

A velocity space hybridization-based Boltzmann equation solver

G. Oblapenko^{a,*}, D. Goldstein^b, P. Varghese^{a,b}, C. Moore^c

^a*Oden Institute for Computational Engineering and Sciences, The University of Texas at Austin, 2201 E 24th St, Stop C0200, Austin, TX 78712*

^b*ASE-EM Department, The University of Texas at Austin, 2617 Wichita St., Stop C0600, Austin, TX 78712*

^c*Sandia National Laboratories, Albuquerque, NM, USA*

Abstract

In the present work, a new method for simulation of rarefied gas flows is proposed, a velocity-space hybrid of both a DSMC representation of particles and a discrete velocity quasi-particle representation of the distribution function. The hybridization scheme is discussed in detail, and is numerically verified for two test-cases: the BKW relaxation problem and a stationary Maxwellian distribution. It is demonstrated that such a velocity-space hybridization can provide computational benefits when compared to a pure discrete velocity method or pure DSMC approach, while retaining some of the more attractive properties of discrete velocity methods. Further possible improvements to the velocity-space hybrid approach are discussed.

Keywords: Rarefied gas dynamics, Direct Simulation Monte Carlo, Boltzmann equation, discrete velocity method

1. Introduction

Perhaps the most widely adopted approach to simulation of rarefied gas flows is the Direct Simulation Monte Carlo (DSMC) method developed by

*Corresponding author

Email addresses: georgii.oblapenko@austin.utexas.edu (G. Oblapenko), david@ices.utexas.edu (D. Goldstein), varghese@mail.utexas.edu (P. Varghese), chmoore@sandia.gov (C. Moore)

Preprint submitted to Journal of Computational Physics

January 23, 2020

4 G.A. Bird [1, 2]; however, it suffers from statistical noise and requires sig-
5 nificant computational and memory resources to accurately simulate trace
6 populations, such as those of the tails of a distribution function, of high-lying
7 internal energy states, or of trace species present in the mixture. As such,
8 DSMC can have problems when the trace species are either the main species
9 of interest (for example, heavy molecules in a light carrier gas in substrate
10 deposition processes [3]); when low populations can have a large influence on
11 the flow, for example, high-velocity tails in dissociation or ionization reac-
12 tions; in a beam skimmer set-up [4]; or in the case of a high-velocity beam
13 colliding with a cold gas [5].

14 Several modifications to DSMC have been proposed to improve the res-
15 olution of trace species [3, 6] and reduce the statistical variance of the
16 method [7, 8, 9]. However, the developed weighting approaches improve
17 the accuracy of modeling of trace species only, and do not improve the reso-
18 lution of the stochastic representation of the non-trace species' distribution
19 function.

20 The discrete velocity family of methods has been developed as an alter-
21 native to DSMC [10, 11, 12], and various interpolation schemes are used
22 in order to conserve mass, momentum and energy during the collision pro-
23 cess [12, 13, 14, 15, 16]. An overview of discrete velocity methods can be
24 found in [17] and [5]. Discrete velocity methods are capable of accurate rep-
25 resentation of weakly-populated tails of the distribution function, and have
26 been successfully applied to high-fidelity modeling of populations of high-
27 lying internal energy states [18, 19].

28 Lattice Boltzmann Equation methods [20, 21, 22], while also restricting
29 the set of possible velocities, differ from the above discrete velocity methods
30 in their use of the BGK linear collision operator [23], whereas discrete velocity
31 methods referenced in the previous paragraph can model the full non-linear
32 collision operator rather than a simplified BGK operator.

33 Discrete velocity methods can also benefit greatly from variance reduc-
34 tion techniques [24, 25], which not only decrease the noise in the distribution
35 function, but also provide a significant speed-up of the simulations. These
36 techniques were originally developed in the context of removing stiffness in
37 explicit time integration of discrete velocity models in the limit of a low
38 Knudsen number [26] and later adapted to DSMC simulations [7]. Outside
39 of DSMC and discrete velocity methods, variance reduction has been ap-
40 plied to Particle-in-Cell methods [27, 28], as well as to Fokker-Planck-type
41 equations [29, 30, 31]. Variance reduction may also be applied to flows of

42 gases with internal degrees of freedom [5]. Other notable features of the dis-
43 crete velocity method include the ability to adapt resolution in velocity space
44 (and have independent velocity grids for different species) [18], adapt the
45 noise level (without incurring additional memory cost) [15], accurately model
46 chemical reactions even when reaction probabilities are extremely low [32],
47 and determine the complete distribution function (including internal states)
48 even over many orders of magnitude of variability.

49 However, the discrete velocity method has its own set of drawbacks:
50 it requires a numerical scheme for the convection step (compared to the
51 straightforward convection step in DSMC); treatment of boundary condi-
52 tions may require ghost cells and calculation of number density fluxes; and
53 velocity space discretization can be a source of “ripples” in the values of
54 computed macroscopic quantities in case of discontinuous boundary condi-
55 tions [33, 34, 35, 36, 37]. Also, discrete velocity methods have not yet incor-
56 porated the vast range of physical models and capabilities available to DSMC
57 (such as state-specific inelastic cross-sections, including those of chemical re-
58 actions [38, 39], ionization [40] and radiation processes [41, 42]; coupling with
59 Particle-in-Cell (PIC) codes [43, 44]; the possibility of coupling with quasi-
60 classical trajectory simulations [45]), and in general have seen less application
61 to large-scale simulations.

62 Research in the area of hybrid methods involving the Boltzmann equa-
63 tion has been mostly focused on the hybridization of DSMC and CFD solvers,
64 using continuum methods in flow regions with a low Knudsen number and
65 utilizing DSMC to accurately model the rarefied parts of the flow; examples
66 of such approaches can be found in [46, 47, 48, 49]. While this hybrid ap-
67 proach allows one to save computational effort (compared to a pure DSMC
68 computation of the same flow), the interface between the DSMC and CFD
69 regions, and the noise associated with DSMC (which is a source of noise in
70 the macroscopic boundary conditions for the continuum region) remains an
71 outstanding issue. It is worth mentioning that the lower noise of the dis-
72 crete velocity method might make it more amenable to such a physical space
73 hybridization with continuum solver [50], especially when the trace species
74 and/or tails of the velocity distribution function are important to the flow
75 dynamics as can be the case in atmospheric pressure gas discharges [51].

76 However, to the authors’ knowledge, few works have attempted a hy-
77 bridization in velocity space. In [52], an approach based on solving Euler
78 equations in a region of velocity space (assuming a Maxwellian distribution)
79 and a BGK model equation outside of that “equilibrium” region was devel-

80 oped. In [53] several approaches based on discrete velocity models applied to
 81 the BGK model equation were considered, including a velocity-space hybrid
 82 approach, in which the tails of the distribution function were represented by
 83 particles with arbitrary velocities, while the bulk of the distribution func-
 84 tion was represented by particles on a discrete velocity grid. However, such
 85 a hybridization loses some of the more attractive properties of the discrete
 86 velocity methods — namely, the ability to accurately simulate the tails of
 87 the distribution function. In [54], the DSMC method was used to represent
 88 the bulk of the distribution function, and a discrete Galerkin approach was
 89 used to model the the tails of the distribution; however, no extensive studies
 90 of the errors and computational efficiency of such a scheme were undertaken.
 91 In [55], a similar approach was investigated, but the tails of the distribution
 92 function were modeled using the BGK collision operator, rather than the full
 93 Boltzmann equation.

94 Therefore, it is of interest to develop and study a computational scheme
 95 that is hybrid in its velocity-space representation and does not rely on the
 96 assumptions of a near-equilibrium distribution function and the BGK col-
 97 lision operator: the majority of the mass in a cell in physical space being
 98 represented by a few large DSMC-type particles (whose velocities are not
 99 restricted to a discrete grid), while the tails of the distribution function are
 100 simulated using the discrete velocity method. The expected benefits of such
 101 an approach are an improvement in computational efficiency compared to a
 102 pure discrete velocity method, while retaining the flexibility of DSMC, and
 103 ability to accurately resolve the tails of the distribution functions, as well as
 104 the populations of trace species.

105 The paper is structured as follows: In Sect. 2, a brief overview of our
 106 present discrete velocity method is given for a single-component gas with no
 107 internal degrees of freedom. Then, in Sect. 3, the approach to hybridization
 108 in velocity space is described in detail. After that, the hybrid method is
 109 applied to two spatially homogeneous problems, results of numerical simula-
 110 tions are presented, and various approaches to modeling of rarefied gas flows
 111 are compared.

112 **2. The Boltzmann equation and the discrete velocity method**

113 Let $\phi(\mathbf{r}, \boldsymbol{\eta}, t)$ be the density of the mathematical expectation of the num-
 114 ber of particles in an element of phase space $(\mathbf{r}, \mathbf{r} + d\mathbf{r})$, $(\boldsymbol{\eta}, \boldsymbol{\eta} + d\boldsymbol{\eta})$ at time
 115 t .

116 The definitions of macroscopic quantities such as the bulk velocity \mathbf{u} ,
 117 number density n and temperature T through the function ϕ are well-known
 118 and not given here, the reader is referred to works such as [56, 57, 2]).

119 Defining a reference temperature T_r , mass m_r , number density n_r and a
 120 cross-section σ_r , and a characteristic length L , one can define the reference
 121 velocity as

$$\eta_r = \sqrt{\frac{2kT_r}{m_r}}, \quad (1)$$

122 where k is Boltzmann's constant. One can also introduce a reference mean
 123 free path λ_r :

$$\lambda_r = \frac{1}{\sigma_r n_r}, \quad (2)$$

124 and define a Knudsen number as $Kn = \frac{\lambda_r}{L}$.

125 Now we can introduce dimensionless quantities:

$$\hat{\boldsymbol{\eta}} = \boldsymbol{\eta}/\eta_r; \hat{t} = t\eta_r/L; \hat{\sigma}_t = \sigma_t/\sigma_r; \hat{\mathbf{r}} = \frac{\mathbf{r}}{L}; \bar{\phi} = \phi/n_r, \quad (3)$$

126 where σ_t is the total collision cross-section, integrated over all scattering
 127 angles.

128 One can then write the Boltzmann equation in a scaled form [24]:

$$\frac{\partial \bar{\phi}}{\partial \hat{t}} + \hat{\boldsymbol{\eta}} \cdot \nabla_{\hat{\mathbf{r}}} \bar{\phi} = \frac{1}{Kn} \int_{\hat{\boldsymbol{\zeta}}} \left[\bar{\phi}(\hat{\boldsymbol{\eta}}') \bar{\phi}(\hat{\boldsymbol{\zeta}}') - \bar{\phi}(\hat{\boldsymbol{\eta}}) \bar{\phi}(\hat{\boldsymbol{\zeta}}) \right] \hat{g} \hat{\sigma}_t d\hat{\boldsymbol{\zeta}}, \quad (4)$$

129 where $\hat{\boldsymbol{\zeta}}$ is the scaled velocity of the colliding particle, $\hat{g} = \left| \hat{\boldsymbol{\zeta}} - \hat{\boldsymbol{\eta}} \right|$ is the
 130 magnitude of the relative collisional velocity, and primed variables denote
 131 the post-collisional quantities.

132 Equation (4) can be discretized by defining the velocity distribution func-
 133 tion $\bar{\phi}$ on a finite set of points in velocity space. For a uniform discrete
 134 velocity grid, it takes on the following form [24]:

$$\frac{\partial \bar{\phi}}{\partial \hat{t}} + \hat{\boldsymbol{\eta}} \cdot \nabla_{\hat{\mathbf{r}}} \bar{\phi} = \frac{\beta^3}{Kn} \sum_{\hat{\boldsymbol{\zeta}} \neq \hat{\boldsymbol{\eta}}} \left[\bar{\phi}(\hat{\boldsymbol{\eta}}') \bar{\phi}(\hat{\boldsymbol{\zeta}}') - \bar{\phi}(\hat{\boldsymbol{\eta}}) \bar{\phi}(\hat{\boldsymbol{\zeta}}) \right] \hat{g} \hat{\sigma}_t. \quad (5)$$

135 Here $\beta = \Delta\eta/\eta_r$ is the scaled step size in velocity space. We can define
 136 $\hat{\phi} = \beta^3 \bar{\phi}$, with $\hat{\phi}(\hat{\boldsymbol{\zeta}})$ being the density of the mathematical expectation of the

137 number of particles in a volume of size β^3 centered around the velocity $\hat{\zeta}$.
 138 Equation (5) then takes on the following form:

$$\frac{\partial \hat{\phi}}{\partial \hat{t}} + \hat{\eta} \cdot \nabla_{\hat{\mathbf{r}}} \hat{\phi} = \frac{1}{Kn} \sum_{\hat{\zeta} \neq \hat{\eta}} \left[\hat{\phi}(\hat{\eta}') \hat{\phi}(\hat{\zeta}') - \hat{\phi}(\hat{\eta}) \hat{\phi}(\hat{\zeta}) \right] \hat{g} \hat{\sigma}_t. \quad (6)$$

139 While the collision operator can be evaluated in a deterministic fashion,
 140 this requires $\mathcal{O}(N_v^2)$ operations, where N_v is the number of points on the
 141 discrete velocity grid, and thus is very expensive from a computational view.
 142 However, the collision operator can be computed using a stochastic Monte-
 143 Carlo scheme [24, 5], which is the method used in the present work. The
 144 number of collisions to be performed is computed as

$$N_{coll} = nint \left(\frac{\Delta \hat{t} \hat{T}^{2/3}}{Kn C_{RMS}^2 \hat{m}_r^{2/3} \beta^3} \left(\frac{\hat{n} - 2\hat{n}_{neg}}{\hat{n}} \right)^2 \right), \quad (7)$$

145 where $nint$ denotes rounding to the nearest integer, \hat{m}_r is the scaled collision-
 146 reduced mass, \hat{n}_{neg} is the amount of negative mass in the numerical repre-
 147 sentation of the velocity distribution function at that particular location in
 148 physical space (explained below), and C_{RMS} is a parameter that is propor-
 149 tional to the user-defined level of acceptable noise in the distribution function
 150 at equilibrium [5]. Decreasing the value of C_{RMS} leads to lower noise in the
 151 distribution function, at the cost of having to compute more collisions.

152 A collision is computed by selecting a pair of locations on the discrete
 153 velocity grid for depletion. The probability of a location being selected is
 154 proportional to the absolute magnitude of the distribution function at that
 155 point:

$$P(\hat{\eta}) = \frac{|\hat{\phi}(\hat{\eta})|}{\sum_{\hat{\zeta}} |\hat{\phi}(\hat{\zeta})|}. \quad (8)$$

156 Once two velocity locations are chosen for a collision, the density at those
 157 locations is depleted by the quantity

$$\Delta \hat{\phi} = \Delta \hat{t} \frac{(\hat{n} - 2\hat{n}_{neg})^2}{2Kn N_{coll}} \text{sign} \left(\hat{\phi}(\hat{\eta}) \hat{\phi}(\hat{\zeta}) \right) \hat{g} \hat{\sigma}_t. \quad (9)$$

158 The removed mass is then immediately replenished — its velocity is com-
 159 puted using standard collision mechanics [2]; however, since the post-collision
 160 velocity does not necessarily lie on the discrete velocity grid, a remapping

161 scheme is used [13, 15], which conserves mass, momentum and energy. The
 162 remapping scheme utilizes a 7-point stencil; and the scheme delivers negative
 163 mass to 3 of the stencil points in order to yield a net conservation of energy.
 164 This can result in small amounts of net negative mass at a few velocity space
 165 locations, typically near the tails of the distribution function (see [5]). The
 166 presence of negative mass is accounted for in Eqns. (7) and (9).

167 3. Description of the velocity-space hybrid method

168 Let $\mathcal{D} \subset \mathbb{R}^3$ denote the set of discrete velocities used for the computation.
 169 Let us also define a region $\mathcal{F} \subset \mathbb{R}^3$ of velocity space in which particles can
 170 have any velocity (i.e. a particle whose velocity lies in \mathcal{F} is not restricted to
 171 the discrete velocity grid), as well as define a region $\mathcal{Q} = \mathcal{D} \setminus \mathcal{F}$ in which the
 172 distribution function is defined only on the discrete velocity grid.

173 Fig. 1 shows a schematic of such a hybrid representation of the velocity
 174 distribution function (for some fixed value of v_z).

175 The basic velocity-space hybrid algorithm for simulation of rarefied gas
 176 flows consists of the following steps (performed during a single simulation
 177 timestep):

- 178 1. Collide particles in region \mathcal{F} using DSMC collision mechanics, that is,
 179 a No-Time-Counter (NTC) [2] scheme or some other DSMC collision
 180 scheme, such as the Majorant Frequency method [58];
- 181 2. Collide mass in region \mathcal{Q} with mass in \mathcal{Q} and particles in \mathcal{F} using
 182 discrete velocity collision mechanics;
- 183 3. Remap any particles that lie outside region \mathcal{F} to the discrete velocity
 184 grid in \mathcal{Q} ;
- 185 4. Perform merging of particles in region \mathcal{F} (the necessity of this procedure
 186 will be discussed in the next subsections);
- 187 5. Perform convection, using a finite difference or finite volume scheme
 188 for the distribution function defined on the discrete velocity points in
 189 \mathcal{Q} and using Lagrangian convection for the particles in \mathcal{F} .

190 During step 1, a post-collision velocity of a particle may lie in \mathcal{Q} , in
 191 which case it is remapped onto the discrete velocity grid (this is done after
 192 all the collision steps have been performed, so as to allow a DSMC particle
 193 to undergo multiple collisions during a single timestep). Conversely, during
 194 step 2, a post-collision velocity of the replenishment quantity may lie inside
 195 \mathcal{F} , in which case the mass is not remapped onto the discrete velocity grid,




Figure 1: Schematic showing the velocity-space hybrid representation of the velocity distribution function (for a fixed value of the z -velocity v_z). 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 red circles denote the DSMC particles (with their size corresponding to their computational weight, see below for details) in region \mathcal{F} (shown here as the dark grey rectangle).

196 but rather immediately created as a new particle in \mathcal{F} . Finally, during the
 197 remapping procedure, one or more points of the remapping stencil may lie
 198 inside \mathcal{F} , in which case the mass replenished at these points is also created
 199 as new particles which are not restricted to the discrete velocity grid. It is
 200 important to note that all particles in \mathcal{F} have a different weight associated
 201 with them, unlike standard DSMC, where particles usually have identical
 202 weights. Even if one started the simulation with equal-weight particles in \mathcal{F} ,
 203 due to the remapping scheme, particles with different computational weights
 204 will be created in \mathcal{F} .

205 While variable weight DSMC algorithms have been developed that combine
 206 the splitting and merging steps [3, 59], they are generally not conservative,
 207 and utilize various workarounds in order to ensure energy conservation;
 208 in order to avoid this issue, in the present work the collisions of variable
 209 weight particles and the merging step are performed separately in order to
 210 construct a conservative scheme.

211 *3.1. Collision algorithm*

212 As described previously, three types of collisions are possible:

- 213 1. Collisions between particles in region \mathcal{F} ;
- 214 2. Collisions between particles in region \mathcal{F} and mass in region \mathcal{Q} ;
- 215 3. Collisions between mass in region \mathcal{Q} .

216 Since one of the aims of discrete velocity methods is to have a low amount
 217 of noise in the tails of the distribution function, collisions of types 2 and 3 are
 218 treated using the Monte Carlo collision scheme as detailed in [5]. Collisions
 219 between particles in region \mathcal{F} are treated using DSMC collision mechanics,
 220 and are performed first — since collisions of types 2 and 3 might create a large
 221 number of DSMC particles. In the present work, the No-Time-Counter col-
 222 lision method will be utilized [2] (although the Majorant Frequency method
 223 also works well). However, a modified version of the method has to be used
 224 to account for the fact that, in general, particles in \mathcal{F} have different weights.

225 The required number of collisions (for a single-species gas) is computed
 226 as follows [60]:

$$N_{coll,\mathcal{F}} = \frac{1}{2} \frac{N_p(N_p - 1)(F_{num}\sigma g)_{max}\Delta t}{V}, \quad (10)$$

227 where N_p is the number of particles in \mathcal{F} , V is the volume of the cell in
 228 physical space in which the collisions are considered, F_{num} is the number of
 229 real gas particles represented by a single model DSMC particle in \mathcal{F} , and Δt
 230 is the timestep of the simulation. $(F_{num}\sigma g)_{max}$ denotes the maximum value
 231 of the product $F_{num} \cdot \sigma \cdot g$ (taken over the whole region \mathcal{F}).

232 Re-writing (10) in terms of the scaled quantities, we obtain

$$N_{coll,\mathcal{F}} = \frac{1}{2Kn} N_p(N_p - 1)(\hat{n}\hat{\sigma}\hat{g})_{max}\Delta\hat{t}, \quad (11)$$

233 where $(\hat{n}\hat{\sigma}\hat{g})_{max}$ is the maximum value of $\hat{n} \cdot \hat{\sigma} \cdot \hat{g}$ taken over the region \mathcal{F} .

234 When a particle pair is tentatively chosen for a collision, it is accepted
 235 with a probability of

$$P = \frac{\hat{n}_1\hat{\sigma}\hat{g}}{(\hat{n}\hat{\sigma}\hat{g})_{max}}. \quad (12)$$

236 Here \hat{n}_1 denotes the weight (dimensionless number density) of the heavier of
 237 the two colliding particles.

238 When a pair is accepted for a collision, the heavier particle (whose weight
239 is denoted \hat{n}_1) is split into two particles: one with weight \hat{n}_2 (the weight of
240 the lighter collision partner) and one with weight $\hat{n}_1 - \hat{n}_2$. The particles with
241 weights \hat{n}_2 collide and their post-collision velocities are computed based on
242 the cross-section model used in the simulation, while the particle with weight
243 $\hat{n}_1 - \hat{n}_2$ does not have its velocity affected.

244 Due to the fact that the remapping scheme causes negative mass to be cre-
245 ated at the “exterior” stencil points, negative mass particles may be present
246 in region \mathcal{F} . When colliding a negative mass particle from \mathcal{F} with another
247 particle from \mathcal{F} , if one uses the particle splitting scheme as described above,
248 an additional negative mass particle is created as a result of the splitting.
249 Another option is to treat such collisions using the discrete velocity method
250 Monte Carlo collision scheme [24] (which, due to the small depletion quan-
251 tity involved, might not create as much negative mass in the system), and
252 use the the DSMC collision scheme only for collisions between positive mass
253 particles from \mathcal{F} . This is the approach used in the present work.

254 It should be noted that the developed algorithm places no constraints on
255 the size of \mathcal{F} or where it lies in velocity space (e.g., in Fig. 1 it is offset from
256 the center); moreover, for a multi-species flow, different-sized regions \mathcal{F} may
257 be defined for each species. As a limiting case, one can also do species-wise
258 velocity hybridization, treating some chemical species in a flow using the
259 discrete velocity method, and modeling others with DSMC.

260 In the present work, we consider the case where \mathcal{F} contains the bulk of
261 the distribution function. As such, the values of the lower-order moments
262 (number density, velocity, temperature) are governed by the particles in \mathcal{F} ,
263 and therefore, it is expected that the performance of the hybrid scheme in
264 terms of dependence on Knudsen and Mach numbers will be similar to that
265 of standard DSMC. Possible improvements at low Knudsen numbers might
266 be obtained by using a Fokker-Planck-based DSMC scheme [61, 62].

267 *3.2. Merging algorithm*

268 It can be seen that the number of particles in \mathcal{F} will constantly grow, both
269 due to particle splitting and due to creation of new particles as a result of
270 both the remapping scheme and some post-collision replenishment quantities
271 appearing in \mathcal{F} . Additionally, the total amount of negative mass will also
272 keep increasing, which leads to an increase in the computational cost of the
273 collision loop (due to a larger number of collisions as defined by (7)), as well

274 as to an increase in the level of noise in the representation of the distribution
 275 function. Therefore, a strategy for merging particles in \mathcal{F} is also needed.

276 While merging algorithms for performing a 2 : 1 merge have been exten-
 277 sively developed [63, 64, 65, 66], they are non-conservative by their nature,
 278 and employ various techniques to reduce errors and enforce conservation. In
 279 order to preserve mass, momentum and energy, a $N : 2$ merge needs to be
 280 performed (with $N \geq 3$) [67, 68, 69]. Moreover, in order to avoid excessive
 281 thermalization of the distribution function, particles should be merged to-
 282 gether based on their proximity in velocity space [69]. It should be noted
 283 that performing a $N : 3$ or $N : 4$ merge, while providing the possibility of
 284 conserving of all second-order moments (and some third-order moments in
 285 the case of the $N : 4$ merge), is complicated by having to solve a system
 286 of non-linear equations numerically [70], and thus, such approaches are not
 287 investigated in the present paper.




Figure 2: Example of a merging grid for a fixed value of \hat{V}_z ; blue symbols denote the points on the discrete velocity grid, grey dots denote the DSMC particles, black lines denote the boundaries of \mathcal{F} , and the red lines show the merging grid cells — $M = 16$ in this \hat{V}_z plane, $M = 64$ in total. The left subplot shows the pre-merge particle, the right — the post-merge particles (a merge in a velocity-space cell does not necessarily produce particles that lie in the same merging cell, therefore, the number of post-merge particles in a velocity-space merging cell is not necessarily equal to 2).

288 In the present work a uniform grid (unrelated to the discrete velocity grid)
 289 is overlaid over \mathcal{F} , and all particles within a single grid cell are merged into
 290 two particles ($N : 2$ merge). An example of such a grid is shown in Fig. 2.

291 Due to the possibility of a particle having negative mass, the merge within a
292 single such cell might fail, if either one of the following conditions is found to
293 be true: 1) the total mass in the cell is negative or 2) the total energy in the
294 cell is negative. However, the frequency of such events is relatively low, and
295 if a total negative mass and/or energy is accumulated in a merging cell at
296 a time-point, in most cases, during the next timestep, enough positive mass
297 and/or energy will be input into the cell for the merge to be performed.

298 The merging algorithm is linear with respect to the number of particles
299 N_p , and requires $\mathcal{O}(M)$ additional memory, where M is the number of cells
300 in the velocity-space merging grid. We note, however, that such extra mem-
301 ory is used for only one spatial cell at a time and so is trivial for realistic
302 multidimensional problems. [Therefore, the memory required by the hybrid
303 scheme scales linearly with the number of physical dimensions and spatial
304 grid cells, as in standard DVM and DSMC methods. The amount of mem-
305 ory required by the merging procedure is constant and independent of the
306 number of physical grid cells, since collisions and merging are done on a
307 cell-by-cell basis.](#)

308 The algorithm iterates over the particles three times, first computing and
309 storing the number density in each merging grid cell, as well as computing
310 and storing the bulk velocity of the particles in each cell. On the second pass,
311 the kinetic energy (relative to the bulk velocity) in each cell is computed and
312 stored, and after that, the mass and velocity components of the two-post
313 merge particles in each cell are computed.

314 While it would be reasonable to assume that this merging procedure limits
315 the number of particles in \mathcal{F} to $2M$, the number M of cells in the merging
316 grid does not set a strict limit on the number of particles being simulated
317 – as discussed above, due to the presence of negative mass, it is possible
318 that no merging will be performed at all in some merging cells. However, on
319 average, the number of particles is indeed limited to approximately $2M$ (see
320 results of numerical simulations below). Use of adaptive merging algorithms
321 in order to better conserve higher-order moments and avoid issues related to
322 the presence of negative mass will be explored in future work.

323 Fig. 3 presents the developed hybrid approach in a flow chart form for
324 convenience.

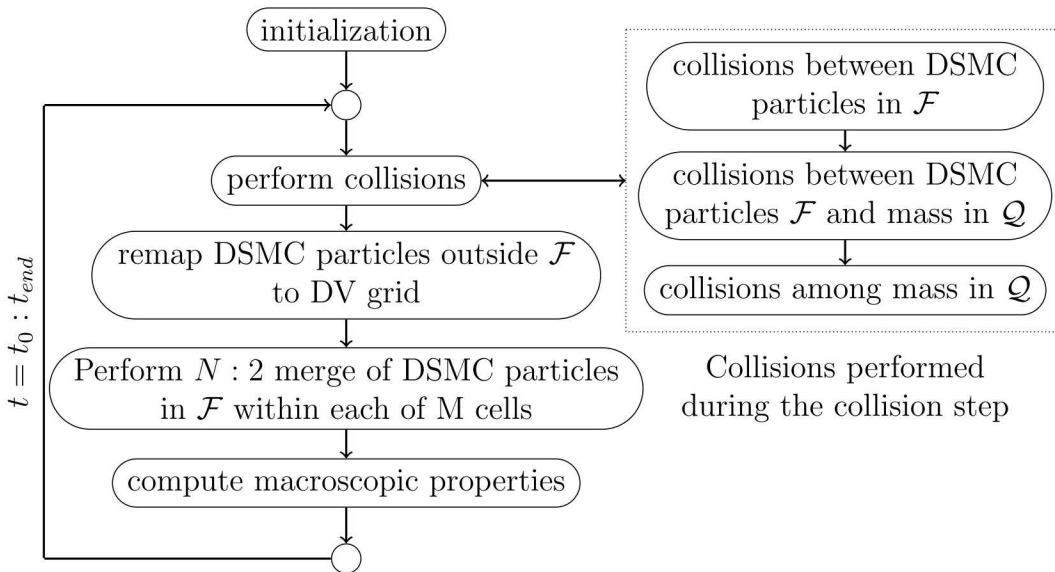


Figure 3: Flow chart of the hybrid velocity-space algorithm, t_0 is the start time of the simulation and t_{end} is the end time.

325 4. Numerical results and discussion

326 In order to verify the velocity-space hybrid method, two spatially homo-
 327 geneous test cases were considered: the BKW relaxation problem [71, 72]
 328 and the case of an initial Maxwellian distribution. For both test cases, the
 329 reference temperature was set to 273 K, the reference number density was
 330 equal to 10^{23} m^{-3} .

331 The extent of the discrete velocity grid was ± 3.5 for both test cases (that
 332 is, the velocity grid extends to ± 3.5 times the mean thermal speed η_r
 333 in each Cartesian direction), and two grids were considered: a “fine” grid with
 334 $29 \times 29 \times 29 = 24,389$ points (with a grid spacing of $\beta = 0.25$) and a “coarse”
 335 grid with $15 \times 15 \times 15 = 3,375$ points (with a grid spacing of $\beta = 0.5$). The
 336 region \mathcal{F} was considered to be a cube centered around the origin in velocity
 337 space. Two different extents of the region \mathcal{F} were considered: ± 1.375 and
 338 ± 1.875 (that is, ± 1.375 and ± 1.875 times the mean thermal speed η_r
 339 in each Cartesian direction). The two different-sized regions \mathcal{F} will be denoted as
 340 \mathcal{F}_{small} and \mathcal{F}_{large} , correspondingly. On a fine grid, when the gas is initialized
 341 with a Maxwellian distribution, approximately 86% of the mass is contained
 342 in \mathcal{F}_{small} , and 98% – in \mathcal{F}_{large} ; these regions occupy 1331 and 3375 points of
 343 the discrete velocity grid, respectively. On a coarse grid, the corresponding

344 values are similar at 80% and 96%; \mathcal{F}_{small} occupies 125 points of the discrete
 345 velocity grid and \mathcal{F}_{large} occupies 343 points.

346 Three different merging grids were used in the hybrid simulations, all
 347 with an equal number of merging cells in each velocity direction: with $M =$
 348 $4 \times 4 \times 4 = 64$ cells (corresponding to ~ 128 post-merge particles), $M =$
 349 $6 \times 6 \times 6 = 216$ cells (corresponding to ~ 432 post-merge particles), and $M =$
 350 $8 \times 8 \times 8 = 512$ cells (corresponding to ~ 1024 post-merge particles).

351 Initialization was performed in the following manner: the initial distribu-
 352 tion function (either the BKW distribution or the Maxwellian distribution)
 353 was evaluated at each grid point in \mathcal{D} , and the value of the distribution
 354 function at each discrete velocity was then normalized by the same factor
 355 so that the total number density was equal to 1. In case of a hybrid code,
 356 all the mass in \mathcal{F} was treated as DSMC-type mass, and a merge was per-
 357 formed at $t = 0$ (otherwise the number of DSMC particles at $t = 0$ will
 358 be equal to the number of velocity grid points inside \mathcal{F} , and can be much
 359 higher than the number of DSMC particles during the following timesteps in
 360 case a fine grid is used). It should be noted that the initialization scheme
 361 used in the simulations only preserves the number density (compared to the
 362 true initial distribution), and leads to some small error in the higher-order
 363 moments; this can be mitigated by using a finer discrete velocity grid. Using
 364 a coarser grid leads to underestimation of the moments, as discussed in [5].
 365 Another approach for the Maxwellian test-case would be to compute the
 366 exact equilibrium discrete distribution, as described in [73, 52], and use it
 367 for initialization, however, for the purposes of this work this is not required.
 368 The problem of how to best represent a smooth analytic probability density
 369 function on a cubic grid for the purpose of simply establishing the initial
 370 conditions for a simulation will not be further examined here.

371 The simulations were run with the following values of C_{RMS} : 7.5×10^{-4} ,
 372 1.5×10^{-3} , 2.5×10^{-3} , 5×10^{-3} , 7.5×10^{-3} , 1×10^{-2} . At higher values of C_{RMS} ,
 373 the discrete velocity method can become unstable to due large amounts of
 374 negative mass in the tails of the distribution function caused by the large
 375 depletion quantity being used during the crudely represented collisions.

376 The test gas considered was argon, and the Variable Hard Sphere (VHS)
 377 cross-section was used for the Maxwellian test case, with a collision diameter
 378 $d_{coll} = 4.17 \text{ \AA}$ and $\omega = 0.81$ (for the BKW relaxation test case, a pseudo-
 379 Maxwell model was used for the collision cross-section). For all the test
 380 cases, a scaled time step equal to 0.1 of the scaled collision time was used:
 381 $\Delta \hat{t} = 0.1 \hat{t}$.

382 4.1. BKW relaxation

383 The BKW relaxation is an analytic time-dependent solution of the Boltz-
 384 mann equation for the case of a pseudo-Maxwell collision cross-section [71,
 385 72]. The analytic expression for the distribution function is given by

$$f(\mathbf{v}, t) = \frac{1}{(\pi C)^{3/2}} \frac{1}{2C} \left(5C - 3 + \frac{2(1-C)v^2}{C} \right) \exp(-v^2/C), \quad (13)$$

386 where $C = 1 - 2/5 \exp(-t/6)$ is defined for all $t \geq 0$.

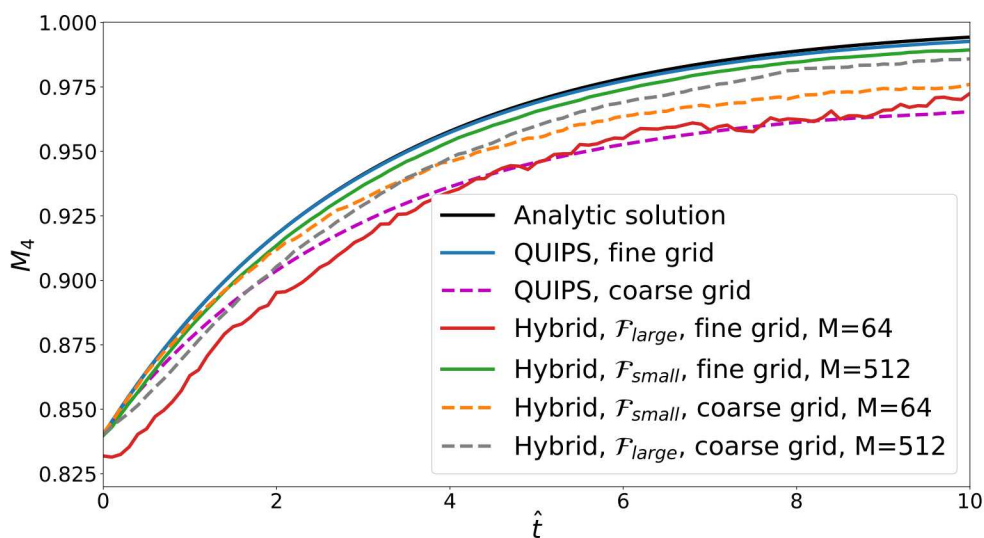


Figure 4: 4th dimensionless moment (averaged over 200 ensembles) as a function of the dimensionless time for different combinations of grids, sizes of hybrid regions and number of merging cells. “QUIPS” denotes the non-hybrid discrete velocity simulations. For all the simulations, $C_{RMS} = 5 \times 10^{-3}$ was used. The 6th and 8th moments show a similar trend and are not presented.

387 Analytic expressions for the moments of the distribution can also be ob-
 388 tained, and take on the following form (moments are scaled so that their
 389 limit as $t \rightarrow \infty$ is equal to 1):

$$\hat{M}_{2l}^{an}(t) = C^l (l - (l - 1)C), \quad (14)$$

390 where \hat{M}_{2l}^{an} is the analytic moment of order $2l$.

391 Figure 4 shows the temporal evolution of the 4th dimensionless moment
 392 as a function of the dimensionless time. The black line denotes the analytic

393 solution as defined by expression (14). We see that in general, increasing the
 394 size of region \mathcal{F} , as well as decreasing the number of merging cells M leads
 395 to an increased amount of noise in the simulation (even with the ensemble
 396 averaging used). It can be seen that moments computed using the hybrid
 397 code relax towards the equilibrium value at approximately the expected rate
 398 for some of the simulation parameters. Using a low number of merging cells
 399 leads to a systematic underestimation of the moments due to 1) a low number
 400 of DSMC particles used, which leads to a negative bias [74, 75] and 2) a bias
 401 introduced by the merging procedure (discussed in detail in Appendix A).

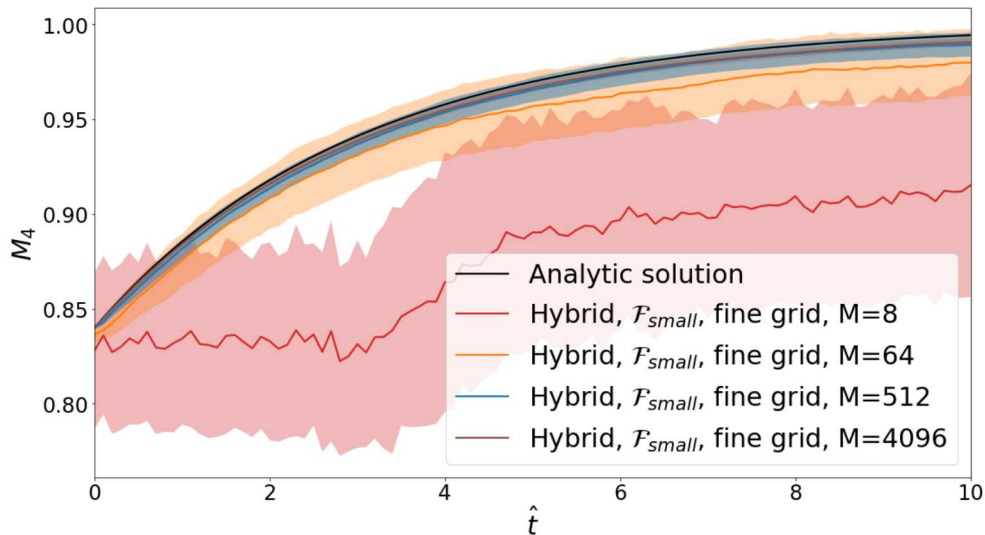


Figure 5: The 4th dimensionless moment computed on a fine grid with \mathcal{F}_{small} for different numbers of merging cells M ; a value of $C_{RMS} = 0.005$ was used. The moments are plotted with the semi-transparent regions showing a range of plus/minus one standard deviation. 200 ensembles were used to compute the means and the standard deviations. The 6th and 8th moments show a similar trend and are not presented.

402 Figure 5 shows convergence of the 4th moment in the number of DSMC
 403 particles used in the simulation (as governed by the number of merging re-
 404 gions M). Using a larger number of DSMC particles leads to a satisfactory
 405 convergence to the expected result. Because of the initialization scheme used
 406 (via a discrete velocity grid), further increase in the number of merging cells
 407 will provide diminishing returns due to the error in the initial distribution.

408 We now pose the question: how can we quantify the relative benefit of
 409 using a hybrid DSMC/discrete velocity approach targeted at improving the

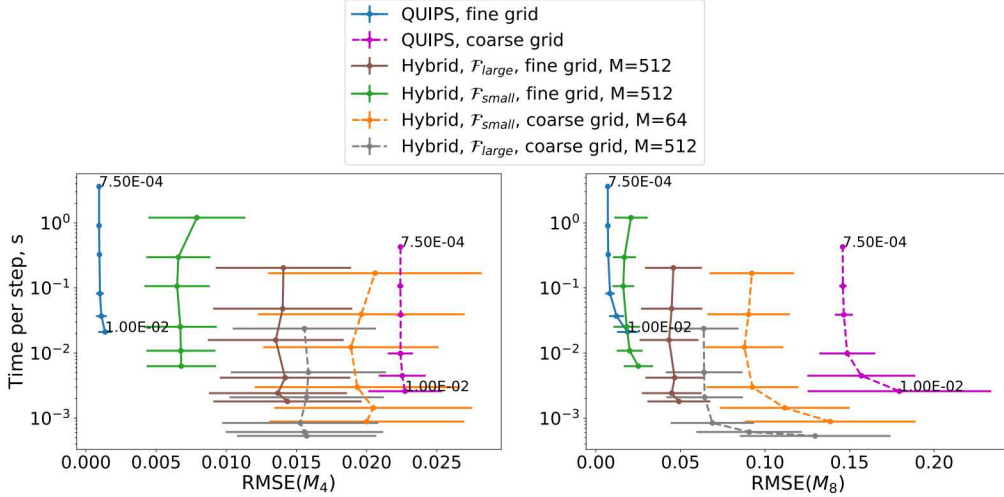


Figure 6: Average time spent on collisions during a single timestep as a function of the RMSE in the 4th and 8th dimensionless moments (when compared to the analytic solution). Different points on the same curve correspond to different values of the C_{RMS} parameter (the pure discrete velocity curves have the lowest and highest values of C_{RMS} annotated).

410 representation of the tails of the distribution function in a non-equilibrium
 411 flow? Let us consider the following metric: computational efficiency vs de-
 412 viation of the moments when compared to the analytic solution. The Root
 413 Mean Square Error (RMSE) of the N^{th} moment, M_N , compared to the ana-
 414 lytic solution is defined as follows:

$$RMSE(\hat{M}_N) = \sqrt{\frac{\sum_{\hat{t}=0}^{10} (\hat{M}_N - \hat{M}_N^{an})^2}{100}}, \quad (15)$$

415 that is, it is computed over the first 100 timesteps (10 collision times). In
 416 equation (15) we consider instantaneous values of the moments, and not
 417 ensemble-averaged quantities. Ensemble averages of the RMSE themselves
 418 yield satisfactorily smooth results. The amount of computational time spent
 419 on calculating the collisions during a single timestep was used as a metric for
 420 the computational efficiency of the considered algorithms.

421 Figure 6 shows the time spent on calculating the collisions plotted against
 422 the error in the moments (compared to the analytic solution). Different
 423 points on each curve correspond to different values of C_{RMS} — lowering the
 424 value of C_{RMS} leads to increased computational cost and lower error in the

425 moments. It is evident that the C_{RMS} parameter does not have a significant
 426 impact on the error in the lower-order moment (the various curves are nearly
 427 vertical and the amount of statistical scatter does not change with varying
 428 C_{RMS} values) for the hybrid code. This is due to the lower influence of
 429 the tails of the distribution function on the lower-order moments. For the
 430 higher-order moment (8^{th} moment), the tails of the distribution function have
 431 a more significant influence, and thus, the error in the moments in the hybrid
 432 simulations is more strongly influenced by the value of C_{RMS} , which mostly
 433 affects the Q region.

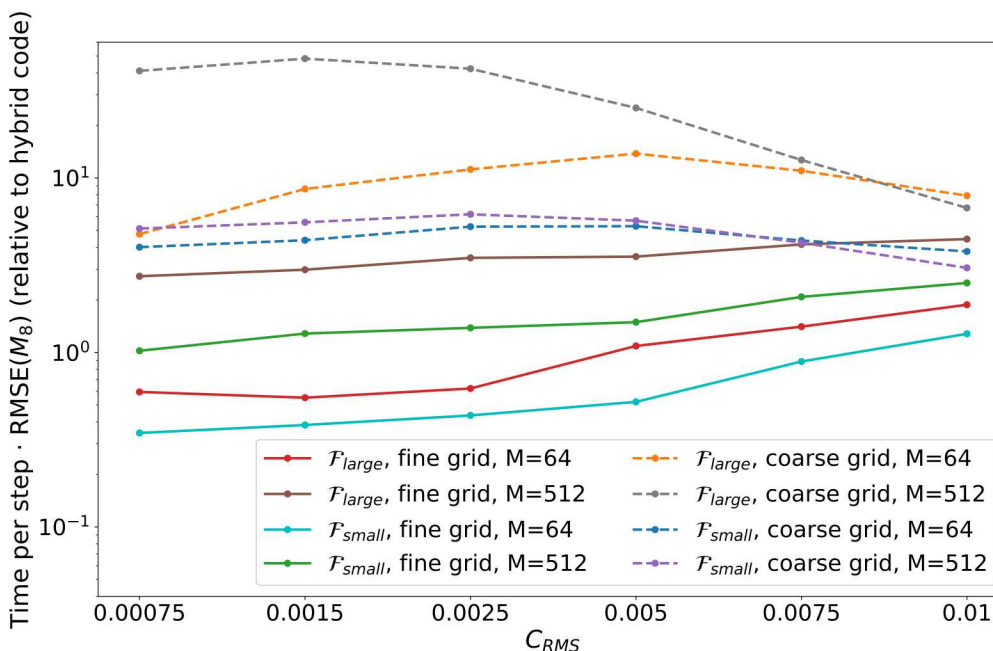


Figure 7: Ratio of average time spent on collisions during a single timestep multiplied RMSE of the 8^{th} moment for the discrete velocity code divided by the same quantity for a hybrid code as a function of C_{RMS} . A ratio of larger than 1 means that the hybrid code is outperforming a pure discrete velocity code in terms of either computational speed, RMSE, or both.

434 To better compare the speed-vs-error performance between the different
 435 codes, we also consider the following ratio:

$$\frac{\left(t_{coll} \cdot RMSE(\hat{M}_8)\right)_{QUIPS}}{t_{coll} \cdot RMSE(\hat{M}_8)}, \quad (16)$$

436 that is, the ratio of the computational time per collision step multiplied by
437 the RMSE of the 8th moment for the discrete velocity code divided by the
438 computational time per collision step multiplied by the RMSE of the 8th
439 moment for a hybrid code (on the same velocity grid). This allows us to
440 simultaneously compare both the noise level and computational efficiency of
441 the various approaches. Fig. 7 shows these ratios as functions of C_{RMS} . In
442 general it can be seen that the hybrid approach can provide an improvement
443 over the purely discrete velocity method, either speed- or error-wise, and this
444 improvement is especially large on a coarse velocity grid. If one only considers
445 computational time, the hybrid code can provide up to a 5-fold speed-up and
446 improved RMSE compared to a pure discrete velocity simulation (see grey
447 and pink curves on Fig. 6). The benefit of the approach of hybridization in
448 velocity-space is especially true for higher-order moments; however, with the
449 considered simulation parameters, the hybrid method cannot achieve as low
450 an error as the discrete velocity method on a fine grid; lowering the error
451 further would require more DSMC particles (more merging cells M).

452 4.2. Maxwellian distribution

453 The second test case considered was the behavior of a gas initialized with
454 a Maxwellian distribution. Due to the stochastic nature of the collision algo-
455 rithm, a certain amount of noise will be present in the numerical simulation;
456 the objective of this test case was to study the influence of various parame-
457 ters on the noise in the tails of the distribution function and to confirm that
458 the hybrid scheme does not force the equilibrium distribution away from
459 Maxwellian. The tails of the distribution function are defined as all the dis-
460 crete velocity points lying outside a cube with extent ± 1.375 , that is, all the
461 points in $\mathcal{Q}_{tail} = \mathcal{D} \setminus \mathcal{F}_{small}$. Since for the larger hybrid region \mathcal{F}_{large} the value
462 of $\hat{\phi}$ is 0 at some of the grid points in \mathcal{Q}_{tail} , in order to perform a consis-
463 tent comparison, during the output step, the particles in \mathcal{F}_{large} are remapped
464 to the discrete velocity grid and the resultant distribution function summed
465 with $\hat{\phi}$ (this is strictly a visualization procedure and the remapping does not
466 affect the actual particles and values of $\hat{\phi}$ used further in the simulation).

467 At each timestep, given the values of $\hat{\phi}_{tail}$ (the set of values of $\hat{\phi}$ in the
468 tails region of velocity space), two metrics are computed, the Root Mean
469 Square Error (RMSE) and Mean Absolute Error (MAE), defined as

$$RMSE(\hat{\phi}_{tail}) = \sqrt{\frac{\sum_{\hat{\boldsymbol{\eta}} \in \mathcal{Q}_{tail}} \left(\hat{\phi}_{tail}(\hat{\boldsymbol{\eta}}) - \hat{\phi}_{tail}^{eq}(\hat{\boldsymbol{\eta}}) \right)^2}{M_{tail}}}, \quad (17)$$

$$MAE(\hat{\phi}_{tail}) = \frac{\sum_{\hat{\boldsymbol{\eta}} \in \mathcal{Q}_{tail}} |\hat{\phi}_{tail}(\hat{\boldsymbol{\eta}}) - \hat{\phi}_{tail}^{eq}(\hat{\boldsymbol{\eta}})|}{M_{tail}}, \quad (18)$$

470 where M_{tail} is the number of points on the velocity grid in region \mathcal{Q}_{tail} ,
 471 and $\hat{\phi}_{tail}^{eq}$ is the equilibrium distribution function of the tails. These metrics
 472 are computed over 200 timesteps and their means calculated, to be used as
 473 measures of noise in the tails of the distribution function.

474 Due to the initialization procedure which does not produce an exact equi-
 475 librium distribution at the initial time $t = 0$ on the truncated discrete grid,
 476 one has to let the simulation code run for a few collision times before be-
 477 ginning to accumulate the errors in the tails, so that the initial coarse rep-
 478 resentation of the Maxwellian distribution can relax to its true equilibrium.
 479 The equilibrium distribution of the tails $\hat{\phi}_{tail}^{eq}$ can be computed as a long
 480 time-average of the tails distribution $\hat{\phi}_{tail}$:

$$\hat{\phi}_{tail}^{eq} = \sum_{\hat{t}=\hat{t}_1}^{\hat{t}_2} \frac{\hat{\phi}_{tail}(\hat{t})}{N_{\hat{t}}}, \quad (19)$$

481 where \hat{t}_1 and \hat{t}_2 are the starting and ending times of the averaging process,
 482 and $N_{\hat{t}}$ is the number of timesteps over which the averaging is performed.

483 In the present work, the computation of the equilibrium distribution function
 484 was begun after 20 collision times (corresponding to 200 timesteps) had
 485 elapsed, and the averaging was performed over 300 timesteps (this number
 486 of timesteps was found to be sufficient).

487 Figure 8 shows the number of particles in region \mathcal{F} as a function of the
 488 dimensionless time for different numbers of merging cells M . We see that the
 489 number of particles does not grow in an unrestricted fashion, but occasional
 490 increases occur due to the presence of negative mass and/or energy in the
 491 merging grid cells.

492 As for the BKW test case, we also want to compare the pure discrete
 493 velocity method with the hybridized approach both in terms of noise and
 494 computational efficiency, and thus also consider the time spent on performing
 495 collisions during a single timestep.

496 To expand the scope of the present work, we also compare to a pure
 497 variable-weight DSMC code based on the hybrid code, as well as the open-
 498 source SPARTA DSMC code [76]. The hybrid code can be run as a pure
 499 DSMC code (with variable-weight particles) if one assumes that $\mathcal{F} = \mathbb{R}^3$;
 500 the initialization step is still based on evaluation of the distribution function

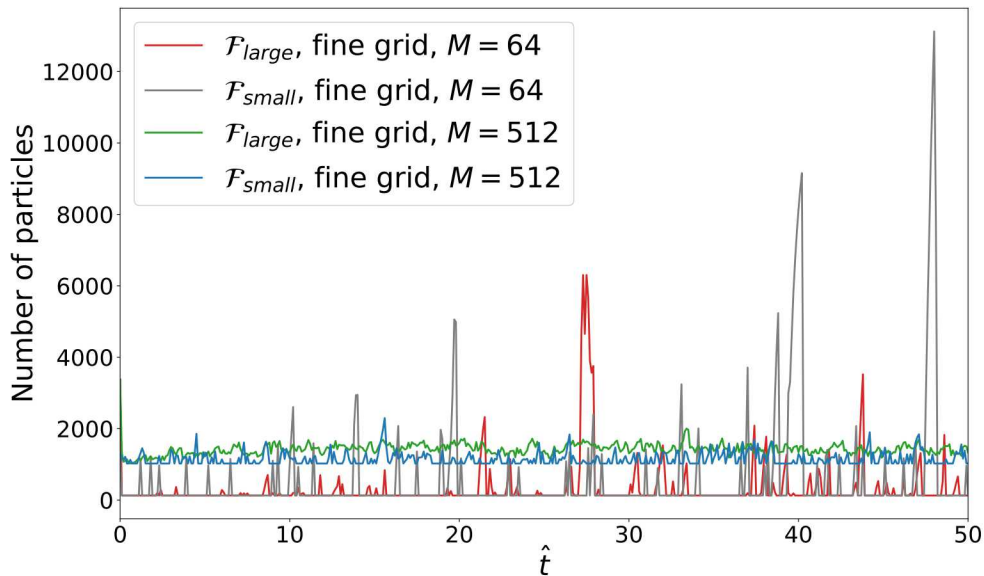


Figure 8: Number of particles in region \mathcal{F} as a function of time for different numbers of merging cells M ; a value of $C_{RMS} = 0.005$ was used.

501 on a discrete velocity grid (which is only required at this initial step). The
 502 total number of DSMC particles is governed by the number of merging cells
 503 M , as before (although without any issues related to negative mass). The
 504 SPARTA code uses particles with equal computational weight, and thus does
 505 not require any merging procedures.

506 Figure 9 shows the time spent on performing collisions during a single
 507 timestep plotted against the RMSE and MAE tail noise metrics. Different
 508 points on the curves correspond to different values of C_{RMS} , which directly
 509 affects the computational cost — the upper-left corner of the plot correspond-
 510 ing to low C_{RMS} (high computational cost, low error) and the lower-right
 511 corner corresponding to high C_{RMS} (low computational cost, high error). In-
 512 creasing the size of the region \mathcal{F} leads to increased noise in the distribution
 513 function, however, the computational effort is reduced only slightly (see or-
 514 ange and red curves). Increasing the number of merging regions M reduces
 515 error without a significant increase in computational cost (since the NTC
 516 collision procedure is much faster than the Monte Carlo evaluation of the
 517 collision integral for the collisions involving mass in \mathcal{Q}). It can be seen that
 518 at high C_{RMS} values, the hybrid approach is marginally more effective than

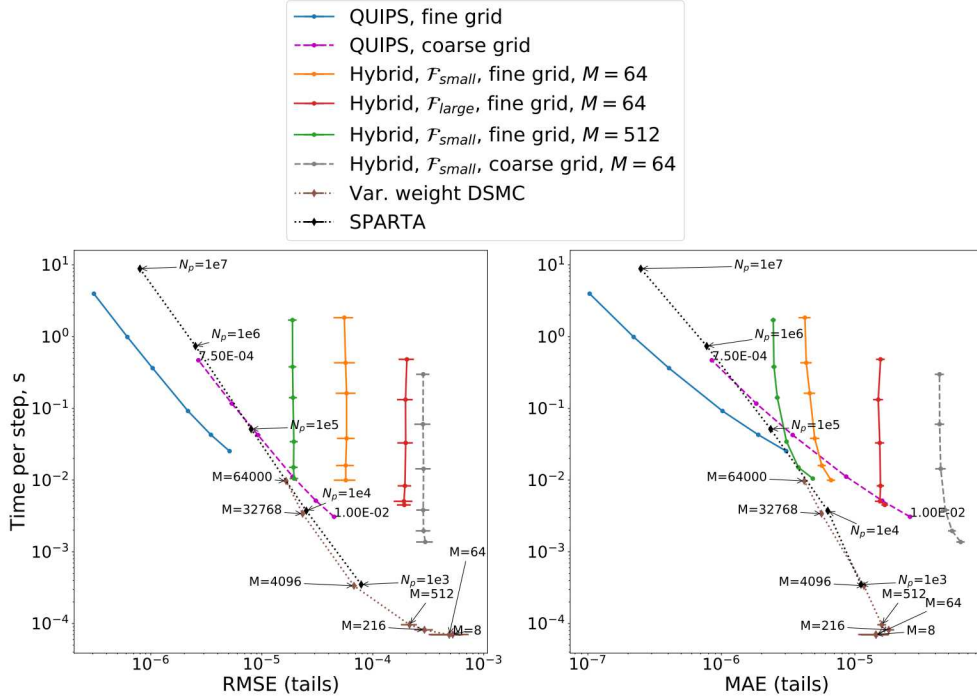


Figure 9: Time spent on performing collisions during a single timestep plotted against RMSE (left) and MAE (right) in the tails. For the hybrid and QUIPS codes, different points on the same curve correspond to different values of the C_{RMS} parameter (the pure discrete velocity curve on a coarse grid has the lowest and highest values of C_{RMS} annotated); for the variable weight DSMC calculation, different points on the curve correspond to different numbers of merging cells M); for the SPARTA computation, different points correspond to different numbers of model particles N_p .

519 the pure discrete velocity method, if one looks at the Mean Absolute Error
 520 (MAE) metric. The fact that the MAE metric, which is not as sensitive to
 521 large errors as is RMSE, is seen to decrease with decreasing values of C_{RMS}
 522 for the hybrid method, supports the hypothesis that the noise introduced by
 523 the merging procedure and hybridization with a low number of DSMC-type
 524 particles is more significant than the noise introduced by the Monte-Carlo
 525 scheme for evaluation of collisions involving particles from region \mathcal{Q} (which
 526 is influenced by the choice of C_{RMS}). It should be noted that our present
 527 hybrid code has not been specifically optimized and thus the metrics pre-
 528 sented only serve as a coarse comparison between the non-hybrid and hybrid
 529 approaches. Moreover, the computational cost of maintaining accuracy of an




Figure 10: Time spent on performing collisions during a single timestep plotted against RMSE in the 4th and 8th dimensionless moments. Different points on the same curve correspond to different values of the C_{RMS} parameter (the pure discrete velocity curve on a coarse grid has the lowest and highest values of C_{RMS} annotated); for the variable weight DSMC calculation, different points on the curve correspond to different numbers of merging cells (M); for the SPARTA computation, different points correspond to different numbers of model particles N_p .

530 already equilibrium distribution is likely different than maintaining accuracy
 531 in an evolving flow in which the tails and/or trace species play an important
 532 role in the dynamics.

533 For achieving very low-noise results, the pure discrete velocity method
 534 can be seen to be computationally more efficient than DSMC or hybrid com-
 535 putations.

536 Figure 10 shows the computational time spent on a collision step plotted
 537 against the root-mean-square error in the 4th and 8th moments. While
 538 the discrete velocity method can achieve significantly lower noise in the 4th
 539 moment than the hybrid or DSMC methods, it does so at a noticeable in-
 540 crease in computational cost. For the higher order (8th) moment, the hybrid
 541 approach requires less computational time for a given noise level than the
 542 pure discrete velocity method. Both the SPARTA and pure DSMC compu-
 543 tations are more cost-efficient than the hybrid and discrete velocity methods

544 for the higher-order moments.

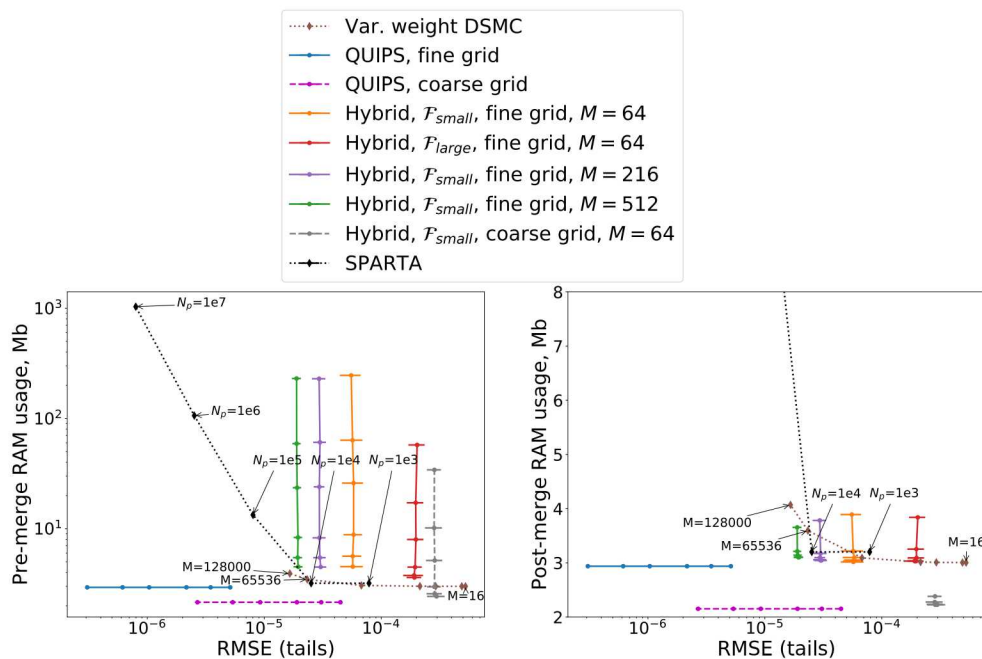


Figure 11: Memory usage in Mb plotted against RMSE in the tails. Left subplot shows pre-merge memory usage, right subplot shows post-merge memory usage. Different points on the same curve correspond to different values of the C_{RMS} parameter (the pure discrete velocity curve on a coarse grid has the lowest and highest values of C_{RMS} annotated; for the pure DSMC calculation, different points on the curve correspond to different numbers of merging cells M); for the SPARTA computation, different points correspond to different numbers of model particles N_p .

545 Another important aspect of computational efficiency is memory usage
 546 — modern supercomputers are often limited in their computational speed by
 547 memory bottlenecks, rather than processing power [77]; this is especially true
 548 for hybrid (GPGPU/CPU) architectures, since passing data between General
 549 Purpose Graphics Processing Units and CPUs is slow, and therefore, a code
 550 with a lower memory usage might be more amenable to large-scale computa-
 551 tions. GPGPU units have been successfully used both for DSMC [78, 79, 80]
 552 and discrete velocity method [81] computations, and therefore, a compari-
 553 son of memory usage of various computational methods might be indicative
 554 of potential scaling issues. It is important to note that for a given velocity
 555 grid, achieving lower noise with a discrete velocity method does not require

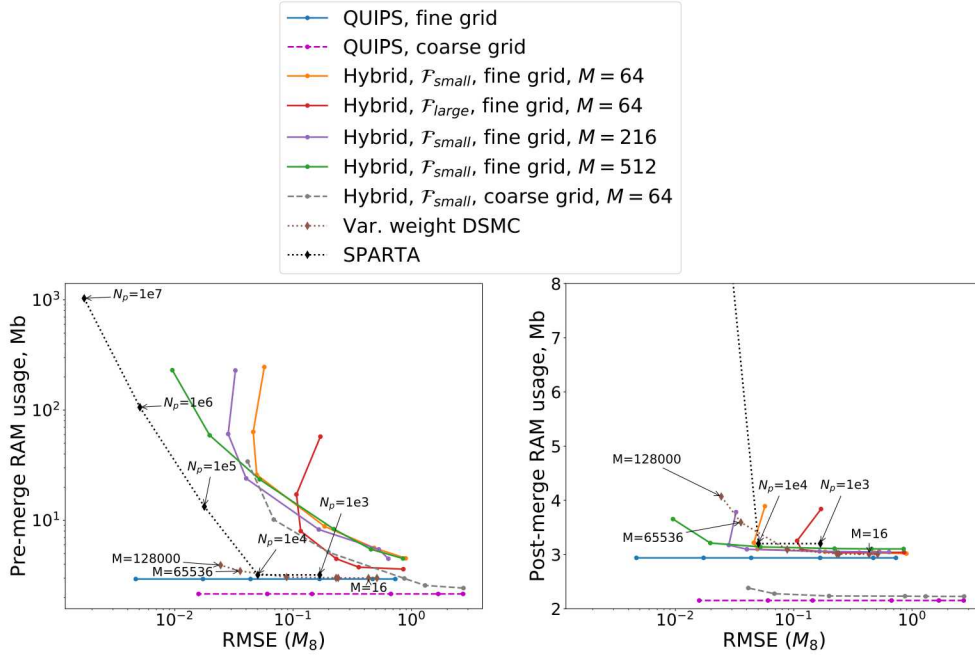


Figure 12: Memory usage in Mb plotted against RMSE of the 8th moment. Left subplot shows pre-merge memory usage, right subplot shows post-merge memory usage. Different points on the same curve correspond to different values of the C_{RMS} parameter (the pure discrete velocity curve on a coarse grid has the lowest and highest values of C_{RMS} annotated; for the variable weight DSMC calculation, different points on the curve correspond to different numbers of merging cells M); for the SPARTA computation, different points correspond to different numbers of model particles N_p .

556 additional memory (only more sampling operations are performed to better
 557 approximate the collision integral), whereas as with DSMC, the noise is in-
 558 versely correlated with the number of simulation particles N_p . That is, to
 559 achieve less noise in DSMC, one has to incur not only an increased compu-
 560 tational cost, but also an increased memory cost. Moreover, for practical
 561 multidimensional problems low noise DSMC requires high memory usage for
 562 all cells (or at least many cells) in physical space. Even if a collision scheme
 563 that requires a relatively low number of particles per cell to be accurate is
 564 used (such as the simplified Bernoulli trials collision procedure [82]), and time
 565 or ensemble averaging is employed, it still is problematic to use fixed-weight
 566 DSMC to correctly model flows with large number density ratios between
 567 different species. For example, if modelling an Ar/e⁻ mixture with a ratio of

568 the number densities of electrons and argon atoms equal to 0.001, approxi-
 569 mately 1000 argon simulator particles are required for each electron model
 570 particle; moreover, if the trace species (in this case electrons) are initialized
 571 with a non-equilibrium distribution, a larger number of particles is required
 572 to avoid initialization bias [75].

573 Figure 11 shows the pre- and post-merge memory usage plotted against
 574 the RMSE metric for the distribution function. The amount of memory used
 575 by the discrete velocity code is dependent only the number of points on the
 576 velocity grid, and is not affected by the choice of C_{RMS} ; for the hybrid code,
 577 the memory usage immediately before the merging procedure (after all the
 578 collisions at a given timestep have been performed and the number of particles
 579 in \mathcal{F} is high) was used to produce the left subplot. It can be seen that for
 580 a given velocity grid and size of \mathcal{F} , the memory usage of the hybrid method
 581 is virtually independent of the number of merging regions M . This is due to
 582 the fact that the plot shows peak memory usage (number of particles before
 583 merging), and the number of particles in \mathcal{F} before merging is mostly governed
 584 by the choice of C_{RMS} . While the pre-merge memory usage of the hybrid code
 585 can be rather high compared to a pure discrete velocity code, it should be
 586 emphasized that this is only needed to store the large amount of particles that
 587 are created in \mathcal{F} due to the depletion-replenishment collision scheme; and this
 588 amount of additional memory would not increase even if one increases the
 589 number of cells in physical space, since the collision and merging steps are
 590 performed on a cell-by-cell basis. This is in contrast with the SPARTA code,
 591 in which memory usage does not change (due to an absence of a merging
 592 step), which can make it unfeasible to run complex simulations with large
 593 numbers of model particles. The pure (variable-weight) DSMC version of the
 594 developed hybrid code is marginally more effective with regards to memory
 595 usage than SPARTA, which might be due to its ability to better resolve the
 596 tails via particles with small computational weights.

597 For the hybrid code, the post-merge memory usage still has a dependence
 598 on C_{RMS} due to failed merges — if a merge in a merging grid sub-cell fails, a
 599 lower C_{RMS} means that more particles remain in that sub-cell, and memory
 600 usage is higher. This problem could be alleviated with a more sophisticated
 601 merging procedure.

602 Finally, Fig. 12 shows the pre- and post-merge memory usage plotted
 603 against the RMSE of the 8th moment of the distribution function. When
 604 one considers the post-merge memory usage, it can be seen that the hybrid
 605 method and variable-weight DSMC approaches, as well as the discrete veloc-

606 ity method, are noticeably more memory efficient than a fixed-weight DSMC
607 code.

608 It can be concluded that with additional optimization of the computa-
609 tional procedures, the developed hybrid scheme in velocity space might be
610 an attractive strategy for a given acceptable error in the tails, especially if
611 variance reduction is implemented (as that can provide benefits both in terms
612 of computational efficiency and noise level).

613 5. Conclusion

614 A new approach to modeling rarefied gas flows based on a hybridization in
615 velocity space has been developed and verified for a single-species monatomic
616 gas.

617 The hybrid approach has been tested on the BKW relaxation problem,
618 and has been shown to provide improvement in the computational effort-
619 to-noise ratio in moments of the distribution function compared to a purely
620 discrete velocity simulation.

621 The influence of various simulation parameters on the noise in the tails of
622 the distribution function as well as noise in the moments of the distribution
623 function has been studied for a Maxwellian distribution.

624 It is shown that that the main source of noise in the tails of the dis-
625 tribution function when using the hybrid approach is the low number of
626 DSMC-type particles, rather than the value of C_{RMS} (the noise parameter
627 in the discrete velocity method).

628 With the basic principle of hybridization having now been successfully
629 demonstrated, further modifications of the hybrid method may include its
630 generalization to flows containing particles with internal energies, multiple-
631 species flows, flows with chemical reactions, and flows in which one imple-
632 ments a variance reduction algorithm [25]; these features have already been
633 proven in DSMC and discrete velocity methods separately, as discussed ear-
634 lier. The use of a more complex adaptive merging scheme to minimize the
635 amount of noise in the distribution function introduced by the merging pro-
636 cess and resolve issues with impossibility of merging due to presence of neg-
637 ative mass will be investigated in future work. The development of other
638 metrics (perhaps more problem-specific) to better understand the trade-offs
639 between various numerical methods and their parameters introduce also re-
640 mains to be investigated.

641 **6. Acknowledgments**

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

- 651 [1] G. Bird, Approach to translational equilibrium in a rigid sphere gas,
652 Phys. Fluids 6 (1963) 1518–1519.
- 653 [2] G. A. Bird, Molecular Gas Dynamics and the Direct Simulation of Gas
654 Flows, Clarendon, Oxford, England, UK, 1994.
- 655 [3] I. D. Boyd, Conservative species weighting scheme for the direct sim-
656 ulation Monte Carlo method, J. Thermophys. Heat Transfer 10 (1996)
657 579–585.
- 658 [4] G. Bird, Transition regime behavior of supersonic beam skimmers, Phys.
659 Fluids 19 (1976) 1486–1491.
- 660 [5] P. B. Clarke, A discrete velocity method for the Boltzmann equation
661 with internal energy and stochastic variance reduction, Ph.D. thesis,
662 University of Texas at Austin, 2015.
- 663 [6] S. Rjasanow, W. Wagner, A stochastic weighted particle method for the
664 Boltzmann equation, J. Comput. Phys. 124 (1996) 243–253.
- 665 [7] L. L. Baker, N. G. Hadjiconstantinou, Variance reduction for Monte
666 Carlo solutions of the Boltzmann equation, Phys. Fluids 17 (2005)
667 051703.
- 668 [8] G. A. Radtke, N. G. Hadjiconstantinou, W. Wagner, Low-noise monte
669 carlo simulation of the variable hard sphere gas, Phys. Fluids 23 (2011)
670 030606.

- 671 [9] C. R. Schrock, A. W. Wood, Distributional Monte Carlo solution tech-
672 nique for rarefied gasdynamics, *J. Thermophys. Heat Transfer* 26 (2012)
673 185–189.
- 674 [10] A. Nordsieck, B. L. Hicks, Monte Carlo evaluation of the Boltzmann
675 collision integral, in: *Proc. of the 5th Intern. Symposium on Rarefied*
676 *Gas Dynamics*, volume 1, pp. 695–710.
- 677 [11] D. Goldstein, B. Sturtevant, J. Broadwell, Investigations of the motion
678 of discrete-velocity gases, *Prog. Astronaut. Aeronaut.* 117 (1989) 100–
679 117.
- 680 [12] Z. Tan, P. L. Varghese, The δ - ε method for the Boltzmann equation, *J.*
681 *Comput. Phys.* 110 (1994) 327–340.
- 682 [13] P. Varghese, Arbitrary post-collision velocities in a discrete velocity
683 scheme for the Boltzmann equation, in: *Proc. of the 25th Intern. Sym-*
684 *posium on Rarefied Gas Dynamics*, pp. 225–232.
- 685 [14] F. Tcheremissine, Solution of the Boltzmann kinetic equation for low
686 speed flows, *Transp. Theory Stat. Phys.* 37 (2008) 564–575.
- 687 [15] A. Morris, P. Varghese, D. Goldstein, Improvement of a discrete velocity
688 boltzmann equation solver with arbitrary post-collision velocities, in:
689 *AIP Conf. Proc.*, volume 1084, AIP Publishing, pp. 458–463.
- 690 [16] O. Dodulad, F. Tcheremissine, Multipoint conservative projection
691 method for computing the Boltzmann collision integral for gas mixtures,
692 in: *AIP Conf. Proc.*, volume 1501, AIP Publishing, pp. 302–309.
- 693 [17] V. V. Aristov, Direct methods for solving the Boltzmann equation and
694 study of nonequilibrium flows, volume 60, Springer Science & Business
695 Media, 2012.
- 696 [18] P. Clarke, P. Varghese, D. Goldstein, A. Morris, P. Bauman,
697 D. Hegermiller, A novel discrete velocity method for solving the Boltz-
698 mann equation including internal energy and non-uniform grids in ve-
699 locity space, in: *AIP Conf. Proc.*, volume 1501, AIP Publishing, pp.
700 373–380.

- 701 [19] P. Clarke, P. Varghese, D. Goldstein, A low noise discrete velocity
702 method for the boltzmann equation with quantized rotational and vi-
703 brational energy, *J. Comput. Phys.* 352 (2018) 326–340.
- 704 [20] G. R. McNamara, G. Zanetti, Use of the Boltzmann equation to simulate
705 lattice-gas automata, *Phys. Rev. Lett.* 61 (1988) 2332.
- 706 [21] F. Higuera, S. Succi, R. Benzi, Lattice gas dynamics with enhanced
707 collisions, *Europhys. Lett.* 9 (1989) 345.
- 708 [22] S. Succi, *The lattice Boltzmann equation: for fluid dynamics and be-*
709 *yond*, Oxford university press, 2001.
- 710 [23] P. L. Bhatnagar, E. P. Gross, M. Krook, A model for collision pro-
711 cesses in gases. I. small amplitude processes in charged and neutral one-
712 component systems, *Phys. Rev.* 94 (1954) 511.
- 713 [24] A. Morris, P. Varghese, D. Goldstein, Monte Carlo solution of the Boltz-
714 mann equation via a discrete velocity model, *J. Comput. Phys.* 230
715 (2011) 1265–1280.
- 716 [25] P. Clarke, P. Varghese, D. Goldstein, Discrete velocity computations
717 with stochastic variance reduction of the Boltzmann equation for gas
718 mixtures, in: *AIP Conf. Proc*, volume 1628, AIP Publishing, pp. 1032–
719 1039.
- 720 [26] F. Cheremisin, Solving the Boltzmann equation in the case of passing
721 to the hydrodynamic flow regime, in: *Doklady Physics*, volume 45,
722 Springer, pp. 401–404.
- 723 [27] A. Y. Aydemir, A unified Monte Carlo interpretation of particle simula-
724 tions and applications to non-neutral plasmas, *Phys. Plasmas* 1 (1994)
725 822–831.
- 726 [28] G. J. Wilkie, W. Dorland, Fundamental form of the electrostatic δf -
727 PIC algorithm and discovery of a converged numerical instability, *Phys.*
728 *Plasmas* 23 (2016) 052111.
- 729 [29] S. Brunner, E. Valeo, J. A. Krommes, Collisional delta-f scheme with
730 evolving background for transport time scale simulations, *Phys. Plasmas*
731 6 (1999) 4504–4521.

- 732 [30] M. H. Gorji, N. Andric, P. Jenny, Variance reduction for Fokker–Planck
733 based particle Monte Carlo schemes, *J. Comput. Phys.* 295 (2015) 644–
734 664.
- 735 [31] B. S. Collyer, C. Connaughton, D. A. Lockerby, Importance sampling
736 variance reduction for the Fokker–Planck rarefied gas particle method,
737 *J. Comput. Phys.* 325 (2016) 116–128.
- 738 [32] Y. Poondla, P. Varghese, D. Goldstein, K. Higdon, Modeling of chemical
739 reactions using Quasi-Particle Simulation (QuiPS), in: *AIP Conf. Proc.*,
740 volume 2132, AIP Publishing, p. 140002.
- 741 [33] K. D. Lathrop, Ray effects in discrete ordinates equations, *Nuclear
742 Science and Engineering* 32 (1968) 357–369.
- 743 [34] W. Fiveland, Discrete-ordinates solutions of the radiative transport
744 equation for rectangular enclosures, *J. Heat Transfer* 106 (1984) 699–
745 706.
- 746 [35] J. C. Chai, H. S. Lee, S. V. Patankar, Ray effect and false scattering in
747 the discrete ordinates method, *Numer. Heat Transfer, Part B* 24 (1993)
748 373–389.
- 749 [36] A. Sekaran, P. Varghese, D. Goldstein, An analysis of numerical conver-
750 gence in discrete velocity gas dynamics for internal flows, *J. Comput.
751 Phys.* 365 (2018) 226–242.
- 752 [37] M. T. Ho, J. Li, L. Wu, J. M. Reese, Y. Zhang, A comparative study of
753 the DSBGK and DVM methods for low-speed rarefied gas flows, *Com-
754 put. Fluids* (2019).
- 755 [38] D. C. Wadsworth, I. J. Wysong, Vibrational favoring effect in DSMC
756 dissociation models, *Phys. Fluids* 9 (1997) 3873–3884.
- 757 [39] J. Kim, I. Boyd, State resolved thermochemical modeling of nitrogen
758 using DSMC, in: *43rd AIAA Thermophysics Conference*, p. 2991.
- 759 [40] D. Liechty, M. Lewis, Treatment of electronic energy level transition
760 and ionization following the particle-based chemistry model, in: *48th
761 AIAA Aerospace Sciences Meeting Including the New Horizons Forum
762 and Aerospace Exposition*, p. 449.

- 763 [41] A. B. Carlson, A. Hassan, Radiation modeling with direct simulation
764 Monte Carlo, *J. Thermophys. Heat Transfer* 6 (1992) 631–636.
- 765 [42] A. Fierro, C. Moore, B. Scheiner, B. T. Yee, M. M. Hopkins, Radiation
766 transport in kinetic simulations and the influence of photoemission on
767 electron current in self-sustaining discharges, *J. Phys. D: Appl. Phys.*
768 50 (2017) 065202.
- 769 [43] V. V. Serikov, S. Kawamoto, K. Nanbu, Particle-in-cell plus direct simu-
770 lation Monte Carlo (PIC-DSMC) approach for self-consistent plasma-gas
771 simulations, *IEEE Trans. Plasma Sci.* 27 (1999) 1389–1398.
- 772 [44] A. Fierro, C. Moore, B. Yee, M. Hopkins, Three-dimensional kinetic
773 modeling of streamer propagation in a nitrogen/helium gas mixture,
774 *Plasma Sources Sci. Technol.* 27 (2018) 105008.
- 775 [45] T. E. Schwartzentruber, M. S. Grover, P. Valentini, Direct molecular
776 simulation of nonequilibrium dilute gases, *J. Thermophys. Heat Transfer*
777 32 (2017) 892–903.
- 778 [46] R. Roveda, D. B. Goldstein, P. L. Varghese, Hybrid Euler/particle
779 approach for continuum/rarefied flows, *J. Spacecr. Rockets* 35 (1998)
780 258–265.
- 781 [47] H. Carlson, R. Roveda, I. Boyd, G. Candler, A hybrid CFD-DSMC
782 method of modeling continuum-rarefied flows, in: 42nd AIAA Aerospace
783 Sciences Meeting and Exhibit, p. 1180.
- 784 [48] H. S. Wijesinghe, N. G. Hadjiconstantinou, Discussion of hybrid
785 atomistic-continuum methods for multiscale hydrodynamics, *Int. J.*
786 *Multiscale Comput. Eng.* 2 (2004).
- 787 [49] T. E. Schwartzentruber, L. C. Scalabrin, I. D. Boyd, Hybrid particle-
788 continuum simulations of nonequilibrium hypersonic blunt-body flow-
789 fields, *J. Thermophys. Heat Transfer* 22 (2008) 29–37.
- 790 [50] V. Kolobov, R. Arslanbekov, Towards adaptive kinetic-fluid simulations
791 of weakly ionized plasmas, *J. Comput. Phys.* 231 (2012) 839–869.
- 792 [51] Y. P. Raizer, *Gas Discharge Physics*, Springer-Verlag Berlin Heidelberg,
793 1991.

- 794 [52] N. Crouseilles, P. Degond, M. Lemou, A hybrid kinetic/fluid model for
795 solving the gas dynamics Boltzmann–BGK equation, *J. Comput. Phys.*
796 199 (2004) 776–808.
- 797 [53] G. Dimarco, L. Pareschi, Hybrid multiscale methods II. Kinetic equa-
798 tions, *Multiscale Model. Simul.* 6 (2008) 1169–1197.
- 799 [54] T.-J. Pan, K. A. Stephani, Investigation of velocity-space coupling ap-
800 proach in DSMC for tail-driven processes, in: *AIP Conf. Proc.*, volume
801 1786, AIP Publishing, p. 050017.
- 802 [55] T.-J. Pan, K. A. Stephani, Investigation of a coupling approach of
803 DSMC and DG methods for tail-driven processes, in: *47th AIAA Ther-
804 mophysics Conference*, p. 4023.
- 805 [56] J. H. Ferziger, H. G. Kaper, H. G. Kaper, *Mathematical theory of trans-
806 port processes in gases*, North-Holland, 1972.
- 807 [57] S. Chapman, T. G. Cowling, D. Burnett, *The mathematical theory of
808 non-uniform gases: an account of the kinetic theory of viscosity, thermal
809 conduction and diffusion in gases*, Cambridge university press, 1990.
- 810 [58] M. Ivanov, S. Rogasinsky, Analysis of numerical techniques of the Direct
811 Simulation Monte Carlo method in the rarefied gas dynamics, *Russian
812 Journal of numerical analysis and mathematical modelling* 3 (1988) 453–
813 466.
- 814 [59] D. Petkow, A. Mirza, G. Herdrich, S. Fasoulas, Treatment of differently
815 weighted particles in reactive re-entry flows with DSMC, in: *AIP Conf.
816 Proc.*, volume 1501, AIP Publishing, pp. 1507–1514.
- 817 [60] D. P. Schmidt, C. Rutland, A new droplet collision algorithm, *J. Com-
818 put. Phys.* 164 (2000) 62–80.
- 819 [61] M. H. Gorji, P. Jenny, Fokker–Planck–DSMC algorithm for simulations
820 of rarefied gas flows, *J. Comput. Phys.* 287 (2015) 110–129.
- 821 [62] E. Jun, M. H. Gorji, M. Grabe, K. Hannemann, Assessment of the cubic
822 Fokker–Planck–DSMC hybrid method for hypersonic rarefied flows past
823 a cylinder, *Comput. Fluids* 168 (2018) 1–13.

- 824 [63] G. Lapenta, J. U. Brackbill, Dynamic and selective control of the number
825 of particles in kinetic plasma simulations, *J. Comput. Phys.* 115 (1994)
826 213–227.
- 827 [64] D. W. Hewett, Fragmentation, merging, and internal dynamics for PIC
828 simulation with finite size particles, *J. Comput. Phys.* 189 (2003) 390–
829 426.
- 830 [65] D. R. Welch, T. C. Genoni, R. E. Clark, D. V. Rose, Adaptive particle
831 management in a particle-in-cell code, *J. Comput. Phys.* 227 (2007)
832 143–155.
- 833 [66] J. Teunissen, U. Ebert, Controlling the weights of simulation particles:
834 adaptive particle management using k-d trees, *J. Comput. Phys.* 259
835 (2014) 318–330.
- 836 [67] G. Lapenta, J. U. Brackbill, Dynamic and selective control of the number
837 of particles in kinetic plasma simulations, *J. Comput. Phys.* 115 (1994)
838 213–227.
- 839 [68] S. Rjasanow, T. Schreiber, W. Wagner, Reduction of the number of
840 particles in the stochastic weighted particle method for the Boltzmann
841 equation, *J. Comput. Phys.* 145 (1998) 382–405.
- 842 [69] R. S. Martin, J.-L. Cambier, Octree particle management for DSMC
843 and PIC simulations, *J. Comput. Phys.* 327 (2016) 943–966.
- 844 [70] R. S. Martin, J.-L. Cambier, Moment preserving adaptive particle
845 weights using octree velocity distributions for PIC simulations, in: *AIP*
846 *Conf. Proc.*, volume 1501, AIP Publishing, pp. 872–879.
- 847 [71] A. V. Bobylev, One class of invariant solutions of the Boltzmann equa-
848 tion, in: *Akademiia Nauk SSSR Doklady*, volume 231, pp. 571–574.
- 849 [72] M. Krook, T. T. Wu, Exact solutions of the Boltzmann equation, *Phys.*
850 *Fluids* 20 (1977) 1589–1595.
- 851 [73] L. Mieussens, Discrete velocity model and implicit scheme for the BGK
852 equation of rarefied gas dynamics, *Math. Models Methods Appl. Sci.* 10
853 (2000) 1121–1149.

- 854 [74] D. Rader, M. Gallis, J. Torczynski, W. Wagner, Direct simulation Monte
855 Carlo convergence behavior of the hard-sphere-gas thermal conductivity
856 for Fourier heat flow, *Phys. Fluids* 18 (2006) 077102.
- 857 [75] W. McDoniel, C. Moore, G. Radtke, K. Cartwright, M. Bettencourt, On-
858 average error from DSMC collisions, in: *Book of Abstracts, International*
859 *Conference on Phenomena in Ionized Gases*, Sapporo, Japan.
- 860 [76] S. Plimpton, M. Gallis, SPARTA direct simulation Monte Carlo (DSMC)
861 simulator, Sandia National Laboratories, USA, see <http://sparta.sandia.gov> (2015).
862
- 863 [77] N. Attig, P. Gibbon, T. Lippert, Trends in supercomputing: The Euro-
864 pean path to exascale, *Comput. Phys. Commun.* 182 (2011) 2041–2046.
- 865 [78] C.-C. Su, C.-W. Hsieh, M. Smith, M. Jermy, J.-S. Wu, Parallel direct
866 simulation Monte Carlo computation using CUDA on GPUs, in: *AIP*
867 *Conf. Proc.*, volume 1333, AIP Publishing, pp. 343–347.
- 868 [79] M. Goldsworthy, A GPU–CUDA based direct simulation Monte Carlo
869 algorithm for real gas flows, *Comput. Fluids* 94 (2014) 58–68.
- 870 [80] A. Kashkovsky, 3D DSMC computations on a heterogeneous CPU-GPU
871 cluster with a large number of GPUs, in: *AIP Conf. Proc.*, volume 1628,
872 AIP Publishing, pp. 192–198.
- 873 [81] E. Malkov, S. Poleshkin, M. Ivanov, Discrete velocity scheme for solving
874 the Boltzmann equation with the GPGPU, in: *AIP Conf. Proc.*, volume
875 1501, AIP Publishing, pp. 318–325.
- 876 [82] S. K. Stefanov, On DSMC calculations of rarefied gas flows with small
877 number of particles in cells, *SIAM Journal on Scientific Computing* 33
878 (2011) 677–702.

879 **Appendix A. Bias introduced by the merging procedure**

880 To estimate the influence of the particle merging procedure on the higher-
881 order moments, and to better understand the amount of noise and bias in-
882 troduced by the merging, the BKW relaxation test case was run with a
883 $29 \times 29 \times 29$ velocity grid with \mathcal{F}_{small} and different values of the number of




Figure A.13: Time evolution of the 4th moment for the BKW relaxation. The black curve is the analytic solution. The red lines connect the pre- and post-merge moments at a fixed timestep, while the blue lines connect a post-merge moment at timestep i with a pre-merge moment at timestep $i + 1$, that is, they show the temporal evolution of the moment. A C_{RMS} value of 0.005 was used, and the number of merging grid cells M was equal to 8.

884 merging cells M . The simulation was run with a scaled timestep of 0.1 for 100
 885 timesteps (10 collision times); and at each time step, the difference between
 886 the post- and pre-merge value of the considered moments was computed:

$$\Delta \hat{M}_l(t) = \hat{M}_l^{post}(t) - \hat{M}_l^{pre}(t). \quad (\text{A.1})$$

887 Here $\hat{M}_l^{pre}(t)$, $\hat{M}_l^{post}(t)$ are the values of the scaled moment of order l at time t
 888 before and after the merging procedure, correspondingly. The time-averages
 889 of these time-series were taken:

$$\Delta \hat{M}_l = \frac{\sum_{\hat{t}=0}^{10} \Delta \hat{M}_l(\hat{t})}{100}, \quad (\text{A.2})$$

890 and ensemble-averaged over 200 ensembles.

891 Figure A.13 shows the temporal evolution of the 4th moment for a single
 892 run with $M = 8$ and explicitly shows the changes in the values of the moment
 893 due to the merging procedure (red lines). We see that for such a low number
 894 of merging cells (8), the merging adds considerable noise.




Figure A.14: Average difference in the 4th moment introduced by the merging procedure for different numbers of merging cells M . Note that all cases lead to a negative change in M_4 , but that with a moderate number of merging cells M the magnitude of the error can be quite small. The error bars show a range of plus/minus one standard deviation.

895 Figure A.14 shows the average bias introduced in the 4th moment by the
 896 merging procedure as a function of number of merging cells M . It can be
 897 seen that merging causes a consistent underestimation of the moments, and
 898 that increasing the number of merging cells leads to a reduction in the bias.
 899 The bias introduced in the 4th moment was also approximated as a function
 900 of M :

$$\log_{10}(-\Delta M_4) \approx 0.44 - 0.917M^{1/5}, \quad (\text{A.3})$$

901 and the analytic fit is also shown on Fig. A.14.