SAND2019-14501C

Velocity-space Hybridization of DSMC and a Boltzmann Solver

G. Oblapenko*, D. Goldstein† and P. Varghese‡
The University of Texas at Austin, Austin, Texas, 78712

C. Moore§
Sandia National Laboratories, Albuquerque, New Mexico, 87185

A new method for modeling rarefied gas flows is discussed in the present work. Based on hybridization of DSMC and Discrete Velocity Method (DVM) representations of the velocity distribution function, it is aimed at improving the resolution of the tails of the distribution function (compared to DSMC), and computational efficiency (compared to DVM). Details of the method, such as the collision algorithm and the particle merging scheme are discussed. The hybrid approach is applied to the study of noise in a Maxwellian distribution, computation of electron-impact ionization rate coefficient, as well as numeric simulation of a supersonic Couette flow, and compared to pure DSMC and DVM approaches in terms of accuracy, computational speed and memory use.

I. Nomenclature

β	=	Scaled velocity grid spacing
$\hat{\boldsymbol{\eta}}, \hat{\boldsymbol{\zeta}}$	=	Scaled velocities
η_r	=	Reference velocity
σ_r	=	Reference collision cross-section
$\hat{\sigma}_t$	=	Scaled total collision cross-section
$\hat{\phi}$	=	Scaled distribution function
C_{RMS}	=	DVM noise parameter
E_{ion}	=	Ionization energy of argon
$\mathcal{F}$	=	DSMC particle region in hybrid approach
$\hat{g}$	=	Magnitude of scaled relative collisional velocity
k	=	Boltzmann's constant
L	=	Reference length
M	=	Number of merging cells
m_e	=	Electron mass
N_p	=	Number of particles
N_s	=	Number of chemical species
N_v	=	Number of discrete velocity grid points
n_r	=	Reference number density
$\hat{t}$	=	Scaled time
$\bar{\mathbf{u}}$	=	Scaled macroscopic velocity
$\mathcal{Q}$	=	DVM region in hybrid approach

*Postdoctoral Research Fellow, Oden Institute for Computational Engineering and Sciences, The University of Texas at Austin, 2201 E 24th St, Stop C0200

†Director of Computational Flow Physics Laboratory, Department of Aerospace Engineering and Engineering Mechanics, 2617 Wichita St., Stop C0600, AIAA Associate Fellow

‡Director of the Center for Aeromechanics Research, Department of Aerospace Engineering and Engineering Mechanics, 2617 Wichita St., Stop C0600, AIAA Associate Fellow

§Principal Member of Technical Staff, Plasma Theory & Simulation, Sandia National Laboratories, 87185-1168

II. Introduction

DSMC (Direct Simulation Monte Carlo) [1] is perhaps the most widely used approach for modelling rarefied gas flows, and has been successfully used to simulate complex flows with internal energies, chemical reactions, and radiation. However, drawbacks of the method include statistical noise, as well as the high number of particles (and, as a consequence, increased computational time and memory usage) required to accurately model trace populations, such as those of high-lying internal energy states, high-velocity tails of the distribution function, or of trace species. Therefore, DSMC may not perform well in cases where trace populations have a signficant influence on the flow, such as can be the case in atmospheric pressure gas discharges [2]; in a beam skimmer set-up [3]; or in the case of a high-velocity beam colliding with a cold gas [4].

In the discrete velocity family of methods, the velocity distribution function is restricted to a defined set of points in velocity space, thus allowing replacement of the Boltzmann collision operator with a summation, and application of finite-difference or finite-volume schemes to compute the change in the velocity distribution function due to convection. However, these methods require interpolation schemes to ensure mass, momentum and energy conservation during the collision process, such as those developed in [5–9]. Discrete velocity methods are able to accurately resolve very sparsely populated tails of the distribution function, populations of high-energy internal states [10], and model chemical reactions with very low reaction probabilities [11]. They are also well-suited to use of variance reduction [12, 13], a technique introduced in DVM modelling of low-Knudsen number flows [14] and later adapted to DSMC with the aim of reducing stochastical noise [15].

Given the high efficiency of DSMC at capturing the main features of rarefied flows, and the well-suitedness of DVM for resolving low populations, it is of interest to investigate a hybrid approach, in which the tails of the distribution function are modelled with DVM, and the bulk is represented by DSMC particles. Similar approaches (coupling DSMC with a discontinuous Galerkin and BGK solvers) have been studied previously [16, 17] and were shown to be better at resolving distribution function tails, but only zero-dimensional single-species flows were considered, and the influence of various simulation parameters on the accuracy of such schemes was not investigated; their efficiency in contrast to standard DSMC or discontinuous Galerkin/BGK solvers was also not studied. In a similar fashion, we have previously investigated the particulars of coupling a discrete velocity method and DSMC for a single-species flow, where a hybrid approach was shown to provide a computational speed-up over a pure discrete-velocity solver, while retaining the possibility to accurately resolve tails of the distribution function [18]. The aim of the present work is to improve upon the developed approach and apply it to multi-species and non spatially homogeneous problems.

III. Description of Velocity-Space Hybridization

By restricting the distribution function to a finite set (grid) of points $\mathcal{D} \subset \mathbb{R}^3$ in velocity space (and assuming that the grid is uniform), one can write the Boltzmann equation in a discretized scaled form [12]:

$$\frac{\partial \hat{\phi}}{\partial \hat{t}} + \hat{\boldsymbol{\eta}} \cdot \nabla_{\hat{\mathbf{r}}} \hat{\phi} = \frac{1}{Kn} \sum_{\hat{\boldsymbol{\zeta}} \neq \hat{\boldsymbol{\eta}}} \left[\hat{\phi}(\hat{\boldsymbol{\eta}}') \hat{\phi}(\hat{\boldsymbol{\zeta}}') - \hat{\phi}(\hat{\boldsymbol{\eta}}) \hat{\phi}(\hat{\boldsymbol{\zeta}}) \right] \hat{g} \hat{\sigma}_t \,., \tag{1}$$

where the Knudsen number is defined as λ_r / L, and the reference mean free path is defined as $1/(\sigma_r n_r)$.

In order to avoid direct evaluation of the collision integral in Eq. (1) (which would require N_v^2 operations), a stochastic Monte-Carlo evaluation scheme is used [4, 12]: points on the velocity grid are sampled from the cumulative distribution function, the density at those points is depleted by a small amount, and a new velocity of the depleted mass is computed using standard collision mechanics. Since the post-collision velocity does not necessarily lie on the velocity grid, a remapping scheme is used to re-distribute this mass across neighbouring grid points in a fashion which conserves mass, momentum, and energy [6, 8]. This approach to the discrete-velocity modelling of the Boltzmann equation has been called the "Quasi-Particle Simulation Method" (QUIPS) [11], and will be referred to as such in the present work.

The number of collisions performed is inversely proportional to C_{RMS}^2, an adjustable parameter which controls the noise in the distribution function due to the stochastic evaluation of the collision integral. Higher values of C_{RMS} produce noisier results, but require less computational resources.

Change in the velocity distribution function due to convection in the present work is computed using a first-order upwind scheme.

In an approach that is hybrid in velocity space, one could consider different options of representing the velocity distribution function: for example, one could represent the bulk of the distribution with the QUIPS solver and the tails with DSMC particles; or have several disjoint regions in velocity space were DSMC particles are used, or perhaps

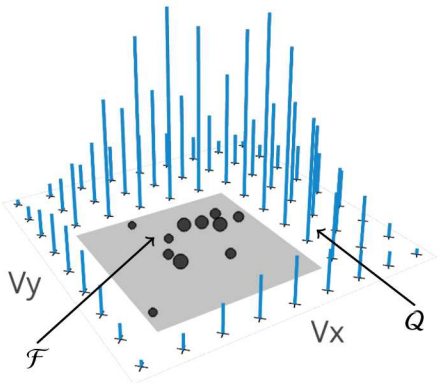

Fig. 1 Schematic of the hybrid representation of the velocity distribution function. Blue bars denote the distribution function defined at the discrete grid points in $\mathcal{Q}$ (with their height corresponding to the value of the distribution function at that point), the black circles denote the DSMC particles (with their size corresponding to their computational weight) in region $\mathcal{F}$ (shown here as the grey rectangle).

even interspersing DSMC particles among the quasi-particles on the velocity grid. However, since our interest is in accurately modelling processes that are affected by high-velocity particles (such as electron-impact ionization), we explore the concept of representing the bulk of the velocity distribution function with a (relatively) small number of DSMC particles and using QUIPS to model the tails of the distribution.

To perform hybridization in velocity space, a region $\mathcal{F} \subset \mathbb{R}^3$ is chosen in which particles can have any velocity and computational weight (DSMC region), and the discrete velocity distribution function is non-zero only outside this region (we will denote the set of discrete velocity grid points outside $\mathcal{F}$ as $\mathcal{Q}$). Fig. 1 shows a schematic of such a hybrid representation of the velocity distribution function (for a fixed value of the Z velocity component).

Collisions between DSMC particles are performed using the No-Time-Counter collision scheme [1], modified for use with variable-weight DSMC particles [19], while collisions involving mass in $\mathcal{Q}$ (that is, DSMC-QUIPS collisions and QUIPS-QUIPS collisions) are computed using the QUIPS Monte-Carlo scheme described above.

At each collision step, new particles are created in $\mathcal{F}$, either as a result of some the points used in the remapping procedure lying inside $\mathcal{F}$, or as a result of a post-collision velocity being located inside $\mathcal{F}$. In order to avoid rapid growth in the number of particles, at each computational step, a merge is performed in $\mathcal{F}$. The DSMC particle region is split into M equal-volume cells, and in each cell in velocity space, all the particles are merged down to 2 particles (in order to conserve mass, momentum, and energy) [20, 21]. Since the remapping scheme used in the discrete velocity method produces negative mass at some of the points, a net negative mass and/or energy may occur in a merging cell; in case this happens, a random neighboring cell is selected and an attempt to perform a merge of all the particles in the two cells is performed.

The merging routine has $O(N_p)$ complexity, and requires $O(M^3)$ additional storage. It should be noted however, that this required additional memory is not dependent on the number of physical cells used in the simulation. Collisions are performed in a cell in physical space and merging is done immediately afterwards, before collisions in the next cell are computed. This allows us to re-use the merging-associated storage space in the next cell in physical space. Performing merging immediately after collisions are performed in a cell also allows us to avoid a significant increase in the total number of particles – while some additional memory is required to store the particles created during the collision step, it is freed after the merging is performed, and thus the number of physical cells used in the simulation does not influence the amount of additional memory required.

IV. Numerical Results

In this section, numerical results for several test cases are presented. Computations were performed with the pure QUIPS solver, the developed hybrid solver, and the open-source fixed-weight DSMC code SPARTA [22]. In the QUIPS solver, initialization is performed as follows: the initial velocity distribution function is evaluated at each point on the discrete grid, and the values are then scaled so that the correct density is obtained. Such an approach does not give the exact values of the momentum and energy and in case of an initial Maxwellian velocity distribution function

does not produce a truly equilibrium distribution (since a Maxwellian evaluated on a discrete velocity grid is not the equilibrium distribution function for a discrete velocity method [23]), but with a sufficiently wide and fine grid the discretization error is small. For all the cases considered in the present work, the relative error in the energy at $t = 0$ was less than 10^{-5} %. For the hybrid code, the same initialization scheme is used, but mass at the discrete velocity grid nodes inside $\mathcal{F}$ is created as DSMC particles. Immediately after initialization and before the first timestep, particle merging is performed, since the number of discrete velocity grid points that are inside $\mathcal{F}$ may be larger than the target number of DSMC particles.

A. BKW relaxation

In order to verify the collision procedures in the developed hybrid scheme, the BKW (Bobylev-Krook-Wu) relaxation test case was considered [24, 25]. The BKW problem is an analytic time-dependent solution of the Boltzmann equation for the case of a pseudo-Maxwell collision cross-section ($\hat{\sigma} \propto 1/\hat{g}$). A $29 \times 29 \times 29$ grid extending to $\pm 3.5\eta_r$ was used (that is, the velocity grid extends to ±3.5 times the mean thermal speed η_r in each Cartesian direction). The region $\mathcal{F}$ was considered to be a cube centered around the origin in velocity space, extending from $-1.375\eta_r$ to $1.375\eta_r$ in each Cartesian direction. A timestep of 0.1 the mean collision time was used.

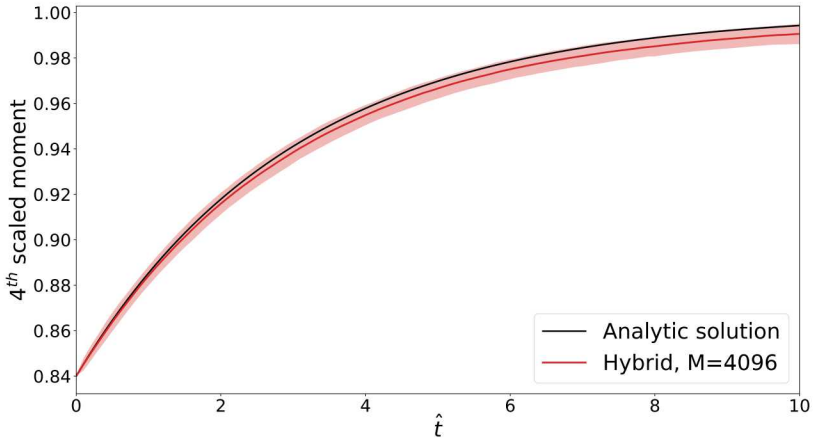

Fig. 2 Value of the 4^{th} scaled moment as a function of the scaled time. Black curve shows the analytic result for the temporal evolution of the scaled 4^{th} moment and the red curve shows the ensemble-averaged result obtained with the hybrid code, with the semi-transparent region corresponding to a range of ± two standard deviations (obtained through ensemble averaging).

Figure 2 shows the temporal evolution of the 4^{th} scaled moment computed analytically (black curve) and using the hybrid scheme with a large number (4096) of merging cells (red curve), corresponding to 8192 DSMC particles. The hybrid code results are averaged over 200 ensembles, and a range of plus/minus two standard deviations is shown by the semi-transparent region. It can be seen that the numerical solution relaxes towards equilibrium at approximately the expected rate; the difference between the analytic and numerical solution at larger times is due to bias introduced by the merging procedure and error due to the discretization of velocity space; both could be remedied with improved resolution (more merging cells and thus more DSMC particles).

To assess the necessity of an adaptive merging scheme, a single-species gas initialized with a Maxwellian distribution was considered. Figure 3 shows the number of particles in $\mathcal{F}$ after a merge is performed when using either the non-adaptive (if a merge within a cell fails due to negative mass and/or energy, nothing is done) or adaptive (if a merge within a cell fails, it and a neighboring cell are merged together) merging schemes. It can be seen that use of a non-adaptive approach leads to occasional increases in the number of particles inside $\mathcal{F}$, which can be an issue of multi-dimensional computations with large number of physical grid cells. Use of an adaptive merging scheme, however, resolves these issues (at the expense of a small loss of fidelity (smearing) in the representation of the distribution function).

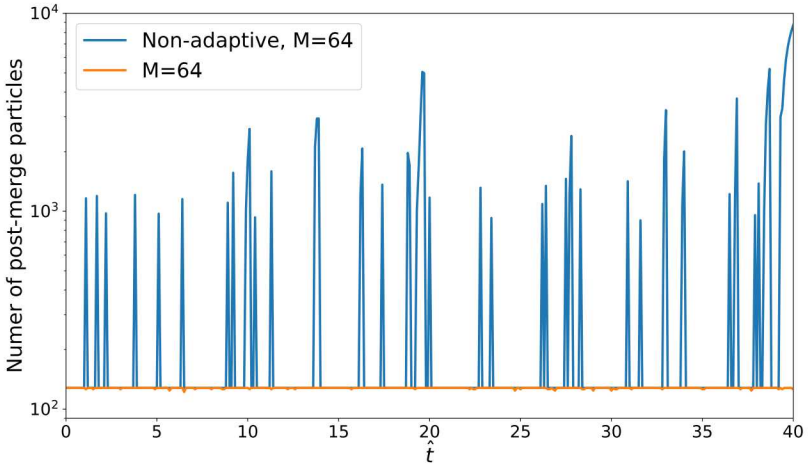

Fig. 3 Comparison of post-merge number of particles when using non-adaptive and adaptive merging schemes.

B. Single-species Maxwellian

To assess the influence of the C_{RMS} parameter on the noise in the tails of the distribution function, the following spatially homogeneous test case was considered: An argon gas was initialized with a Maxwellian dsitribution at 273 K, and allowed to evolve over 100 collision times (1000 timesteps). Since our interest is in maintaining quiet tails of the velocity distribution function, we considered the noise in the 8^{th} moment of the distribution function as a measure of the accuracy of a particular code in capturing the high-velocity tails, since this metric gives more weight to the higher-velocity tails. For the purposes of this test case, the region $\mathcal{F}$ was considered to be a cubic region extending from $-1.375\eta_r$ to $1.375\eta_r$ in each Cartesian direction in velocity space. Two grids were considered: a "fine" $29 \times 29 \times 29$ grid, and a "coarse" $15 \times 15 \times 15$ grid.

In order to compare compare the codes both in terms of computational efficiency and degree of accuracy, the metric of computational time per collision step (which for the hybrid code includes the time spent merging particles in $\mathcal{F}$ after a collision) was considered, and plotted against the root mean squared error (RMSE) of the 8^{th} moment of the velocity distribution function. Since the scaled equilibrium values of the higher-order moments might differ from 1 in the QUIPS and hybrid approaches (since that depends on the velocity grid extent and spacing used), time-averaging was performed for 1000 timesteps to obtain the mean value of the scaled moment for each simulation run.

Figure 4 shows the results obtained with the pure QUIPS and hybrid approaches, as well as with the SPARTA fixed-weight DSMC code [22]. It can be seen that for this test case, a fixed-weight DSMC code is computationally more efficient than both the pure QUIPS and hybrid approaches, but requires a significant number of particles (up to 10 million) to achieve a comparably low level of noise, which does not scale to complex 2-D and 3-D geometries due to the large amount of memory required. Decreasing the value of C_{RMS} (and thus reducing the error in the QUIPS and hybrid methods) comes only at the price of an increased computational cost, while the memory usage remains the same, regardless of the value of C_{RMS} used. The hybrid code provides a speed-up by a factor of 2 to 3 compared to the pure QUIPS approach, while retaining a similar amount of noise unless very small values of C_{RMS} are used — at a low enough value of C_{RMS}, the noise due to the coarse representation of the bulk of the distribution function with a few DSMC particles (and the DSMC collision scheme used), as well as the merging procedure, becomes the main source of error, rather than the error in the stochastic evaluation of the collisions involving QUIPS mass (which is controlled by the value of C_{RMS}).

C. Electron-impact ionization

In order to evaluate the developed hybrid scheme for a more physically tangible problem, we considered the test case of electron-impact ionization of argon. We consider equilibrium argon and electron distributions, and compute the electron-impact ionization rate. Similar to the previous section, computational time per collision step is used as a metric of computational efficiency, while the error relative to the analytic rate is used as a metric of accuracy of the code.

The argon gas was kept at 300 K, and a range of electron temperatures from 1 eV to 100 eV was considered. The number density of argon was taken to be 10^{23} m^{-3}, and the number density of electrons was taken to be 10^{20} m^{-3}.

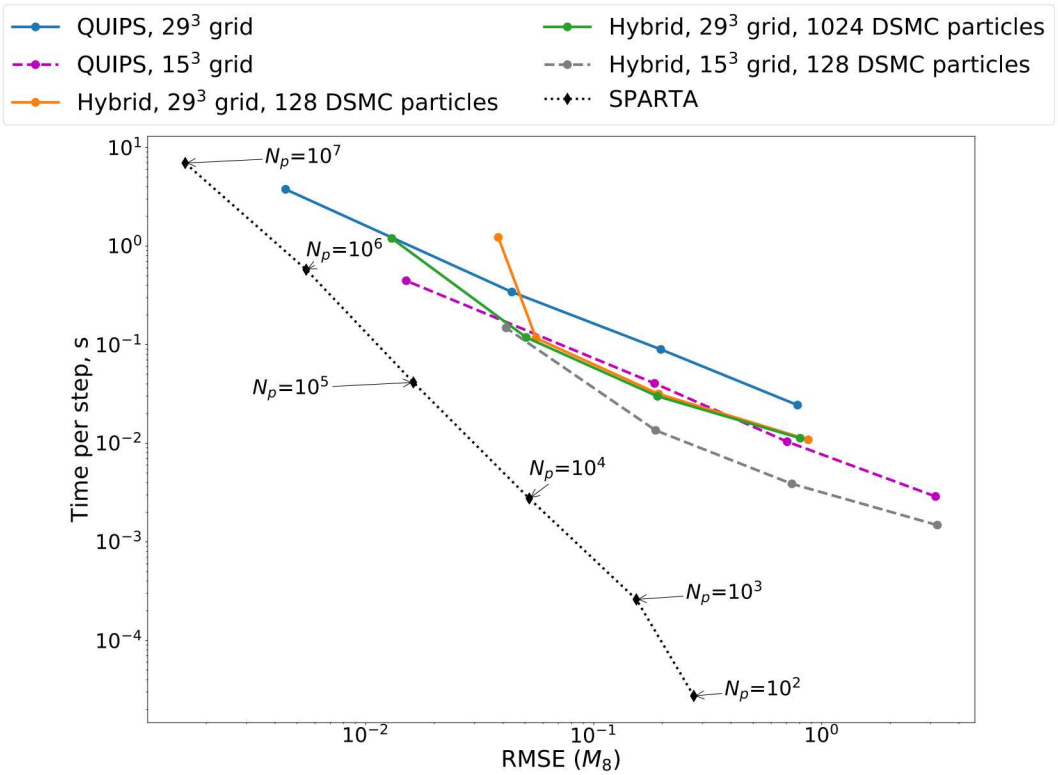

Fig. 4 Computational time per collision step plotted against the RMSE of the 8^{th} moment; different points on the SPARTA curve correspond to different numbers of particles used in the simulation, different points on the curves for the pure QUIPS and hybrid codes correspond to different values of C_{RMS}.

The noise in the electron-impact ionization rate is governed by two factors: noise in the tails of the electron velocity distribution function (for the purposes of this problem, tails correspond to the part of the distribution function which has a kinetic energy higher than the ionization energy of argon, 15.76 eV) and noise in the collision scheme used. A simplified collision algorithm was used in the present study: only argon-argon and electron-electron collisions are actually computed (for the latter, a simple rigid sphere model is used — while not entirely physical for charged-charged particle collisions, in the QUIPS and hybrid methods we only perform electron-electron collisions to introduce noise in the tails of the electron distribution function), which allows us to keep the argon atoms and electrons at constant temperatures. For argon-electron collisions, the collision sampling is performed as usual, and collisions are tested to check whether ionization occurs (the QUIPS and hybrid codes, due to their capabilities to deal with variable mass, can split each collision into elastic and inelastic parts [26], and thus perform an ionization reaction for any electron-argon collision with sufficient relative translational energy), but no actual collisions are performed. A similar approach is performed in the SPARTA code — for argon-electron collisions, the selected collision pairs are tested to see whether an ionization reaction occurs, but the collisions are not actually performed. Thompson's model is used for the ionization reaction cross-section [27].

The choice of hybridization options in multi-species flows poses an additional challenge — each species can be treated using a pure QUIPS method, a hybrid method, or a variable-weight DSMC method, leading to 3^{N_s} possible options just for the choice of schemes to use, where N_s is the number of chemical species considered. For the present problem, the quantity of interest (the ionization rate) is governed by the electron distribution function (since the velocities of the argon atoms are significantly lower and do not influence the ionization probability); thus it makes sense to treat the argon atoms with a variable-weight DSMC scheme, and treat the electrons with either a hybrid scheme or a pure QUIPS method. For the hybrid scheme, it makes sense to treat any electrons with a kinetic energy higher than the ionization energy of argon with the QUIPS method (in order to minimize noise in the part of the distribution relevant to our problem), and use DSMC to model electrons with lower kinetic energies.

Since our region $\mathcal{F}$ is cubic-shaped, in order to ensure that no DSMC electron particle in that region has an energy

higher than the ionization energy of argon, the bounds of the region are from $-v_{ion}/\sqrt{3}$ to $v_{ion}/\sqrt{3}$ in each Cartesian direction, where $v_{ion} = \sqrt{2E_{ion}/m_e}$. Since such a choice of DSMC region is independent of electron temperature, at sufficiently high electron temperatures, it will contain very little of the mass of the electron velocity distribution function, and velocity-space hybridization (at least, with such a fixed choice of the region $\mathcal{F}$) might not be beneficial any more. For argon, 128 DSMC particles were used; in case the electrons were modelled using the hybrid scheme, either 16 or 128 DSMC particles were used to represent the bulk of the electron velocity distribution function. For the argon particles, in case the QUIPS method is used to simulate them, a 14^3 velocity grid was used, while for the electron particles a 28^3 velocity grid was used.

Due to the large difference between the number densities of argon and electrons, a fixed-weight DSMC code such as SPARTA would require a 1000 computational argon particles for each computational electron particle, which would be unfeasible in a complex 2-D or 3-D simulation. However, for the purpose of this spatially homogeneous study, results obtained with SPARTA are included.

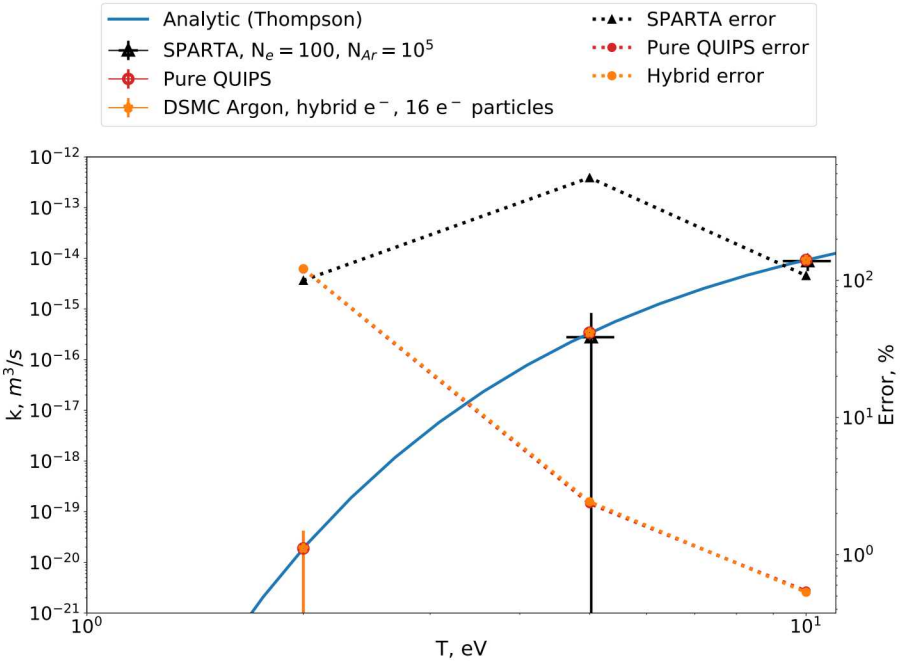

Fig. 5 Ionization rate as a function of temperature and errors in ionization rate computed by the SPARTA, QUIPS and hybrid codes. For the SPARTA computation, 100 electron particles and 10^5 argon particles were used; for the QUIPS and hybrid codes, the value of C_{RMS} used was equal to $1 \cdot 10^{-4}$. The uncertainty bars show a range of ± one standard deviation for the ionization rate coefficient and, for the SPARTA results, for the electron temperature of the sampled electron distribution.

Figure 5 shows the values of the ionization rate coefficient computed with the SPARTA, QUIPS and the hybrid codes for different values of the electron temperature, as well as the error (compared to the analytic rate). For low electron temperatures (2 eV), SPARTA gives a zero value for the rate coefficient due to the fact that no sufficiently high-velocity electron particles are sampled, and thus, the relative error is 100 %. Except for the temperature of 2 eV, the QUIPS code provides an order of magnitude lower error in computing the ionization rate coefficient; the hybrid code gives results very close to those computed with the pure QUIPS approach.

As for the previous test case, we now look at computational time per step plotted against the relative error in the ionization rate coefficient to assess the efficiency of various codes and compare some of the possible hybridization options discussed above.

Figure 6 show the computational time per collision step plotted against the relative error in the ionization rate coefficient for various values of the electron temperature. It can be seen that for all temperatures, it is possible to get a relatively low error when using the QUIPS or the hybrid code, if one uses a sufficiently low value of C_{RMS}. At the higher temperatures, at a certain point decreasing the C_{RMS} does not lead to a less significant decrease in the error (the

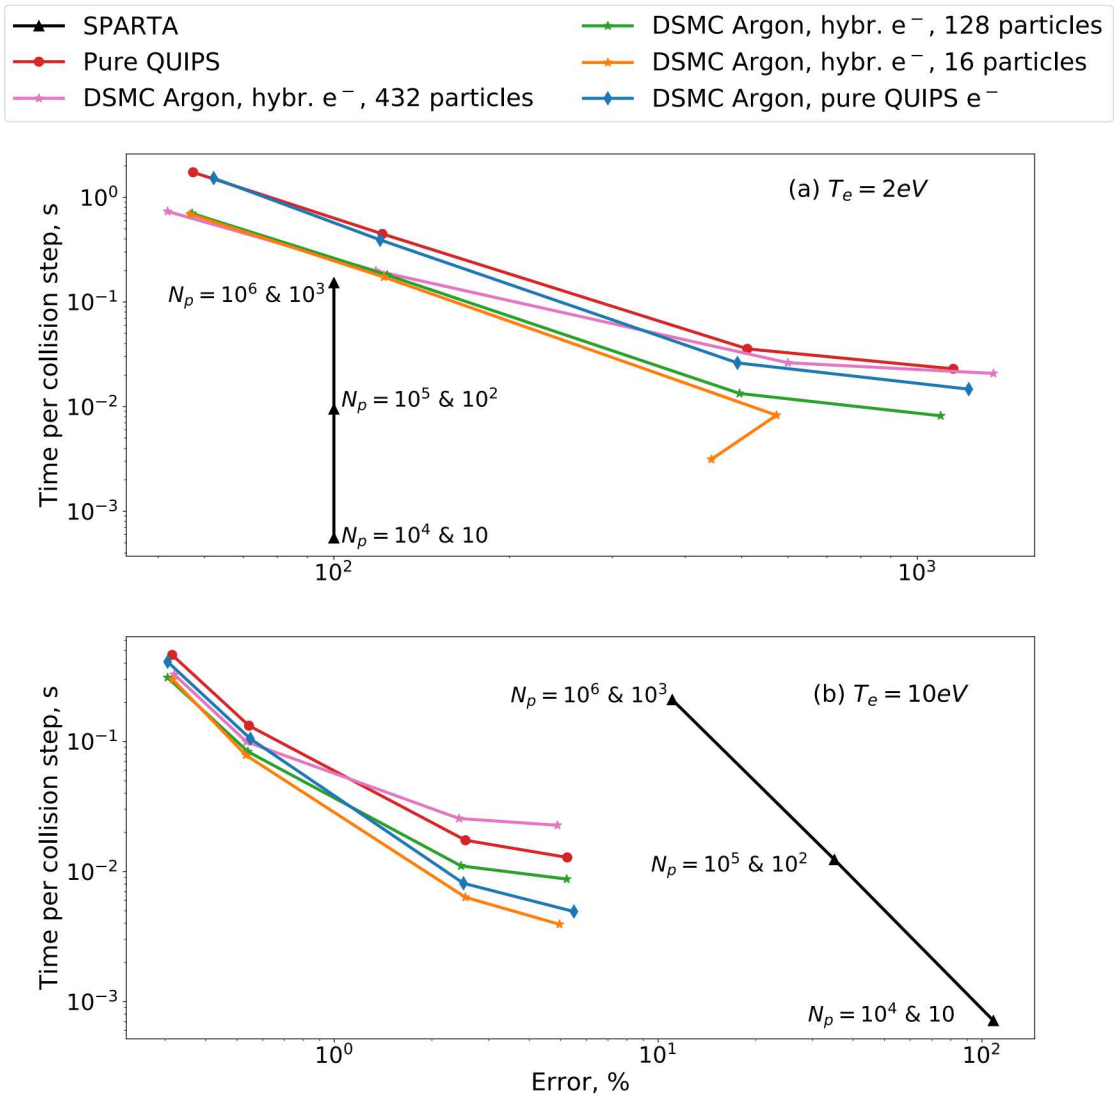

Fig. 6 Computational time per collision step plotted against the relative error in the ionization rate coefficient for electron temperatures of 2 and 10 eV. Different points on the SPARTA curves correspond to different numbers of particles used (first number is the number of argon particles, second is the number of electron particles); different points on the curves for the pure QUIPS and hybrid codes correspond to different values of C_{RMS} for the electron-argon collisions.

pure QUIPS and hybrid curves become steeper), which is due to the velocity space discretization becoming the main source of error, and possibly the noise in the electron-electron collisions.

In case one uses a pure (variable weight) DSMC scheme for the argon atoms and represents the electrons using the QUIPS method, a speed up can be achieved compared to the pure QUIPS code, but only at higher values of C_{RMS}. At lower values of the noise parameter, most of the computational time is spent on the electron-argon collisions, which are treated with the QUIPS Monte-Carlo collision scheme, and no benefit from representing argon with DSMC particles is gained.

However, if one uses DSMC particles to represent the bulk of the electron velocity distribution function, a consistent speed-up can also be achieved for lower temperatures in the whole range of C_{RMS} values considered; reducing the number of electron DSMC particles has for the most part no effect on the error, but increases the speed-up at higher values of C_{RMS}. Conversely, using a large number of DSMC particles for the electrons (432 particles) makes the hybrid code even less efficient than a pure QUIPS approach, especially at higher temperatures. This is due to several factors.

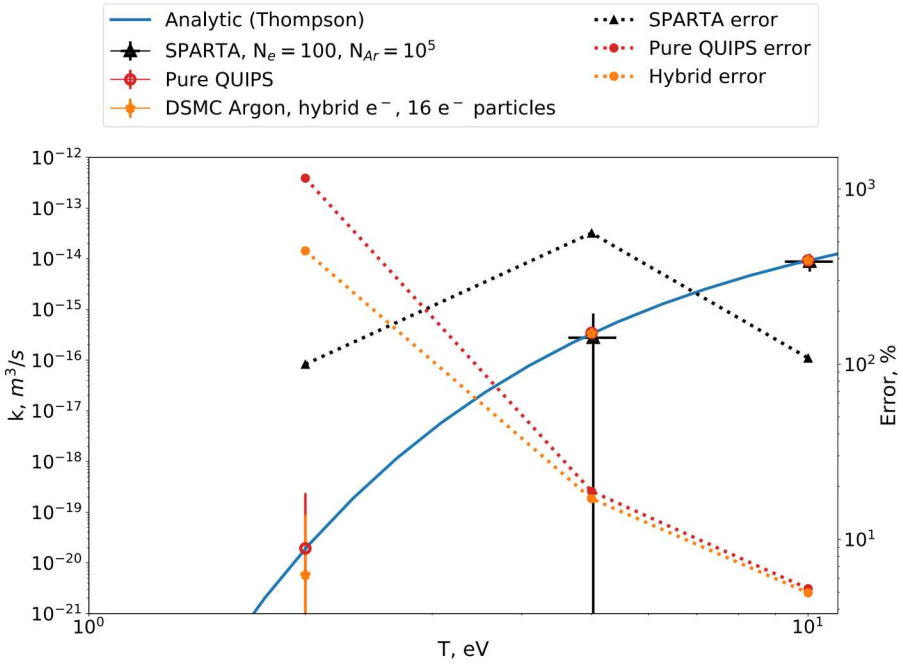

Fig. 7 Ionization rate as a function of temperature and errors in ionization rate computed by the SPARTA, QUIPS and hybrid codes. For the SPARTA computation, 100 electron particles and 10^5 argon particles were used; for the QUIPS and hybrid codes, the value of C_{RMS} used was equal to $1 \cdot 10^{-3}$.

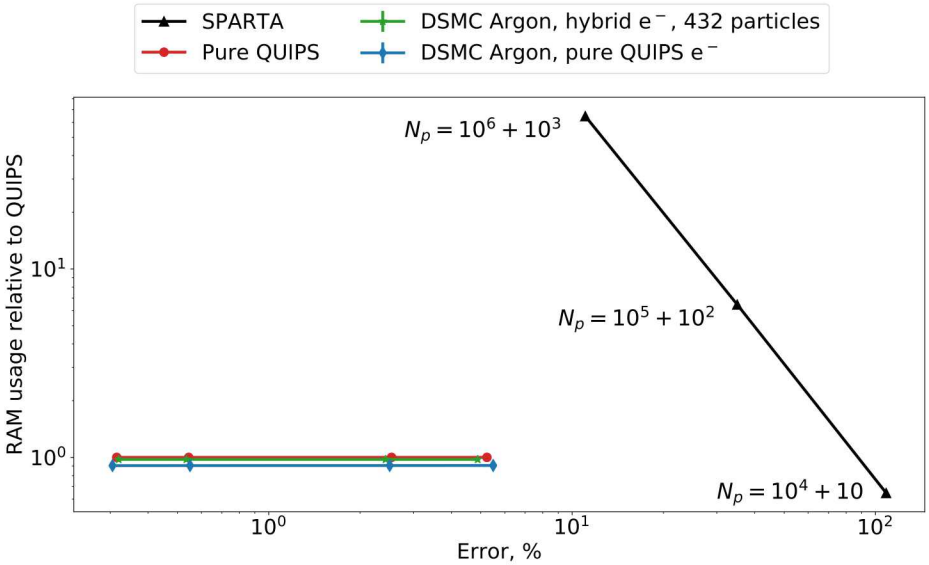

Fig. 8 Memory usage (relative to a pure QUIPS approach) plotted against the relative error in the ionization rate coefficient for an electron temperature of 10eV. Different points on the SPARTA curve correspond to different numbers of particles used (first number is the number of argon particles, second is the number of electron particles); different points on the curves for the pure QUIPS and hybrid codes correspond to different values of C_{RMS} for the electron-argon collisions.

At higher values of the electron temperature, the fixed-size DSMC region contains less and less mass of the electron velocity distribution function, and thus, the hybrid code loses some of its efficiency due to the associated overhead becoming more significant than any speed-ups gained by the use of DSMC collision procedures. Also, at high values of

C_{RMS}, the high level of noise in the collision scheme leads to an increased amount of negative mass particles being created, which cause issues in the merging procedure when the number of merging cells is larger — due to the smaller cell size, the probability of a merge in a cell failing due to negative mass and/or energy increases, which causes an increase in the number of merging adaptations to be performed and leads to an increased computational effort.

Thus, to summarize, the two versions of the hybrid code (DSMC argon/pure QUIPS electrons and DSMC argon/hybrid electrons) provide an improvement of the pure QUIPS approach in terms of efficiency while retaining the same level of accuracy; however, use of the hybrid representation for the electron velocity distribution function requires adapting the choice of DSMC particles/merging cells used to the electron temperature in order to achieve the best speed-up.

It should be noted that on Fig. 6(a), for the case of the DSMC/hybrid code with 16 particles used to represent the bulk of the electron distribution function, a "kink" is observed in the curve at the highest C_{RMS} value. Figure 7 shows the actual rates computed at this value of C_{RMS} — it can be seen that while the error compared to the pure QUIPS code is lower, the actual mean value of the ionization rate coefficient underestimates it by a factor of approximately 2. This is the deficiency of the metric used to estimate the accuracy of the ionization rate coefficient evaluation — even if the ionization rate is severely underpredicted by several orders of magnitude, the error will be less than 100%. Therefore, while the error for this particular hybridization option at this high C_{RMS} value is smaller than the error computed at a lower C_{RMS} value, the actual mean value of the rate coefficient is underestimated; whereas with decreasing the C_{RMS} the correct mean is obtained.

Finally, Fig. 8 shows the memory required by the various methods (relative to the memory requirements of a pure QUIPS approach) plotted against the error in the ionization rate coefficient for an electron temperature of 10 eV. As discussed previously, smaller values C_{RMS} do not incur an additional memory cost; usage of the DSMC representation for argon atoms reduces the overall memory cost, and use of the hybrid scheme for the electron atoms does not require significant amounts of additional memory compared to using a pure QUIPS approach for the electron velocity distribution function. Further memory savings could be achieved by using coarser stretched grids for the electron velocity distribution function, fitted to the area in velocity space which contributes the most to the ionization rate coefficient (so as to retain the level of accuracy in its computation).

D. Couette flow

Finally, a 1-D Couette supersonic flow of argon was considered. The size of the domain was taken to be 0.5 mm, the temperature of the walls to be 300 K, the left wall was assumed to be stationary and the right wall had a y-velocity equal to 1000 m/s, which corresponds to a Mach number of 3.1. The number density of the flow was taken to be $5 \cdot 10^{22}$ m^{-3} (corresponding to a Knudsen number of approximately 0.1). For spatial discretization, 50 cells were used. A timestep of $2.59 \cdot 10^{-9}$ s was used for all the simulations. To produce smooth results, the flow quantities were time-averaged after having reached steady state: after 8000 timesteps, averaging was performed for a further 92000 timesteps. The DSMC region was taken to be a cubic region extending from $\bar{\mathbf{u}}_i - 1.375\eta_r$ to $\bar{\mathbf{u}}_i + 1.375\eta_r$ in each Cartesian direction in velocity space, where $\bar{\mathbf{u}}_i$ is the i-th component of the instantaneous macroscopic velocity in the cell; thus the DSMC region is adaptive and is always centered at the bulk of the distribution function.

Lower values of the number density were also considered (Knudsen number of approximately 1), but due to the lower number of collisions, the velocity-space hybridization scheme did not provide any noticeable computational benefit for this test-case, compared to a pure QUIPS approach.

Figure 9 shows the difference between the computed y-velocity profiles and a linear profile given by $v_y(x) = 1000x/0.005$ as a function of the channel length. It can be seen that all the considered approaches reproduce the velocity slip at the walls and give similar results; however, the velocity profile obtained with the hybrid codes differs slightly from those obtained with pure QUIPS and SPARTA. One possible reason for this discrepancy is the non-conservation of the center of mass in a cell in physical space during the collision step: the mass in $\mathcal{Q}$ is treated as being in the center of a spatial cell, while the particles in $\mathcal{F}$ have continuous-valued spatial positions associated with them. If during a collision mass is transferred from $\mathcal{F}$ to $\mathcal{Q}$, the center of mass in the cell changes. The influence of this effect and possible solutions will be explored in future work.

Figure 10 shows the temperature profiles obtained with the various codes as a function of the channel length. It can be seen that the pure QUIPS code produces a slight over-estimation of the temperature, which might be due to a relatively coarse grid used in the present study. However somewhat better agreement can be seen between the temperature profiles obtained with SPARTA if the hybrid code is used (in which DSMC particles represent the bulk of the distribution).

Finally, as in the previous test cases, we compare the efficiencies of the various codes by considering the computational

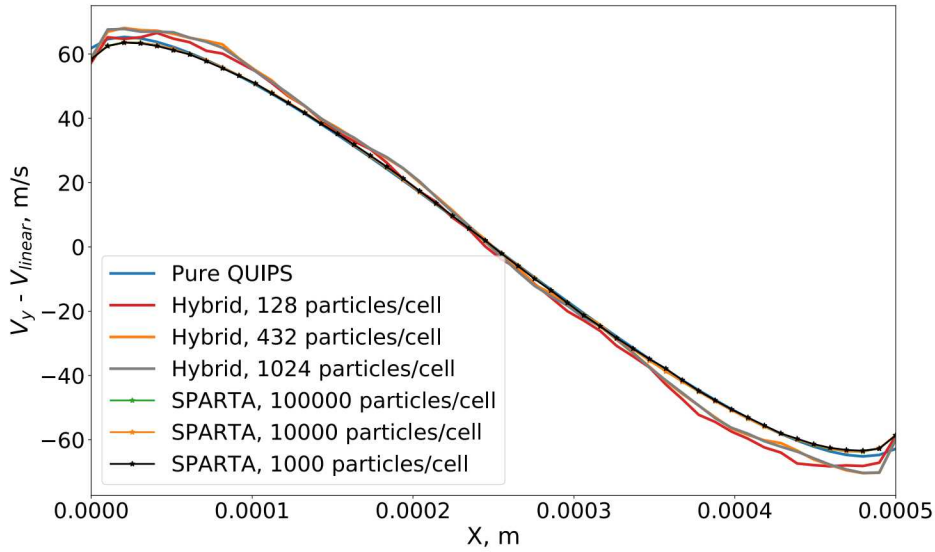

Fig. 9 Difference between the computed profiles of v_y and a linear profile for the Couette flow.

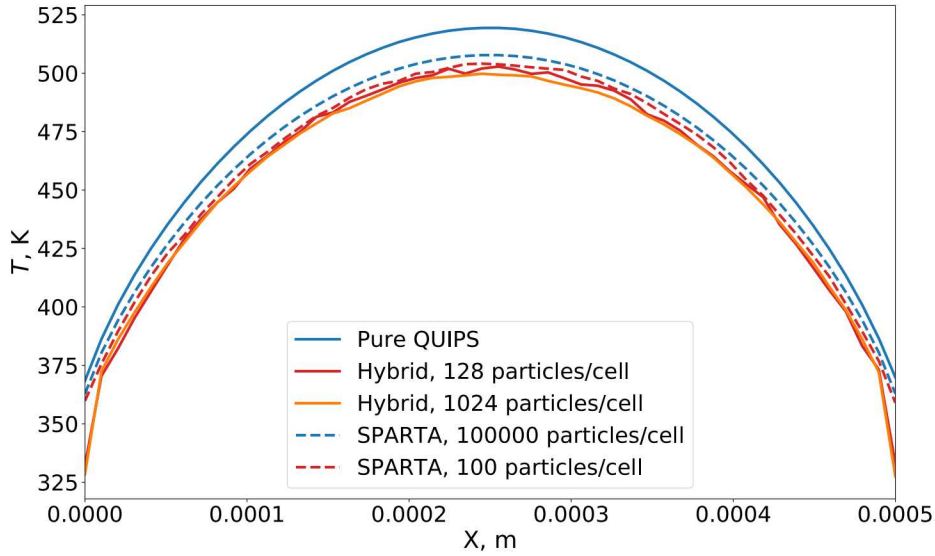

Fig. 10 Temperature profiles for the Couette flow.

time per timestep (which in this case includes the collisions, merging, convection and particle sorting routines) vs the root-mean-squared error in various flow quantities, namely, the temperature and the scaled eighth moment. The pure QUIPS and the hybrid code are run with different values of the C_{RMS} noise parameter, ranging from $1 \cdot 10^{-3}$ to $5 \cdot 10^{-3}$.

Figure 11 shows the computational time per step plotted against the root-mean-squared error of the translation temperature (averaged over the domain). We see that for this low-order moment, the hybridization scheme has little benefit — the error is not affected significantly by the value of the C_{RMS} noise parameter (since the high-velocity tails contribute little to the flow temperature), and is mainly influenced by the number of DSMC particles used. The QUIPS code, with its quieter collision scheme, is able to achieve lower error than the hybrid code even with a relatively large number of particles used. Compared to a pure DSMC code with a similar number of particles, the hybrid approach is somewhat noisier, which might be due to the additional error introduced by the merging procedure.

However, if one considers computational time per step versus error in the scaled 8^{th} moment, which is affected more significantly by the noise in the high-velocity tails, we see that the hybrid scheme provides noise values similar to those

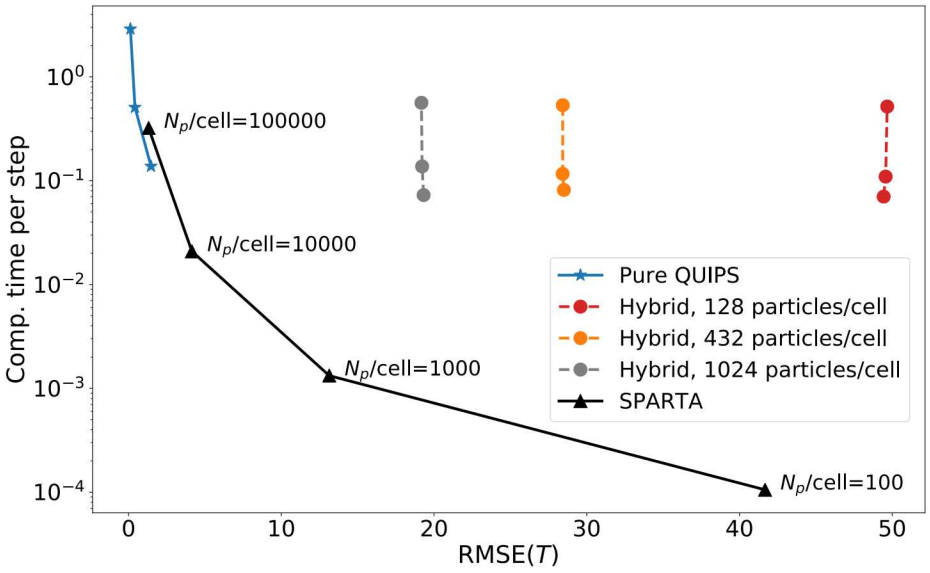

Fig. 11 Computational time per timestep plotted against RMSE (Root-mean-squared error) of the translational temperature. Different points on the SPARTA curve correspond to different numbers of particles per cell used (annotated), different points on the pure QUIPS and hybrid curves correspond to different values of C_{RMS}.

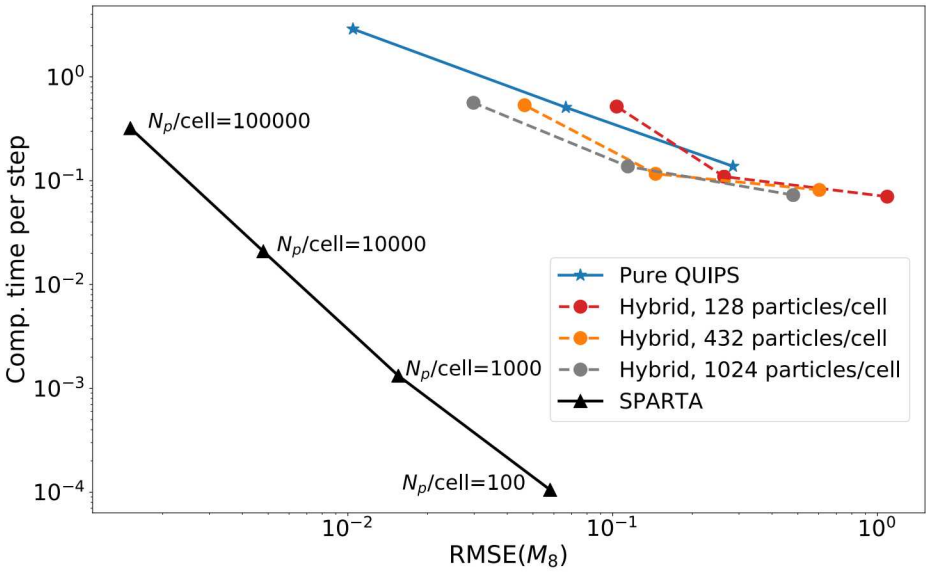

Fig. 12 Computational time per timestep plotted against RMSE (Root-mean-squared error) of the scaled 8^{th} moment. Different points on the SPARTA curve correspond to different numbers of particles per cell used (annotated), different points on the pure QUIPS and hybrid curves correspond to different values of C_{RMS}.

given by the pure QUIPS approach, but at a reduced computational cost. The difference between the results obtained with SPARTA, which is less noisy for this test case, may be due to an insufficiently well-chosen velocity grid (different velocity grid extents and spacings were not considered in the present study). Moreover, such a single-species supersonic flow with no real influence of trace populations is a good area for application of standard fixed-weight DSMC; however, as the ionization test-case shows, for more complicated flows, benefits of applying a discrete velocity-based or a hybrid approach may be very significant.

V. Conclusion

An approach for modelling rarefied gas flows that combines the discrete velocity and DSMC methods in velocity space has been developed and applied to several model problems, such as the BKW relaxation problem, analysis of noise in simulating a stationary Maxwellian distribution, ionization rate coefficient computation and supersonic Couette flow. It has been shown that such a hybrid approach provides better computational efficiency than a pure QUIPS (Quasi-Particle Simulation Discrete Velocity method) approach, while being able to retain accuracy in modelling high-velocity tails of the distribution function. For problems were trace species have a significant impact on the flow physics (such as the considered Ar/e^- mixture), these methods provide both a better computational efficiency and accuracy than standard fixed-weight DSMC, with the ability to improve accuracy without incurring an additional memory cost.

Acknowledgments

This work was supported by Sandia National Laboratories. Sandia National Laboratories is a multimission laboratory managed and operated by National Technology and Engineering Solutions of Sandia, LLC., a wholly owned subsidiary of Honeywell International, Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525. This paper describes objective technical results and analysis. Any subjective views or opinions that might be expressed in the paper do not necessarily represent the views of the U.S. Department of Energy or the United States Government.

References

[1] Bird, G. A., *Molecular Gas Dynamics and the Direct Simulation of Gas Flows*, Clarendon, Oxford, England, UK, 1994.

[2] Raizer, Y. P., *Gas Discharge Physics*, Springer-Verlag Berlin Heidelberg, 1991.

[3] Bird, G., "Transition regime behavior of supersonic beam skimmers," *Phys. Fluids*, Vol. 19, No. 10, 1976, pp. 1486–1491.

[4] Clarke, P. B., "A discrete velocity method for the Boltzmann equation with internal energy and stochastic variance reduction," Ph.D. thesis, University of Texas at Austin, 2015.

[5] Tan, Z., and Varghese, P. L., "The Δ-ε method for the Boltzmann equation," *J. Comput. Phys.*, Vol. 110, No. 2, 1994, pp. 327–340.

[6] Varghese, P., "Arbitrary post-collision velocities in a discrete velocity scheme for the Boltzmann equation," *Proc. of the 25th Intern. Symposium on Rarefied Gas Dynamics*, 2007, pp. 225–232.

[7] Tcheremissine, F., "Solution of the Boltzmann kinetic equation for low speed flows," *Transp. Theory Stat. Phys.*, Vol. 37, No. 5-7, 2008, pp. 564–575.

[8] Morris, A., Varghese, P., and Goldstein, D., "Improvement of a Discrete Velocity Boltzmann Equation Solver With Arbitrary Post-Collision Velocities," *AIP Conf. Proc*, Vol. 1084, AIP, 2008, pp. 458–463.

[9] Dodulad, O., and Tcheremissine, F., "Multipoint conservative projection method for computing the Boltzmann collision integral for gas mixtures," *AIP Conf. Proc*, Vol. 1501, AIP, 2012, pp. 302–309.

[10] Clarke, P., Varghese, P., Goldstein, D., Morris, A., Bauman, P., and Hegermiller, D., "A novel discrete velocity method for solving the Boltzmann equation including internal energy and non-uniform grids in velocity space," *AIP Conf. Proc.*, Vol. 1501, AIP, 2012, pp. 373–380.

[11] Poondla, Y., Varghese, P., Goldstein, D., and Higdon, K., "Modeling of Chemical Reactions Using Quasi-Particle Simulation (QuiPS)," *AIP Conf. Proc.*, AIP Publishing, accepted for publication, 2019.

[12] Morris, A., Varghese, P., and Goldstein, D., "Monte Carlo solution of the Boltzmann equation via a discrete velocity model," *J. Comput. Phys.*, Vol. 230, No. 4, 2011, pp. 1265–1280.

[13] Clarke, P., Varghese, P., and Goldstein, D., "Discrete velocity computations with stochastic variance reduction of the Boltzmann equation for gas mixtures," *AIP Conf. Proc*, Vol. 1628, AIP, 2014, pp. 1032–1039.

[14] Cheremisin, F., "Solving the Boltzmann equation in the case of passing to the hydrodynamic flow regime," *Doklady Physics*, Vol. 45, Springer, 2000, pp. 401–404.

[15] Baker, L. L., and Hadjiconstantinou, N. G., "Variance reduction for Monte Carlo solutions of the Boltzmann equation," *Phys. Fluids*, Vol. 17, No. 5, 2005, p. 051703.

[16] Pan, T.-J., and Stephani, K. A., "Investigation of velocity-space coupling approach in DSMC for tail-driven processes," *AIP Conf. Proc.*, Vol. 1786, AIP Publishing, 2016, p. 050017.

[17] Pan, T.-J., and Stephani, K. A., "Investigation of a Coupling Approach of DSMC and DG Methods for Tail-driven Processes," *47th AIAA Thermophysics Conference*, 2017, p. 4023.

[18] Oblapenko, G., Goldstein, D., Varghese, P., and Moore, C., "A velocity space hybridization-based Boltzmann equation solver," *submitted to J. Comput. Phys.*, ????

[19] Schmidt, D. P., and Rutland, C., "A new droplet collision algorithm," *J. Comput. Phys.*, Vol. 164, No. 1, 2000, pp. 62–80.

[20] Lapenta, G., and Brackbill, J. U., "Dynamic and selective control of the number of particles in kinetic plasma simulations," *J. Comput. Phys.*, Vol. 115, No. 1, 1994, pp. 213–227.

[21] Rjasanow, S., Schreiber, T., and Wagner, W., "Reduction of the number of particles in the stochastic weighted particle method for the Boltzmann equation," *J. Comput. Phys.*, Vol. 145, No. 1, 1998, pp. 382–405.

[22] Plimpton, S., and Gallis, M., "SPARTA direct simulation Monte Carlo (DSMC) simulator," *Sandia National Laboratories, USA, see http://sparta. sandia. gov*, 2015.

[23] Mieussens, L., "Discrete velocity model and implicit scheme for the BGK equation of rarefied gas dynamics," *Math. Models Methods Appl. Sci.*, Vol. 10, No. 08, 2000, pp. 1121–1149.

[24] Bobylev, A. V., "One class of invariant solutions of the Boltzmann equation," *Akademiia Nauk SSSR Doklady*, Vol. 231, 1976, pp. 571–574.

[25] Krook, M., and Wu, T. T., "Exact solutions of the Boltzmann equation," *Phys. Fluids*, Vol. 20, No. 10, 1977, pp. 1589–1595.

[26] Clarke, P., Varghese, P., and Goldstein, D., "A low noise discrete velocity method for the Boltzmann equation with quantized rotational and vibrational energy," *J. Comput. Phys.*, Vol. 352, 2018, pp. 326–340.

[27] Lieberman, M. A., and Lichtenberg, A. J., *Principles of plasma discharges and materials processing*, John Wiley & Sons, 2005.