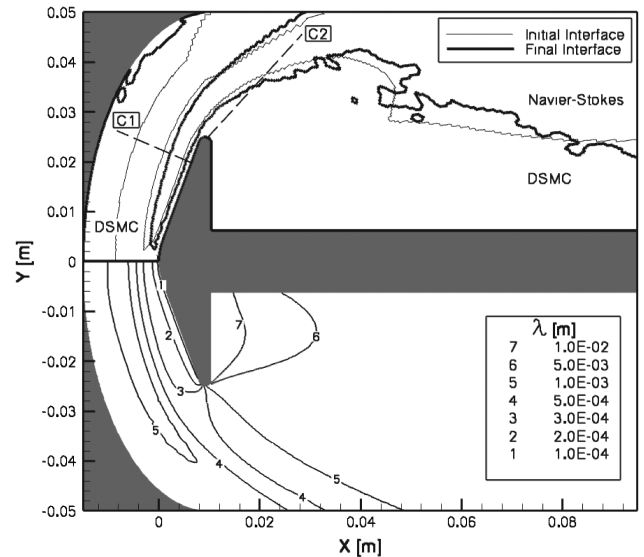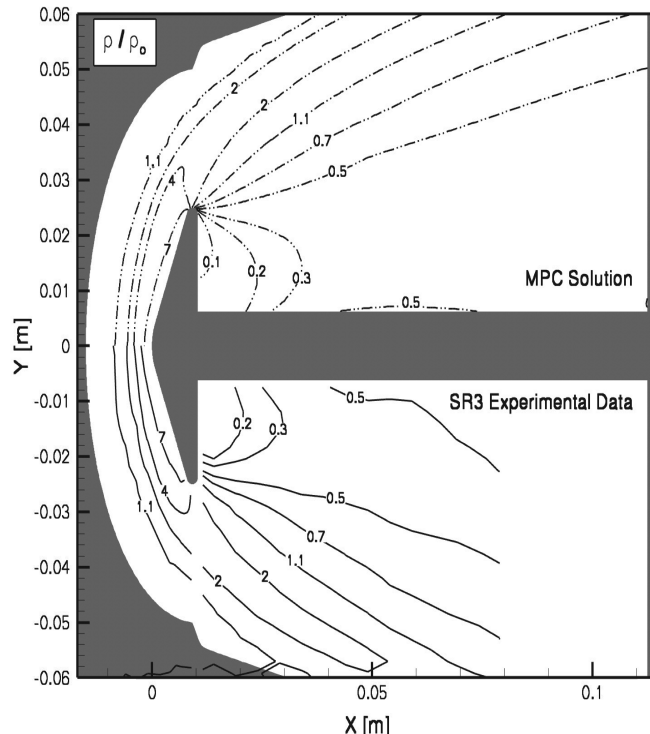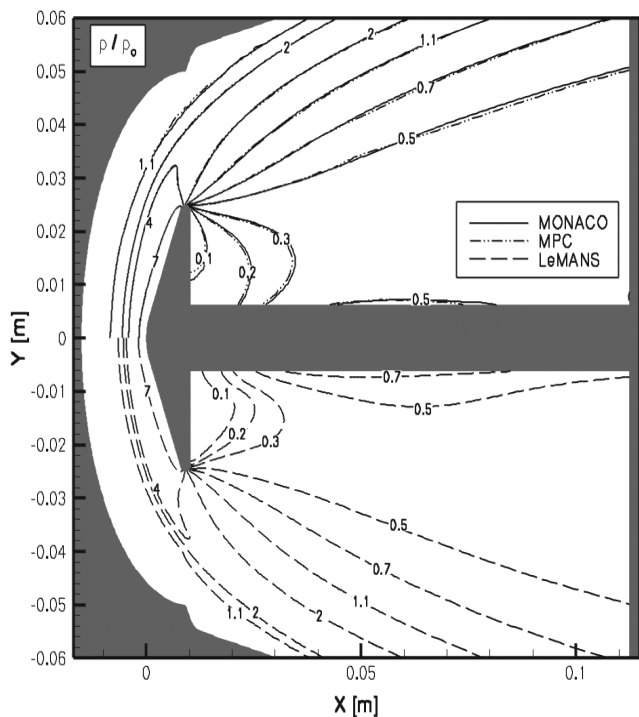
shock and outside of the wake flow where the flow is modeled with the NS equations.

The resulting density contours around the planetary probe predicted by full DSMC simulation (MONACO), full NS simulation (LeMANS), and modular particle–continuum simulation (MPC) are shown in Fig. 6a. The MPC method is seen to reproduce the full DSMC solution with a high degree of accuracy. Compared with full NS simulation, the bow shock is seen to be thicker and the density along the sting is seen to be lower. The MPC solution is compared directly with experimental measurement of the flowfield density in Fig. 6b. The density distribution in the flowfield was experimentally obtained by the electron beam fluorescence technique. The experimental density contours plotted in Fig. 6b represent the ratio of measured density to density at the same location in the absence of the planetary probe model. MPC and full DSMC simulations are seen to predict the experimentally measured shock thickness, whereas the NS simulation predicts a bow shock that is too thin. In addition, MPC and full DSMC simulations are seen to better predict the low densities measured along the sting.
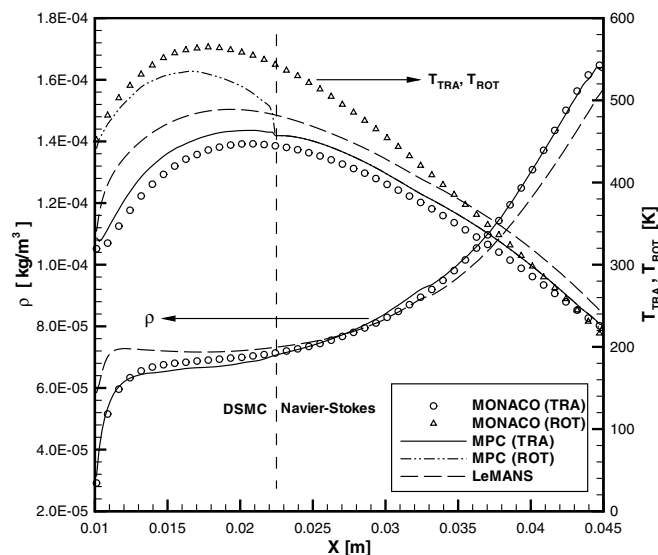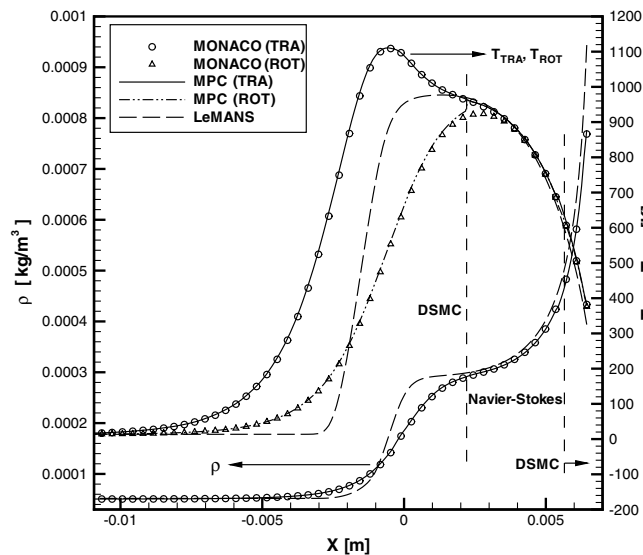
Figures 7a and 7b show profiles of both translational and rotational temperature, as well as density profiles in the forebody and wake regions, respectively. In these figures, the full DSMC (MONACO) simulation models both translational and rotational temperatures separately, whereas the full NS (LeMANS) temperature profile assumes a single translational–rotational temperature. The MPC solution is capable of modeling both translational and rotational temperatures in DSMC particle regions, but assumes a single translational–rotational temperature in continuum NS regions. Simulation data extracted along line C1 (from Fig. 5) are plotted in Fig. 7a. Here, the MPC solution has transitioned from the initial thin NS shock profile into very good agreement with the diffuse shock wave predicted by the full DSMC method. For this flow, the DSMC–NS interface has actually moved upstream so that the DSMC particle region includes the inflow boundary. The MPC method is seen to reproduce both translational and rotational temperatures inside the diffuse shock wave. In the postshock region, where the gas comes back into thermal equilibrium, the MPC method solves the NS equations. At the postshock DSMC–NS interface, Fig. 7a shows that the gas has not quite reached thermal equilibrium, and the translational temperature from the DSMC region is used as the boundary condition for the NS region of the MPC simulation. Next to the probe surface, the MPC simulation switches back to DSMC and is able to capture the same temperature jump predicted by the full DSMC method. Of all flow variables, the discrepancy between MPC and full DSMC solutions is typically largest for temperature due to thermal nonequilibrium. As seen in Fig. 7a, the agreement between

a) Numerical simulation comparisons

b) MPC comparison with experimental result

Fig. 6   Computed and experimentally measured density fields for the SR3 case 2 simulation.

a) Temperature and density along C1

b) Temperature and density along C2

Fig. 7   Flowfield results for the SR3 case 2 simulation.

MPC and full DSMC density profiles is excellent. Figure 7b compares solutions in the rapid expansion region near the capsule shoulder (cut C2 from Fig. 5). As previously mentioned, the continuum breakdown parameter detects compression thermal nonequilibrium ($T_{ROT} < T_{TRA}$) but ignores expansion thermal nonequilibrium ($T_{ROT} > T_{TRA}$). Thus, in the wake, particle and continuum regions are based on the Kn_{GL} parameter and not on the condition of thermal equilibrium. As seen in Fig. 7b, despite the DSMC–NS interface lying close to the probe surface, the MPC method is able to accurately reproduce both translational and rotational temperatures near the surface, despite boundary conditions from the NS regions that are assumed to be in thermal equilibrium. The density predicted by the MPC method near the surface and in the wake is again in excellent agreement with the full DSMC result. It

should be noted that the expansion thermal nonequilibrium that persists in the wake flow does not affect the surface properties, and thus modeling this region with the NS equations is acceptable even if thermal nonequilibrium is present. This is the reason that the continuum breakdown parameter [Eq. (3)] includes regions of compression thermal nonequilibrium only. In a recent study, a separate rotational temperature is added to the NS model in the LeMANS code [26]. Future research should investigate information transfer within the MPC algorithm associated with rotational energy and implement separate rotational temperatures into the MPC method. For example, such research has recently been carried out for the inclusion of vibrational energy in the MPC algorithm [27].

The heating rates computed on the planetary probe surface by full DSMC, full NS, and MPC methods are compared with the
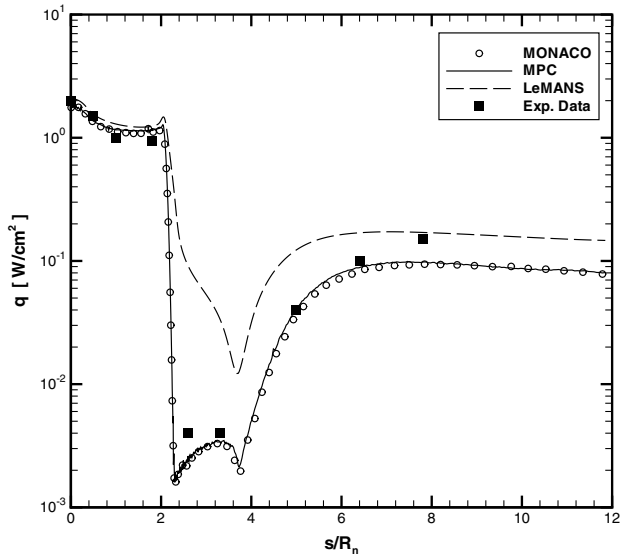
Fig. 8 Heating rate along the probe surface for the SR3 case 2 simulation.
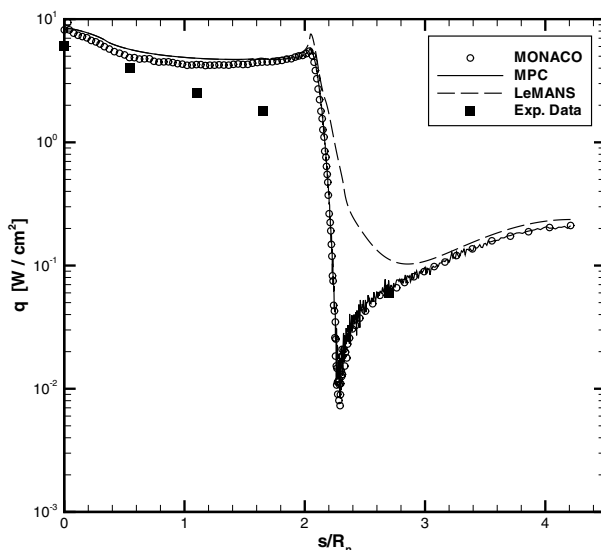
experimental measurements in Fig. 8. It should be noted that the measured signal along the probe base for this experiment (the fifth and sixth data points in Fig. 8) was so small that the heating magnitude could only be characterized as being less than 0.004 W/cm². As seen in Fig. 8, all three methods predict similar heating rates and agree well with the experimental data on the forebody of the planetary probe. As the flow rapidly expands around the shoulder, the NS simulation significantly overpredicts the heating rate around the shoulder and in the rarefied base region. This is due in part to the no-slip boundary condition applied to the NS equations. However, it is also due to the fact that, in assuming more collisionality, the NS equations predict more rapid flow changes (steeper gradients), which correspond to higher rates of heat and momentum flux. Figure 8 clearly shows that the MPC method (using fewer particles) is able to reproduce the heating rates predicted by a full DSMC simulation with a high degree of accuracy.
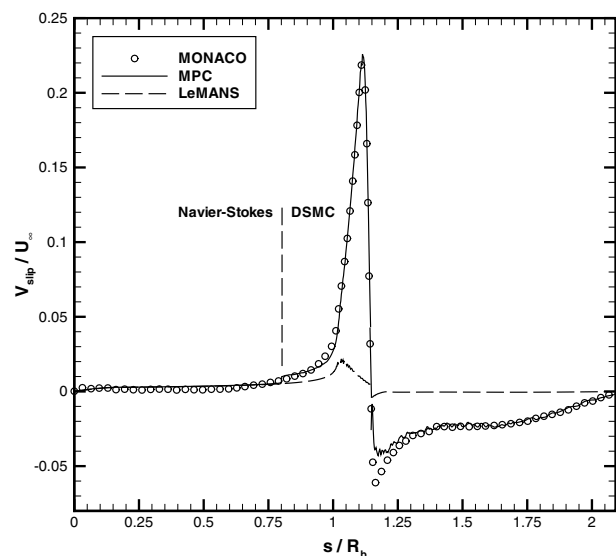
B. SR3 Case 3 Experiment

The final steady-state particle region (using a cutoff value of $Kn_{GL} = 0.05$) for the SR3 case 3 simulation was shown previously

in Fig. 2a. The particle region modeled using the DSMC method 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. 9a and compared with experimental measurements. Because the experiment was sting supported, but the SR3 case 3 simulations are 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. 9a, it can be seen that both predict the same peak heating rate at the stagnation point. Along most of the forebody, the DSMC method 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. Although the reason for this remains unclear, both DSMC and NS results agree very well with simulations performed by other researchers [15]. Around the capsule shoulder and along the capsule base, the DSMC method is seen to predict a much lower heating rate than the NS equations. In Fig. 9a, the MPC simulation is shown to reproduce DSMC results very accurately. In the highly compressed forebody 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 before the shoulder ($s/R_n = 1.6$), where the MPC method switches to DSMC, the heating rate is seen to 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. 9b after being normalized by the freestream velocity. The extraction is performed using Tecplot© where a curve is created parallel to the probe surface at a distance 10 μm normal to the surface. Data from the flowfield are then extracted along this curve, which involves Tecplot linearly interpolating data from the mesh nodes to the curve. In the forebody region, both full DSMC and NS simulations predict virtually zero-slip velocity at the probe surface. However, the DSMC simulation 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, because the velocity magnitude is extracted slightly off the surface, even the NS result will show a velocity magnitude greater than zero in some regions. The results in Fig. 9b agree both qualitatively and quantitatively with those of other researchers [15], 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
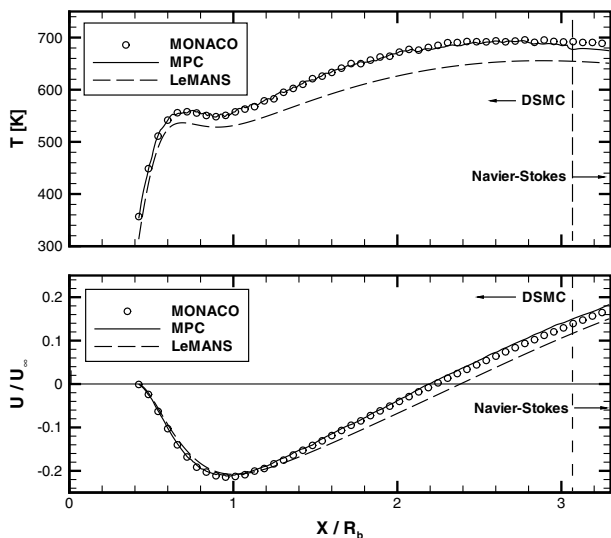

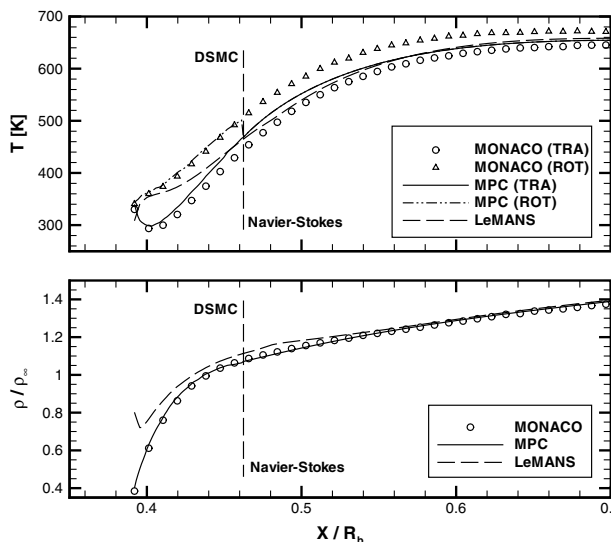a) **Experimental heating rates and numerical predictions**   b) **Velocity slip**

Fig. 9   Heating rates and velocity slip along the planetary probe surface for the SR3 case 3 simulation.

a) Flow properties in the planetary probe wake

b) Flow properties near the planetary probe shoulder

Fig. 10 Flowfield results for the SR3 case 3 planetary probe simulation.

reproduce full DSMC results along the entire planetary probe surface while successfully solving the NS equations in the forebody region, where both DSMC and NS predict no velocity slip. Both MPC results plotted in Figs. 9a and 9b lend further confidence in the use of the gradient-length Knudsen number (Kn_{GL}) and a cutoff value of 0.05 as a conservative measure of continuum breakdown in blunt-body flows.

For the SR3 case 3 simulation (for which the sting is not simulated), the translational temperature and normalized x velocity (U/U_∞) directly in the wake, along the axis of symmetry, are plotted for each method in the top and bottom of Fig. 10a, respectively. The DSMC method 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. 10a, DSMC is seen to predict the vortex to close 10% earlier than the NS simulation. The MPC method captures these flowfield 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 flowfield properties extracted from line C2 normal to the shoulder surface, shown previously in Fig. 2a, are displayed in Fig. 10b. Both translational and rotational temperatures are plotted in the top of Fig. 10b. The DSMC method 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. As seen in the top of Fig. 10b, close to the surface, the MPC method uses the DSMC method 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. Of course, the MPC method is unable to predict the thermal nonequilibrium that 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 deg cut is plotted in the bottom of Fig. 10b. At the surface, the density predicted by the DSMC method 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. Large Energy National Shock Facility Experiment

The initial and final DSMC and NS regions of the MPC planetary probe simulation corresponding to the LENS conditions are shown in Fig. 11. The LENS experimental conditions are similar to the SR3 case 3 conditions in that they result in a low global Knudsen number ($Kn_g = 0.002$) with a 2 order of magnitude variation in λ. For this flowfield, the forebody is well within the continuum regime but, due to the hypersonic velocity, the wake exhibits nonequilibrium effects. As shown in Fig. 11, the sting is now included in the numerical simulations. The temperature and density profiles in the dense forebody flow (along line C1 shown in Fig. 11) are shown in Fig. 12a. Here, the MPC method solves the continuum NS equations in the forebody region. The temperature and density predicted by the NS equations are seen to agree very well with results from a full DSMC simulation, except for the interior of the bow shock wave. The temperature and density profiles near the capsule shoulder (extracted along line C2 in Fig. 11) are shown in Fig. 12b. The trends are similar to the SR3 case 2 wake results shown earlier in Fig. 7b, except that the discrepancy between full NS and full DSMC is smaller because the flow is closer to continuum. Again, the MPC method is able to accurately reproduce full DSMC results for both translational and rotational temperature next to the capsule surface. The remainder of the wake in expanding thermal nonequilibrium is modeled with the
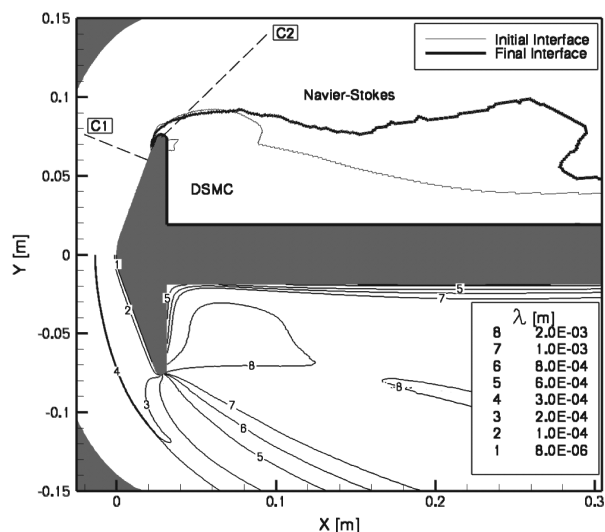

Fig. 11 Interface locations and variation of λ for the LENS simulation.

a) Temperature and density along C1
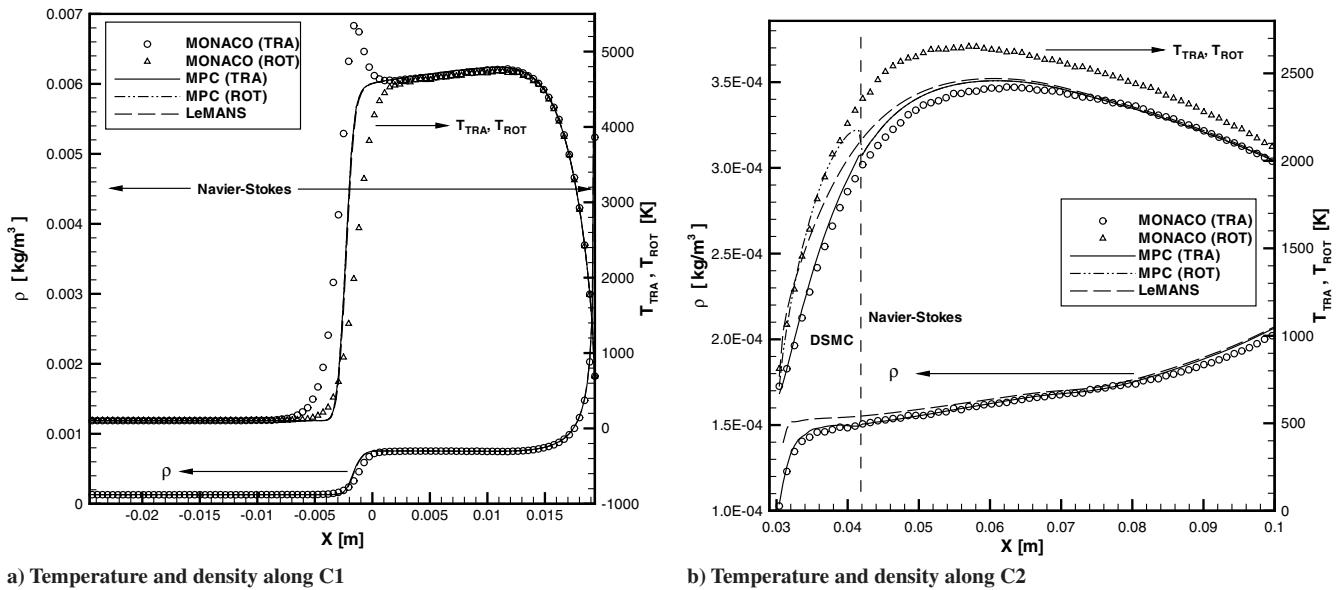
b) Temperature and density along C2

Fig. 12   Flowfield results for the LENS simulation.

NS equations, where the single translational–rotational temperature lies between those predicted by the full DSMC method.

Both the surface pressure and surface heating rates were measured in the LENS facility for these flow conditions. The coefficient of pressure $[C_p = (p - p_\infty)/(0.5\rho_\infty U_\infty^2)]$ for all simulations is compared with experimental data in Fig. 13a. For pressure, all simulations are in excellent agreement and in good agreement with experimental measurements in the wake. The agreement between simulation and experiment in the forebody is not as good. However, it should be noted that the pressure coefficient measured at the stagnation point was recorded as 3.0. This high value indicates that the pressure sensors may not have been appropriate for the ranges experienced in the forebody of this flow [15]. The heating rates along the planetary probe surface are shown in Fig. 13b. Again, all simulations agree well along the forebody of the probe. Recall the MPC method models the entire forebody flow using the NS equations, thereby eliminating significant computational expense while accurately reproducing full DSMC results. As seen in Fig. 13b, the heating rates predicted by the MPC method are in excellent agreement with full DSMC results along the entire probe and sting

surface, and these results are also in better agreement (compared with NS results) with experimentally measured heating rates.

IV. Computational Efficiency

In limiting the DSMC method only to localized regions of nonequilibrium flow, the MPC method reproduces full DSMC results while using a fraction of the particles. The precise manner in which MPC and DSMC simulation times are compared is detailed in [11]. The most relevant measures of efficiency are listed in Table 2 for each of the three simulations studied in this article. To summarize, 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. II.C (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 speedup factor is then the ratio of MPC to full DSMC simulation
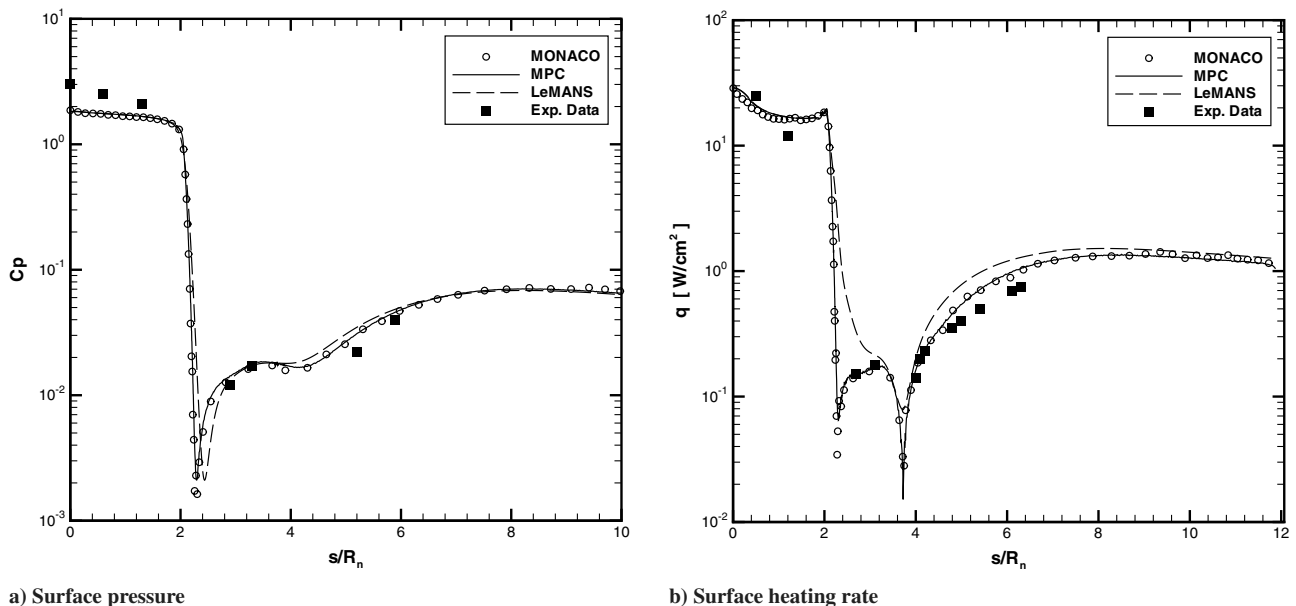

a) Surface pressure

b) Surface heating rate

Fig. 13   Planetary probe surface properties for the LENS simulation.

Table 2 Computational efficiency of the MPC method

SR3 case 2	Speedup	Memory	Particles	Initial NS	Time step	Steady state	Sampling
DSMC	100%	100%	100%	N/A	1	100%	100%
MPC	67% (1.52 times)	109%	64% (ratio of 1.57)	8.3%	20	34%	100%
SR3 case 3	Speedup	Memory	Particles	Initial NS	Time step	Steady state	Sampling
DSMC	100%	100%	100%	N/A	1	100%	100%
MPC	8% (12.5 times)	20%	14% (ratio of 7.1)	0.75%	150	30%	100%
LENS	Speedup	Memory	Particles	Initial NS	Time step	Steady state	Sampling
DSMC	100%	100%	100%	N/A	1	100%	100%
MPC	7.6% (13.1 times)	12%	8.2% (ratio of 12.2)	1.2%	133	62%	100%

times. It should be noted that the time required for a full DSMC simulation scales directly with the total number of particles in the simulation. Thus, if the initial NS solution and NS updates within an MPC simulation take negligible time compared to the simulation of particle regions, an MPC simulation can be approximated as a DSMC simulation that uses far fewer particles. The speedup factor achieved by the MPC method will thus scale directly with the number of particles eliminated (replaced with a continuum description) by the MPC simulation. The memory requirements also scale with the DSMC to MPC particle ratio, but also depend on the amount of DSMC-cell geometry replaced with large continuum NS-cell geometry. The time required for the initial NS solution and NS updates in an MPC simulation may not always be negligible compared with the time spent in DSMC regions. 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 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 the full DSMC solution. 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.

For the SR3 case 2 simulation, by solving the NS equations in the relatively small dense region after the shock and outside the wake, the MPC method requires 64% of the particles and produces accurate results 1.52 times faster than full DSMC simulation. It should be noted that, for the SR3 case 2 conditions, the mean-free path is sufficiently large such that no mesh refinement is required. That is, the same mesh density is used in both particle and continuum regions of the MPC simulation and, as a result, the MPC simulation actually requires more memory (109%) than a full DSMC simulation. However, the implicit time step used to update the NS regions is 20 times larger than the time step used in DSMC regions, and the MPC simulation reaches steady state in 34% of the iterations required by the full DSMC simulation. The modest speedup factor achieved for the SR3 case 2 flow is a result of the fact that this is a highly rarefied flow which can be simulated efficiently by DSMC. More specifically, the time required for initial NS solution and the NS updates within the MPC cycle is not negligible compared with the time spent simulating DSMC regions. On the other hand, the higher density planetary probe simulations (SR3 case 3 and LENS) are computationally intensive for full DSMC simulation. Because of the extremely small values of λ and τ_c in the forebody, mesh refinement within the MPC method is able to eliminate many computational cells and many particles. MPC simulation of SR3 case 3 conditions uses 14% of the particles and implicit NS time steps are 150 times larger than the time step used by DSMC in the forebody region. As a result, hybrid-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. Because 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 (12.5 times) than would be achieved by the elimination of particles alone (7.1 times). MPC simulation of the LENS experiment results in similar computational savings. Because of the even larger variation in λ (compare Figs. 2a and 11), the MPC method is able to eliminate more particles and produce accurate results 13.1 times faster than full DSMC simulation while requiring only 12% of the memory. The inclusion of the sting for the LENS simulation increases the time required for the MPC method to reach steady state and thus the speedup factor achieved is relatively close to the ratio of particles eliminated by the continuum NS regions.

V. Conclusions

Compared with previous MPC research, application of the MPC method to the planetary probe configurations (where the global Knudsen number is low) demonstrates additional capabilities of the method. New capabilities include substantial variation in both mesh density and simulation time step between continuum and particle regions, as well as the ability to handle axisymmetric flows requiring cell-weighting factors in DSMC regions.

For all simulations presented in this article, the MPC method reproduces the flowfield and surface properties computed by full DSMC simulations with high accuracy. The MPC method is shown to successfully capture thermal nonequilibrium inside the shock wave and next to the planetary probe surface for the SR3 case 2 simulation and inside the rapidly expanding wake regions for all simulations considered. The MPC method is also shown to accurately reproduce the heating rates, surface pressure, temperature jump, and velocity slip predicted by full DSMC simulation. Overall, DSMC and NS solutions are shown to be in close agreement with each other and with experimental measurement in the dense forebody flow. However, DSMC and MPC methods more accurately predict experimental results in the rarefied capsule shoulder and base regions, as well as along the sting. In addition, the generality of the MPC method is demonstrated in that highly accurate results are obtained even for a highly rarefied flow, where a relatively small region is simulated using the continuum NS equations (SR3 case 2 conditions).

The gradient-length Knudsen number is shown to position the particle-continuum interfaces in near-equilibrium regions for the blunt-body flows studied in this article. This parameter accurately detects the postshock region where the continuum NS equations are valid and also accurately detects the surface location (near the capsule shoulder) where the DSMC method is required to accurately predict the surface properties (heating rates and velocity slip).

The higher density (low global Knudsen number) planetary probe experiments studied in this article involve 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 forebody enables the planetary probe problem to be solved efficiently by a modular particle-continuum numerical

method. For the first time, a hybrid particle-continuum numerical method is shown capable of achieving an order of magnitude decrease in both computational time and memory compared with full DSMC simulation. Specifically, for the SR3 case 3 experiment, the MPC method is shown to reproduce full DSMC flowfield and surface property results with a high level of accuracy, using 14% of the particles and obtaining the solution approximately 12.5 times faster than full DSMC, while requiring only 20% of the memory. For the LENS experiment, the MPC method is shown to accurately reproduce full DSMC results, using 8.2% of the particles and obtaining the solution approximately 13.1 times faster than full DSMC, while requiring only 12% of the memory. Application of the modular particle-continuum method to the planetary probe benchmark problems clearly shows 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. These variations in spatial and temporal scales become larger for higher reentry velocities experienced during return missions from the moon and Mars. Finally, the three-dimensional DSMC simulation of such multiscale flows will require enormous computational resources. Application of a hybrid particle-continuum method to such three-dimensional, high-speed problems is expected to result in even more significant improvements in performance.

Acknowledgments

References

[1] Bird, G. A., *Molecular Gas Dynamics and the Direct Simulation of Gas Flows*, Oxford Univ. 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.
doi:10.2514/3.26546

[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.
doi:10.2514/3.781

[6] 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*, AIAA Paper 96-0353, 1996.

[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.
doi:10.2514/2.3349

[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.
doi:10.2514/2.3647

[9] Wang, W. L., and Boyd, I. D., "Hybrid DSMC-CFD Simulations of Hypersonic Flow over Sharp and Blunted Bodies," *36th AIAA Thermophysics Conference*, AIAA Paper 03-3644, 2003.

[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, No. 5, 2004, pp. 768–777.
doi:10.1115/1.1792275

[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*, 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*, AIAA Paper 07-613, 2007.

[15] 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.
doi:10.2514/2.6252

[16] Gnoffo, P., and White, J., "Computational Aerothermodynamic Simulation Issues on Unstructured Grids," *37th AIAA Thermophysics Conference*, AIAA Paper 04-2371, 2004.

[17] Candler, G. V., "Unstructured Grid Approaches for Accurate Aeroheating Simulations," *39th AIAA Thermophysics Conference*, AIAA Paper 07-3959, 2007.

[18] 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*, AIAA Paper 06-3773, 2006.

[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, No. 1, 1989, pp. 135–150.
doi:10.1016/0045-7930(89)90012-1

[20] 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.
doi:10.1063/1.868720

[21] Wang, W. L., and Boyd, I. D., "Predicting Continuum Breakdown in Hypersonic Viscous Flows," *Physics of Fluids*, Vol. 15, No. 1, 2003, pp. 91–100.
doi:10.1063/1.1524183

[22] Garcia, A. L., and Alder, B. J., "Generation of the Chapman–Enskog Distribution," *Journal of Computational Physics*, Vol. 140, No. 1, 1998, pp. 66–70.
doi:10.1006/jcph.1998.5889

[23] 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

[24] Boyd, I. D., "Analysis of Rotational Nonequilibrium in Standing Shock Waves of Nitrogen," *AIAA Journal*, Vol. 28, No. 11, 1990, pp. 1997–1999.
doi:10.2514/3.10511

[25] 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

[26] Holman, T., and Boyd, I. D., "Numerical Investigation of the Effects of Continuum Breakdown on Hypersonic Vehicle Surface Properties," *40th AIAA Thermophysics Conference*, AIAA Paper 08-3928, 2008.

[27] Deschenes, T., Boyd, I. D., and Schwartzentruber, T. E., "Incorporating Vibrational Excitation in a Hybrid Particle-Continuum Method," *40th AIAA Thermophysics Conference*, AIAA Paper 08-4106, 2008.