

LOW-VARIANCE DIRECT MONTE CARLO SIMULATIONS USING IMPORTANCE WEIGHTS

HUSAIN A. AL-MOHSEN¹ AND NICOLAS G. HADJICONSTANTINO¹

Abstract. We present an efficient approach for reducing the statistical uncertainty associated with direct Monte Carlo simulations of the Boltzmann equation. As with previous variance-reduction approaches, the resulting *relative* statistical uncertainty in hydrodynamic quantities (statistical uncertainty normalized by the characteristic value of quantity of interest) is small and independent of the magnitude of the deviation from equilibrium, making the simulation of arbitrarily small deviations from equilibrium possible. In contrast to previous variance-reduction methods, the method presented here is able to substantially reduce variance with very little modification to the standard DSMC algorithm. This is achieved by introducing an auxiliary equilibrium simulation which, *via* an importance weight formulation, uses the same particle data as the non-equilibrium (DSMC) calculation; subtracting the equilibrium from the non-equilibrium hydrodynamic fields drastically reduces the statistical uncertainty of the latter because the two fields are correlated. The resulting formulation is simple to code and provides considerable computational savings for a wide range of problems of practical interest. It is validated by comparing our results with DSMC solutions for steady and unsteady, isothermal and non-isothermal problems; in all cases very good agreement between the two methods is found.

Mathematics Subject Classification. 60H30, 76P05.

Received June 2nd, 2009. Revised January 2nd, 2010.
Published online August 26, 2010.

1. INTRODUCTION

Recent [16] interest in small scale, low-speed, gaseous flows has renewed the need for efficient methods for solving Boltzmann-type kinetic equations [10]. Unfortunately, in this regime, traditional Monte Carlo solution methods, the most prevalent of which is direct simulation Monte Carlo (DSMC) [7], become very inefficient due to the overwhelming cost associated with resolving low-signal flows. As a result of this limitation, variance reduction approaches have recently received considerable attention: Baker and Hadjiconstantinou [4] recently showed that solving for the deviation from equilibrium drastically reduces statistical uncertainty and thus enables the simulation of arbitrarily small deviations from equilibrium. The same authors also showed that variance-reduced formulations can be developed for both particle methods [5] and PDE-type approaches [4,6]; a particle method that is equivalent to the one detailed in [5] has also been proposed by Chun and Koch [11]. Unfortunately, in these particle methods, particle cancellation in the collision dominated regime is required for stability [5,11]; this adversely affects both accuracy and efficiency.

Keywords and phrases. DSMC, variance reduction, microscale gas flow.

¹ Mechanical Engineering Dept., MIT, 77 Massachusetts Ave., Cambridge, MA 02139, USA. husain@mit.edu; ngh@mit.edu

More recently, Homolle and Hadjiconstantinou [18,19] have shown that the above limitation can be removed by using a special form of the linear part of the hard-sphere collision operator, originally derived by Hilbert [9], in which the integration over the scattering angle has been performed analytically. Using this approach, they developed an efficient particle method for simulating the hard-sphere gas known as LVDSMC (Low-Variance Deviational Simulation Monte Carlo) [18,19], which does not require a particle cancellation scheme. LVDSMC differs from DSMC only in ways necessary to simulate the deviation from equilibrium (see [4,19]), and so retains the strengths of particle methods, such as DSMC, while exhibiting substantially reduced variance. The LVDSMC method has been recently theoretically analyzed by Wagner [28], who also proposed algorithms for simulating the Variable Hard-Sphere (VHS) model [7]. Moreover, LVDSMC has been extended to treat the relaxation-time approximation by Radtke and Hadjiconstantinou [23].

In the present paper we present an *alternative* particle approach for achieving variance reduction. Instead of using a particle scheme that differs from DSMC to some extent, the method described here achieves the variance reduction by using a combination of a DSMC simulation and an auxiliary equilibrium simulation. Our framework is based on the general idea of correlated sampling [24] that has been successfully used in a number of diverse fields [14,21]. More specifically, it relies on the observation that the statistical uncertainty of non-equilibrium hydrodynamic fields can be significantly reduced by formulating the desired result in terms of a difference between the non-equilibrium (molecular) data and some *correlated* data whose moments are known (*e.g.* an equilibrium simulation).

The advantage of such an approach is that the non-equilibrium simulation (DSMC) remains essentially *unaltered*, which, in addition to historical reasons, it might be preferable in a number of situations. For example, the present formulation can be readily extended to other molecular-interaction models [2]; in contrast, extension of LVDSMC to other molecular models – other than the relaxation-time approximation [23] – is more challenging [28]. Another situation where the present formulation holds an advantage is more complex collision processes, such as chemically reacting flows.

One key challenge associated with the proposed approach, which we will refer to as VRDSMC, lies in developing a framework that allows the auxiliary simulation to be run in parallel and using the same molecular data as the main DSMC simulation (so that the two data streams remain correlated), while at the same time the former represents a chosen equilibrium state. Our formulation uses the concept of importance weights [12] to achieve this. In this paper we discuss how importance weights defined by initial and boundary conditions evolve in time and space under the action of Boltzmann dynamics. Unfortunately, since the present method is not using the Hilbert collision operator used in LVDSMC, it exhibits stability problems in the collision dominated regime that are analogous to those reported in [5,11]. An extensive discussion of how these problems are overcome in this work is given in Section 3.2. Our formulation is validated by comparison to DSMC results for a number of benchmark flows in Section 4.

2. VARIANCE REDUCTION USING IMPORTANCE WEIGHTS: BASIC CONCEPTS

Let $\langle R \rangle_\varphi = \int f_\varphi(\mathbf{c})R(\mathbf{c})d\mathbf{c}$ denote the expectation value of the hydrodynamic variable $R(\mathbf{c})$ over a distribution $f_\varphi(\mathbf{c})$, while $\bar{R}_\varphi = N_{\text{eff}}^{-1} \sum_{i=1}^N R(\mathbf{c}_i)$ denotes an estimate of $\langle R \rangle_\varphi$ obtained using N samples drawn from $f_\varphi(\mathbf{c})$, where N_{eff} is the number of physical particles represented by each simulation particle. In the interest of simplicity, and without loss of generality, here and in the remainder of the paper, the spatial and temporal dependence of the distribution and resulting hydrodynamic quantities will be implied if not explicitly shown.

The basic idea behind the present approach is to produce a variance-reduced estimator, denoted by \bar{R}^{VR} , by writing

$$\bar{R}^{VR} = \bar{R} - \bar{R}_{\text{eq}} + \langle R \rangle_{\text{eq}} \quad (2.1)$$

and ensuring that \bar{R} and \bar{R}_{eq} are estimated using correlated molecular data, while the distribution f_{eq} (assumed here to be an equilibrium one) is chosen such that $\langle R \rangle_{\text{eq}}$ is known. This concept, which has been used in polymer simulation for a number of years [21], is illustrated in Figure 1 for a gas relaxation problem; the figure shows

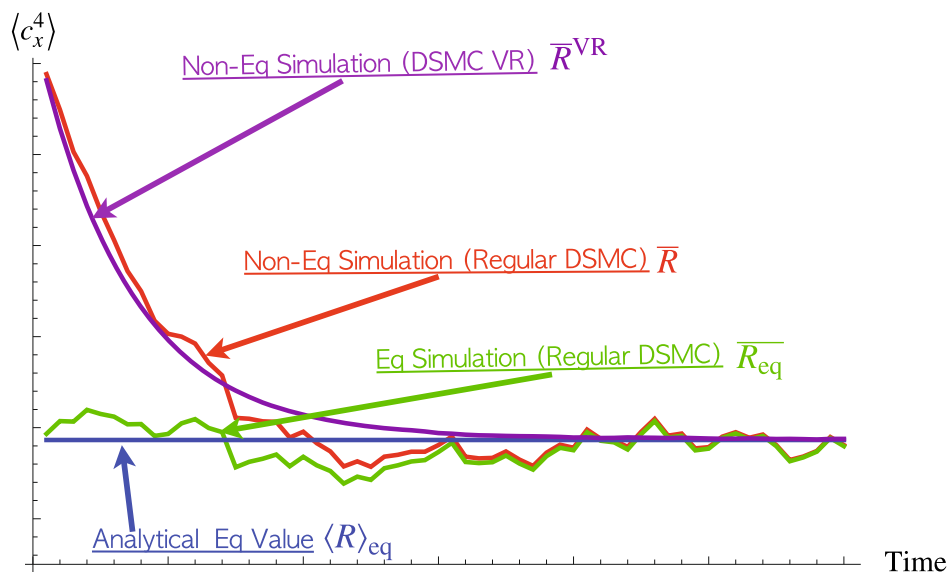


FIGURE 1. Illustration of the variance reduction principle for a molecular relaxation problem [3]. The variance of \overline{R}^{VR} is significantly reduced by replacing the “noisy” estimate \overline{R}_{eq} with its exact expected value $\langle R \rangle_{eq}$.

how actual simulation data [3] of \overline{R} , \overline{R}_{eq} , and $\langle R \rangle_{eq}$, with $R \equiv c_x^4$, can be combined to yield the low-uncertainty estimator \overline{R}^{VR} .

The major challenge associated with implementing this approach lies with the development of a framework which provides molecular data that sample $f_{eq}(\mathbf{c})$, while at the same time are correlated with the (non-equilibrium) DSMC data that sample $f(\mathbf{c})$, the non-equilibrium single-particle distribution function [7]. To achieve this we have chosen to use an importance weight formulation and proceed by defining the importance weight $W(\mathbf{c})$ as

$$W(\mathbf{c}) = \frac{f_{eq}(\mathbf{c})}{f(\mathbf{c})}. \quad (2.2)$$

Using this we can write

$$\langle R \rangle_{eq} = \int \frac{f_{eq}(\mathbf{c})}{f(\mathbf{c})} f(\mathbf{c}) R(\mathbf{c}) d\mathbf{c}$$

and thus

$$\overline{R}_{eq} = N_{eff} \sum_{i=1}^N W_i R(\mathbf{c}_i)$$

where $W_i = W(\mathbf{c}_i)$ and \mathbf{c}_i is drawn from $f(\mathbf{c})$. In other words, \overline{R}_{eq} may be evaluated using samples from f with the weights W_i providing a correction for the relative frequency of each sample \mathbf{c}_i in the two distributions.

Using this formulation, (2.1) becomes

$$\overline{R}^{VR} = N_{eff} \sum_{i=1}^N (1 - W_i) R(\mathbf{c}_i) + \langle R \rangle_{eq} \quad (2.3)$$

and can be evaluated by explicitly sampling the non-equilibrium distribution *only*, provided the weights W_i are known. In the next section we present a prescription for initializing the simulation weights, as well as the rules that govern for the dynamical evolution of these weights based on the governing Boltzmann equation.

It is straightforward to verify that the variance of the estimator \overline{R}^{VR} is significantly smaller than the variance of \overline{R} when $|W_i - 1| \ll 1$, as will be the case when f and f_{eq} are close [24], *e.g.* $f(\mathbf{c}) \approx f_{\text{eq}}(\mathbf{c}) + Ma\phi(\mathbf{c})$, for $Ma \ll 1$ [26]. In this sense, a Maxwell-Boltzmann equilibrium distribution

$$f_{\text{MB}}(\mathbf{c}) = \frac{n_{\text{MB}}}{(\pi c_{\text{MB}}^2)^{3/2}} \exp\left(-\frac{|\mathbf{c} - \mathbf{u}_{\text{MB}}|^2}{c_{\text{MB}}^2}\right)$$

is a reasonable and convenient “reference equilibrium” choice, since in cases where variance reduction is needed, *i.e.* when the deviation from equilibrium is small, the parameters n_{MB} , c_{MB} and \mathbf{u}_{MB} can usually be intuitively chosen such that $|W_i - 1| \ll 1$. In this work, we take the reference equilibrium (f_{eq}) to be an appropriately chosen global equilibrium distribution

$$f_{\text{eq}}(\mathbf{c}) = f_{\text{MB},0}(\mathbf{c}) = \frac{n_0}{(\pi c_0^2)^{3/2}} \exp\left(-\frac{|\mathbf{c}|^2}{c_0^2}\right)$$

where n_0 is the global equilibrium number density, $c_0 = \sqrt{2kT_0/m}$ is the most probable molecular speed based on the global equilibrium temperature T_0 , k is Boltzmann’s constant and m is the molecular mass.

3. VARIANCE REDUCTION USING IMPORTANCE WEIGHTS: VRDSMC

In the interest of simplicity, we focus on a hard-sphere gas of molecular diameter d , in the absence of external fields. We also define $Kn = \lambda/L$ to be the Knudsen number based on the molecular mean free path $\lambda = (\sqrt{2}\pi n_0 d^2)^{-1}$ and the characteristic hydrodynamic lengthscale L . As will be clear from the analysis below, both assumptions (hard-sphere gas and lack of external fields) can be easily relaxed. In fact, the relative simplicity associated with extending this method to other interaction models (*e.g.* variable hard sphere [7]) is one of the main advantages of the proposed approach. A discussion of external-field implementations can be found in [20].

The DSMC method effectively integrates the Boltzmann equation in time by simulating molecular motion in a series of timesteps, each of length Δt , during which a collisionless advection substep and collision substep are performed [7]. The advection substep updates the distribution function due to the action of the advection operator

$$\frac{\partial f}{\partial t} + \mathbf{c} \cdot \frac{\partial f}{\partial \mathbf{x}} = 0 \quad (3.1)$$

while the collision substep updates the distribution function due to the action of the collision operator

$$\left[\frac{\partial f}{\partial t}\right]_{\text{Coll}} = \frac{1}{2} \iiint (\delta'_1 + \delta'_2 - \delta_1 - \delta_2) f_1 f_2 c_r \sigma d\Omega d\mathbf{c}_1 d\mathbf{c}_2. \quad (3.2)$$

A convergence proof for this algorithm can be found in [27]; an analysis of the error associated with the timestep discretization can be found in [13,15]. Here, we use a form of the collision integral that is convenient for discussing particle methods [4,5,19]. In this form, σ is the differential collision cross-section, $f_1 = f(\mathbf{x}, \mathbf{c}_1, t)$, $f_2 = f(\mathbf{x}, \mathbf{c}_2, t)$, $\delta_1 = \delta(\mathbf{c}_1 - \mathbf{c})$, $\delta_2 = \delta(\mathbf{c}_2 - \mathbf{c})$, $\delta'_1 = \delta(\mathbf{c}'_1 - \mathbf{c})$ and $\delta'_2 = \delta(\mathbf{c}'_2 - \mathbf{c})$. Also, $\mathbf{c}_1, \mathbf{c}_2$ are the pre-collision velocities, $c_r = |\mathbf{c}_1 - \mathbf{c}_2|$ is the magnitude of the relative velocity vector, and $\mathbf{c}'_1, \mathbf{c}'_2$ are the post-collision velocities, related to the pre-collision velocities through the scattering angle Ω . In this section and in the remainder of the paper, unless otherwise stated, integration in velocity space extends over \mathbb{R}^3 , while the solid angle integration is over the surface of the unit sphere.

3.1. Weight update rules

As explained above, our formulation utilizes an auxiliary simulation that uses the same samples (particles) as the non-equilibrium DSMC simulation, but instead samples an equilibrium distribution. The difference

between the two distributions is captured by the set of weights W_i which give the relative frequency of finding particle i of velocity \mathbf{c}_i in the equilibrium simulation relative to the non-equilibrium simulation (see (2.2)). As a result of this formulation, a separate (equilibrium) simulation is not actually performed; instead the evolution of the equilibrium simulation is tracked by updating the weights W_i . Below, we develop rules describing the initialization and evolution of these weights.

3.1.1. Initialization

The initial values of the weights can be readily determined since *initially* both f and f_{eq} are explicitly known (*i.e.* the initial condition is known exactly for both the equilibrium and non-equilibrium simulation). Specifically, when DSMC particles are initialized at the start of the simulation, the weight corresponding to each particle is directly calculated using the definition (2.2).

3.1.2. Boundary conditions

Boundary conditions are in some ways similar to initial conditions in the sense that, in typical applications of interest, particle velocities (and thus weights) are redrawn from a pre-specified distribution when interacting with a system boundary. In many cases, the boundary conditions associated with the Boltzmann equation are of the no-flux type [8], *e.g.* solid wall; other types of boundary conditions, such as open boundaries, are also straightforward to treat.

Particles colliding with a wall are returned to the computational domain such that no net mass crosses the boundary. The velocity distribution from which the velocities of particles returning from the wall are drawn depends on the assumed gas-wall interaction. One of the most popular models is known as the Maxwell accommodation model [8], in which a fraction α of the molecules colliding with the wall are diffusely reflected, while the remaining $(1 - \alpha)$ fraction is specularly reflected. Here we will discuss the fully diffuse case ($\alpha = 1$) which highlights the essence of our approach.

In the diffuse case, the boundary condition for particles that come in contact with the wall is

$$f_w(\mathbf{c}) = \frac{n_w}{(\pi c_w^2)^{3/2}} \exp\left(-\frac{|\mathbf{c} - \mathbf{u}_w|^2}{c_w^2}\right) = n_w F_w(\mathbf{c}) \quad (3.3)$$

parameterized by the wall properties, namely the “wall number density” n_w , the most probable speed c_w based on the wall temperature T_w and the wall velocity \mathbf{u}_w . The wall number density may be thought of as the number density of a gas that would be in equilibrium with the boundary; it is determined from mass conservation at the wall, given here for the case $\mathbf{n} \cdot \mathbf{u}_w = 0$

$$\int_{\mathbf{c} \cdot \mathbf{n} < 0} \mathbf{c} \cdot \mathbf{n} f d\mathbf{c} = -n_w \int_{\mathbf{c} \cdot \mathbf{n} > 0} \mathbf{c} \cdot \mathbf{n} F_w d\mathbf{c} \quad (3.4)$$

where \mathbf{n} denotes the wall normal pointing into the gas.

The weights of re-emitted particles can be readily calculated from definition (2.2) provided n_w/n_0 is known (this ratio arises in (2.2) because the boundary condition for the auxiliary simulation is $f_{\text{MB},0}$). In the special case of isothermal constant-density flows (*e.g.* low-speed), n_w/n_0 can be taken to be equal to unity². In the more general case, n_w/n_0 needs to be calculated from the simulation data using (3.4) applied to *both* the equilibrium and non-equilibrium simulations.

In our implementation, this is achieved [2] by the following process. At the beginning of every timestep we assume a value for n_w/n_0 , *e.g.* $n_w/n_0 = 1$, and particle weights for particles colliding with the wall are assigned based on this value. At the end of the timestep the total weight of particles that collided with the wall during that timestep is tallied and the weight of particles re-emitted from the wall rescaled, such that there is no net change in the total equilibrium particle weights due to interaction with the wall.

²Here we assume that the (constant) gas density is chosen as reference condition n_0 .

3.1.3. *The advection substep*

During the advection substep, the auxiliary (equilibrium) simulation needs to integrate

$$\frac{\partial f_{\text{eq}}}{\partial t} + \mathbf{c} \cdot \frac{\partial f_{\text{eq}}}{\partial \mathbf{x}} = f \left(\frac{\partial W}{\partial t} + \mathbf{c} \cdot \frac{\partial W}{\partial \mathbf{x}} \right) + W \left(\frac{\partial f}{\partial t} + \mathbf{c} \cdot \frac{\partial f}{\partial \mathbf{x}} \right) = 0.$$

Using (3.1) we conclude that, during the advection step, the weights need to satisfy

$$\frac{\partial W}{\partial t} + \mathbf{c} \cdot \frac{\partial W}{\partial \mathbf{x}} = 0$$

or in other words, advect with their corresponding particle.

3.1.4. *Collision substep*

Next we present a derivation of weight update rules for the collision substep based on the hard-sphere form of the collision integral; extension to other collision models (*e.g.* variable hard sphere [7]) directly follows. An alternative derivation based on conditional probability arguments is presented in [2].

During the collision substep, the *auxiliary* simulation integrates the equilibrium version of (3.2), namely

$$\left[\frac{\partial f_{\text{eq}}}{\partial t} \right]_{\text{Coll}} = \frac{1}{2} \iiint (\delta'_1 + \delta'_2 - \delta_1 - \delta_2) W_1 W_2 f_1 f_2 c_r \sigma d\Omega d\mathbf{c}_1 d\mathbf{c}_2.$$

To facilitate the interpretation of this equation within the context of the DSMC collision algorithm, we rearrange it in the following form:

$$\begin{aligned} \left[\frac{\partial f_{\text{eq}}}{\partial t} \right]_{\text{Coll}} &= \frac{\text{MX}}{2} \iiint \left(\delta'_1 + \delta'_2 - \frac{\delta_1}{W_2} - \frac{\delta_2}{W_1} \right) W_1 W_2 f_1 f_2 \sigma \frac{c_r}{\text{MX}} d\Omega d\mathbf{c}_1 d\mathbf{c}_2 \\ &+ \frac{\text{MX}}{2} \iiint \left(-\delta_1 - \delta_2 + \frac{\delta_1}{W_2} + \frac{\delta_2}{W_1} \right) \left(\frac{c_r}{1 - \frac{c_r}{\text{MX}}} \right) W_1 W_2 f_1 f_2 \sigma \left(1 - \frac{c_r}{\text{MX}} \right) d\Omega d\mathbf{c}_1 d\mathbf{c}_2 \end{aligned} \quad (3.5)$$

where MX is an upper bound for $W c_r$. Using an importance sampling interpretation [4,5,18,19], the first term of this equation samples events occurring with probability $\hat{c}_r = c_r/\text{MX}$, while the second samples events occurring with probability $(1 - \hat{c}_r)$. In other words, provided collisions are accepted with probability \hat{c}_r (and thus rejected with probability $1 - \hat{c}_r$) in the DSMC calculation, (3.5) provides a means of connecting the events of collision acceptance and rejection routine in the “main” DSMC collision routine with weight evolution in the auxiliary equilibrium calculation.

To make this more concrete, consider a collision-candidate particle pair with velocities \mathbf{c}_1 and \mathbf{c}_2 and weights W_1 and W_2 , respectively. If the collision is accepted in DSMC, according to *the first term* in (3.5), a particle pair with velocities \mathbf{c}'_1 and \mathbf{c}'_2 and weights $W_1 W_2$ should be created, in addition to a pair of *negative* particles with velocities \mathbf{c}_1 and \mathbf{c}_2 and weights W_1 and W_2 , respectively. Note that, by design, the negative particles cancel the colliding particles and thus the collision proceeds by the update $\mathbf{c}_1 \rightarrow \mathbf{c}'_1$, $W_1 \rightarrow W_1 W_2$ and $\mathbf{c}_2 \rightarrow \mathbf{c}'_2$, $W_2 \rightarrow W_1 W_2$. Since the update $\mathbf{c}_1 \rightarrow \mathbf{c}'_1$, $\mathbf{c}_2 \rightarrow \mathbf{c}'_2$ is part of the original DSMC algorithm, we conclude that *if the collision is accepted in DSMC* the weight update is $W_1, W_2 \rightarrow W_1 W_2$. In the case of a rejected collision (in DSMC), *the second term* in (3.5) implies that $W_1 W_2 \hat{c}_r / (1 - \hat{c}_r)$ negative particles with velocities \mathbf{c}_1 and \mathbf{c}_2 , as well as $W_1 \hat{c}_r / (1 - \hat{c}_r)$ particles with velocity \mathbf{c}_1 and $W_2 \hat{c}_r / (1 - \hat{c}_r)$ particles with velocity \mathbf{c}_2 need to be created. Combining these with the colliding particles we obtain a net of $W_1 (1 - W_2 \hat{c}_r) / (1 - \hat{c}_r)$ at \mathbf{c}_1 and $W_2 (1 - W_1 \hat{c}_r) / (1 - \hat{c}_r)$ at \mathbf{c}_2 .

In summary, if the DSMC collision is accepted, the colliding particle weights are updated as $W_1, W_2 \rightarrow W_1 W_2$; if the DSMC collision is rejected, the candidate particle weights are updated as $W_1 \rightarrow W_1 (1 - W_2 \hat{c}_r) / (1 - \hat{c}_r)$ and $W_2 \rightarrow W_2 (1 - W_1 \hat{c}_r) / (1 - \hat{c}_r)$.

3.2. Stability considerations

Implementation of this algorithm reveals a numerical stability issue. Under the action of the above weight evolution rules, for times longer than a few mean-free-times between collisions ($\sqrt{\pi}\lambda/(2c_0)$), the variance of particle weights is observed to diverge (*i.e.* individual particle weights tend to either 0 or ∞); this ultimately results in loss of variance reduction. Despite significant differences in formulation, this behavior bears many similarities to the stability problems observed in other variance-reduced particle methods [5,11] – for example, the instability appears in collision dominated flows ($Kn < 1$), while collisionless and near collisionless ($Kn \gg 1$) calculations are stable in the presence of an accommodating boundary (for a more precise description of the stability limits of the present method see below). In our experience, this issue is one of the biggest challenges associated with this approach as well as previous variance reduction approaches [5,11].

We believe that, in the present formulation, this behavior is a result of the particular collision weight update rules used. In particular, the present implementation only satisfies (3.5) in an average sense (over many samples) and as a result, particle weights generated by this procedure are only meaningful when averaged over many samples. In collisionless flows this issue simply does not arise, while in near-collisionless flows ($Kn \gg 1$) in which n_w is known analytically, particles reach the walls (where their weights are reassigned) before their weights diverge.

Based on this interpretation, we have developed an approach for stabilizing the calculation. This approach uses Kernel Density Estimation (KDE) [25] to reconstruct the distribution function at the end of each timestep and is explained in the next sub-section.

3.2.1. Kernel density estimation of density functions

The KDE procedure allows us to reconstruct an estimate $\hat{f}(\mathbf{c})$ of $f(\mathbf{c})$ by only using N samples drawn from f . To proceed, let $\hat{f}(\mathbf{c})$ denote the “reconstructed” distribution function, obtained through $\hat{f}(\mathbf{c}) = \int K(\mathbf{c} - \mathbf{c}')f(\mathbf{c}')d\mathbf{c}' \simeq f(\mathbf{c})$, where $K(\mathbf{c} - \mathbf{c}')$ is an *appropriate* smoothing kernel. Using this approximation, from definition (2.2) it follows that we can define a reconstructed weight \hat{W}

$$\hat{W}(\mathbf{c}) = \frac{\int K(\mathbf{c} - \mathbf{c}')W(\mathbf{c}')f(\mathbf{c}')d\mathbf{c}'}{\int K(\mathbf{c} - \mathbf{c}')f(\mathbf{c}')d\mathbf{c}'} \tag{3.6}$$

In the work presented here we have used the normalized kernel

$$K(\mathbf{c} - \mathbf{c}') = \begin{cases} \frac{1}{4/3\pi\varepsilon^3c_0^3} & \text{if } |\mathbf{c} - \mathbf{c}'| < \varepsilon c_0 \\ 0 & \text{otherwise.} \end{cases}$$

Using importance sampling and (3.6), we obtain for each particle i

$$\hat{W}_i = \frac{1}{\|S_i\|} \sum_{k \in S_i} W_k. \tag{3.7}$$

Here, S_i denotes the set of particles that are within a sphere of radius εc_0 centered on particle i in velocity space; we denote the average number of such particles $\overline{\|S_i\|}$. Particles within S_i could be found using a K-D Tree [22]. In our implementation we find the set S_i using a non-uniform velocity space binning that is slightly faster and simpler to code.

This procedure introduces a new discretization parameter, ε , that affects the results of our simulation. A large ε means that we average over many particles which improves the stability of the calculation but also introduces numerical error (bias in the reconstruction of f , since $f = \hat{f}$ only when $\varepsilon \rightarrow 0$). We have found that the best tradeoff between stability and bias is to reconstruct the weights of only the particles that are accepted for collision³. Applying (3.7) to all particles, or all collision candidates (accepted for collision and

³This is a variation of the stabilization approach used by the authors in [3].

rejected for collision) improves the calculation stability albeit at the cost of additional numerical error (bias). A quantitative discussion of stability is given in Section 4.

3.3. Variable reference equilibrium state

Although the bias introduced by the KDE decreases as $\varepsilon \rightarrow 0$, the bias magnitude at a finite ε depends on the reference equilibrium state (f_{eq}) chosen during the collision substep. Our numerical experiments have shown that, while having a negligible effect on the stability of the method for a given $\|\overline{S_i}\|$ [2], choosing the reference equilibrium distribution to be the *local* equilibrium distribution (*i.e.* $f_{\text{eq}} = f_{\text{MB,loc}}$) results in a substantial reduction in the KDE bias. This is in qualitative agreement with previous work [19,23], which reports performance improvements when a variable equilibrium distribution is used.

This feature is *very important* because, for the same accuracy (bias) level, it allows the use of a larger ε resulting in a calculation that requires a smaller number of particles for stability (see Fig. 3). In fact, we find that, for the same bias, a local equilibrium reference state reduces the number of particles required for stability by about one order of magnitude at low Knudsen numbers and thus alleviates one of the most important limitations of the proposed method.

Fortunately, the use of a local equilibrium reference state can be limited to the collision substep thus requiring no modification to the remaining parts of the algorithm as described above. This is achieved by temporarily effecting a change of reference state ($f_{\text{eq}} \rightarrow f_{\text{MB,loc}}$) before performing the collision/KDE step and switching back to a global equilibrium state ($f_{\text{eq}} \rightarrow f_{\text{MB,0}}$) when the latter is finished.

The change of reference state, $f_{\text{eq}} \rightarrow f_{\text{MB,loc}}$ can be easily effected by defining a transformed set of weights W'_i such that

$$W'_i = \gamma_i W_i \tag{3.8}$$

where for each particle i

$$\gamma_i = \frac{f_{\text{MB,loc}}(\mathbf{c}_i)}{f_{\text{MB,0}}(\mathbf{c}_i)}.$$

Once the collision step is performed, the inverse change of reference state $f_{\text{eq}} \rightarrow f_{\text{MB,0}}$ can be achieved using

$$W'_i = \gamma_i^{-1} W_i. \tag{3.9}$$

3.4. Summary of the VRDSMC method

A flowchart summarizing the VRDSMC algorithm is shown in Figure 2; the original DSMC algorithm is shown in dark shading, while the additional steps leading to the VRDSMC algorithm are shown in light shading. The key substeps can be summarized as follows:

- (1) **Advection substep:** Identical to a DSMC simulation with weights following the particles. Initial conditions and wall interactions are processed using (2.2) with $f_{\text{eq}} = f_{\text{MB,0}}$.
- (2) **Change the equilibrium reference to local MB $f_{\text{MB,loc}}$:** Update all particle weights so that they correspond to the local MB reference state using (3.8). The local equilibrium distribution parameters are estimated using the variance reduced estimators (see step 5).
- (3) **Collision substep:** Standard DSMC procedure with MX chosen as an upper bound for Wc_r .
 - (a) **Accepted particles:** Accepted collision particles i, j are scattered using standard DSMC procedures [7]. Post-collision weights of particles i and j are updated using $W'_i, W'_j = \hat{W}_i \hat{W}_j$ where \hat{W} denotes weights estimated using a KDE with a kernel of diameter εc_0 at the particle's pre-collision velocity.
 - (b) **Rejected particles:** Weights are updated as in Section 3.1.4 ($W'_i = W_i (1 - W_j \hat{c}_r) / (1 - \hat{c}_r)$, $W'_j = W_j (1 - W_i \hat{c}_r) / (1 - \hat{c}_r)$); no KDE is used.
- (4) **Change equilibrium reference state back to global reference state $f_{\text{eq}} = f_{\text{MB,0}}$:** Reverse step 2 using (3.9); since the collision step conserves mass, momentum and energy, we use the local equilibrium parameters from step 2.

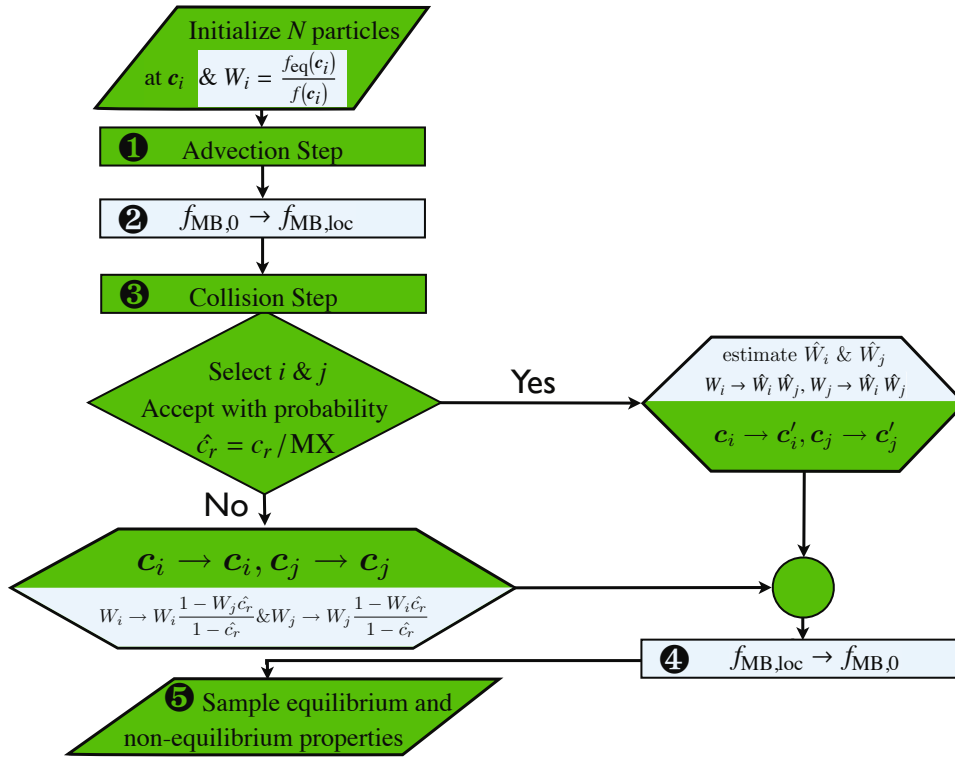


FIGURE 2. Flow chart of the VRDSMC method. Modifications to the original DSMC algorithm are highlighted in light color.

- (5) **Sampling:** DSMC sampling step using the *modified variance reduced estimators* of hydrodynamic properties (2.3) [2].

4. VALIDATION AND COMPUTATIONAL PERFORMANCE

4.1. Numerical stability

In this section we provide a quantitative discussion of the method stability. Because the stability of a calculation is primarily a function of $\|S_i\|$, while it is only a weak⁴ function of N_{cell} , it is more convenient to discuss stability by specifying the total number of particles per cell, N_{cell} , and the average number of particles in a sphere of radius εc_0 in velocity space, $\|S_i\|$.

Figure 3 shows the effect of $\|S_i\|$ on the calculation stability by showing the normalized average weight variance, $\overline{\sigma^2 \{W_i\}}$, as a function of $\|S_i\|$ for various Knudsen numbers, in a simple shear (Couette) flow; the weight

⁴Our numerical experiments show that a 3 order of magnitude increase in N_{cell} only results in a $\sim 20\%$ change in the average weight variance, $\overline{\sigma^2 \{W_i\}}$, in a typical one-dimensional flow problem.




FIGURE 3. Steady-state weight variance in a Couette flow as a function of $\overline{\|S_i\|}$ and Kn . The stability measure is normalized by the weight variance of a $Kn = 10$ simulation performed at $\overline{\|S_i\|} = 20$. High Kn flows require very small KDE intervention.

variance is averaged over the whole computational domain. Stability, defined as $\overline{\sigma^2 \{W_i\}} / \overline{\sigma^2 \{W_i\}}_{Kn=10, \overline{\|S_i\|}=20} \lesssim O(1)$, is a strong function of $\overline{\|S_i\|}$ and Kn . As expected, increasing $\overline{\|S_i\|}$ improves the calculation stability; moreover, flows characterized by $Kn > 1$ are stable in a wide range of values of $\overline{\|S_i\|}$, while flows with $Kn < 1$ typically require $\overline{\|S_i\|} \gg 1$ for stability. Because low KDE error (bias) requires small ε , the latter feature results in the most important weakness of the present method, namely the need for a large number of particles per cell for accurate simulations. This effect is significantly alleviated using a variable reference equilibrium state, as explained in Section 3.3. A discussion of the effect of ε (and $\overline{\|S_i\|}$) on the accuracy of simulation results is presented in Section 4.4. We also note that in cases where n_w is known analytically (not shown in this figure), simulations for $Kn \gg 1$ are unconditionally stable requiring no weight reconstruction (KDE).

4.2. Validation

In this section we use a number of one-dimensional problems to validate the proposed VRDSMC method⁵. Validation of a VRDSMC variant for a homogeneous relaxation problem can be found in a previous publication [3].

We consider a dilute gas of density n_0 between two fully accommodating, parallel plates (walls) a distance L apart ($Kn = \lambda/L$); the gas is initially at equilibrium at a temperature T_0 . Let x denote one of the directions parallel to the plates and y the direction normal to the plates. Our results are compared to DSMC solutions with identical discretization parameters (N_{cell} , cell size Δy , and timestep Δt). Unless otherwise stated, VRDSMC simulations were performed using a variable reference equilibrium state.

⁵A copy of the VRDSMC code is available upon request.




FIGURE 4. Steady state DSMC results (dots) and VRDSMC results (solid lines) for Couette flow. Left: $Kn = 1.0$ with $\overline{\|S_i\|} = 10$ and $N_{\text{cell}} = 200$; Right: $Kn = 0.1$ with $\overline{\|S_i\|} = 10$, $N_{\text{cell}} = 3000$.

4.2.1. Example 1: 1D steady state Couette flow problem

We first consider a simple Couette flow in which two (fully accommodating) plates move in the x direction with velocities $U_{\text{wall}} = \pm 0.085c_0$. The wall velocity magnitude was chosen to minimize the DSMC cost; as shown below and before [4,5,19], variance reduced techniques can resolve arbitrarily small flow disturbances at fixed cost.

Figure 4 shows the resulting flowfield for $Kn = 1$ and $Kn = 0.1$. The discrepancy between the VRDSMC and DSMC results is less than 1% for both cases. For $Kn = 1$, the VRDSMC calculation uses $\varepsilon \approx 0.57$ and $N_{\text{cell}} = 200$ which corresponds to $\overline{\|S_i\|} = 10$. For the same $\overline{\|S_i\|}$, we find that for $Kn = 0.1$ we are able to reproduce the DSMC solution to within $\sim 1\%$ using $\varepsilon \approx 0.23$, which corresponds to $N_{\text{cell}} = 3000$. An indication of the benefit associated with the variable equilibrium reference state can be obtained by noting that with $f_{\text{eq}} = f_{\text{MB},0}$, for $Kn = 0.1$ we are only able to reproduce the DSMC solution to within $\sim 1\%$ using $\varepsilon \approx 0.09$, which corresponds to $N_{\text{cell}} = 50\,000$.

We also note that low-speed ($U_{\text{wall}} \ll c_0$) Couette flows are isothermal and thus, to a good approximation, $n_w/n_0 = 1$. In our simulations, however, the more general approach outlined in Section 3.1.2 is used.

4.2.2. Example 2: 1D unsteady boundary heating problem

In this example we consider the transient response of the gas to an impulsive boundary temperature change. Specifically, at time $t = 0$, the wall temperatures impulsively change from T_0 to $T_0 \pm 0.033T_0$. In this problem the variable temperature induces density gradients and gas motion in the direction normal to the plates; as a result, n_w/n_0 is not known analytically and needs to be evaluated using the more general process outlined in Section 3.1.2, which necessitates weight reconstruction (albeit very mild), even for high Knudsen numbers.

Figure 5 shows the normalized results for the temperature (T), the heat-flux in the wall-normal direction (q_y), the density (ρ) and the y component of flow velocity (u_y) for $Kn = 10$; the heat flux is normalized by $q_{y,0} = \rho c_0^3$. For this large Knudsen number a minimal amount of KDE is required to keep the calculation stable; in this example we used $\varepsilon \approx 0.19$ and $\overline{\|S_i\|} = 1$ which amounts to $N_{\text{cell}} = 500$.

Figure 6 shows similar transient results for the $Kn = 1.0$ case, where VRDSMC can reproduce the DSMC solution within $\sim 1\%$ using $\varepsilon \approx 0.19$ and $\overline{\|S_i\|} = 3$ which amounts to $N_{\text{cell}} = 1500$.

4.3. Magnitude of variance reduction

Figure 7 shows a comparison of the relative statistical uncertainty in the flow velocity, σ/U_{wall} , achieved by DSMC and VRDSMC when simulating a steady Couette flow at $Kn = 1$. The figure shows that VRDSMC exhibits a constant relative statistical uncertainty for $U_{\text{wall}}/c_0 \ll 1$, as expected [4,5,19], and in sharp contrast

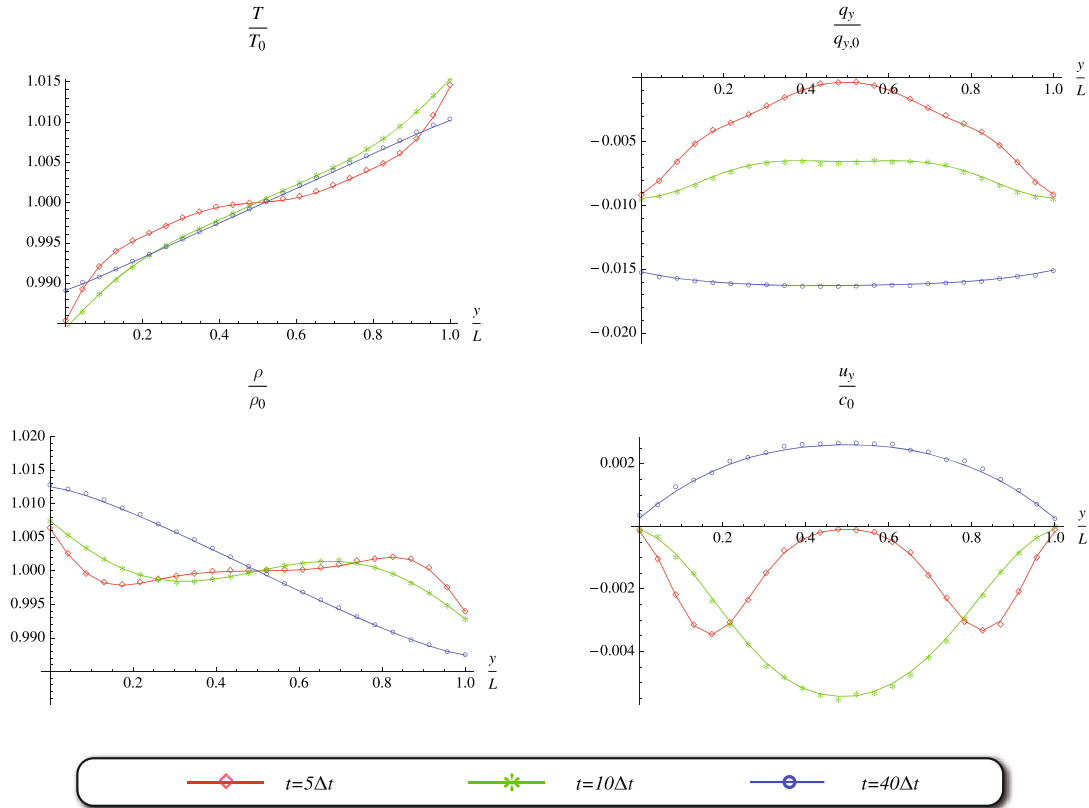


FIGURE 5. Transient results for an impulsive boundary temperature change problem for $Kn = 10$. Solid lines denote VRDSMC results while DSMC results are shown in dots. The snapshots shown correspond to $t = \{5, 10, 40\} * \Delta t$ where $\Delta t = \frac{1}{240} \sqrt{\pi} \lambda / (2c_0)$.

to DSMC whose statistical uncertainty for $U_{\text{wall}}/c_0 \ll 1$ is dominated by equilibrium fluctuations, resulting in $\sigma/U_{\text{wall}} \sim 1/U_{\text{wall}}$ in this limit [17]. We can thus conclude that VRDSMC provides a very considerable amount of variance reduction and thus computational benefit compared to DSMC for the same number of particles per cell.

4.4. Approximation error and limitations

Although VRDSMC is stable under a number of conditions, *e.g.* when n_w is known analytically and $Kn \gtrsim 3$, in general it requires a KDE stabilization step which introduces numerical error (bias). This error can be decreased by reducing the discretization parameter ε , at the expense of requiring a larger N_{cell} for stability (recall that $\|S_i\| \gg 1$ is required for small Kn). In other words, as already explained in the previous two sections, although *accurate* low- Kn calculations are feasible, they do require relatively large numbers of particles per cell ($\sim O(1000)$ for errors of the order of 1%). This requirement is practically non-existent in DSMC, which can provide very accurate solutions with as few as 100 particles per cell, independently of the Knudsen number for most problems of practical interest.

The requirement of large N_{cell} for accuracy is not very limiting in one-dimensional flows, but can be limiting in higher dimensions if the Knudsen number is small for all dimensions. On the other hand, these large numbers of particles do contribute towards reducing the already small statistical uncertainty of the calculation. In other words, if low statistical uncertainty calculations are required and a large number of particles (or ensembles)

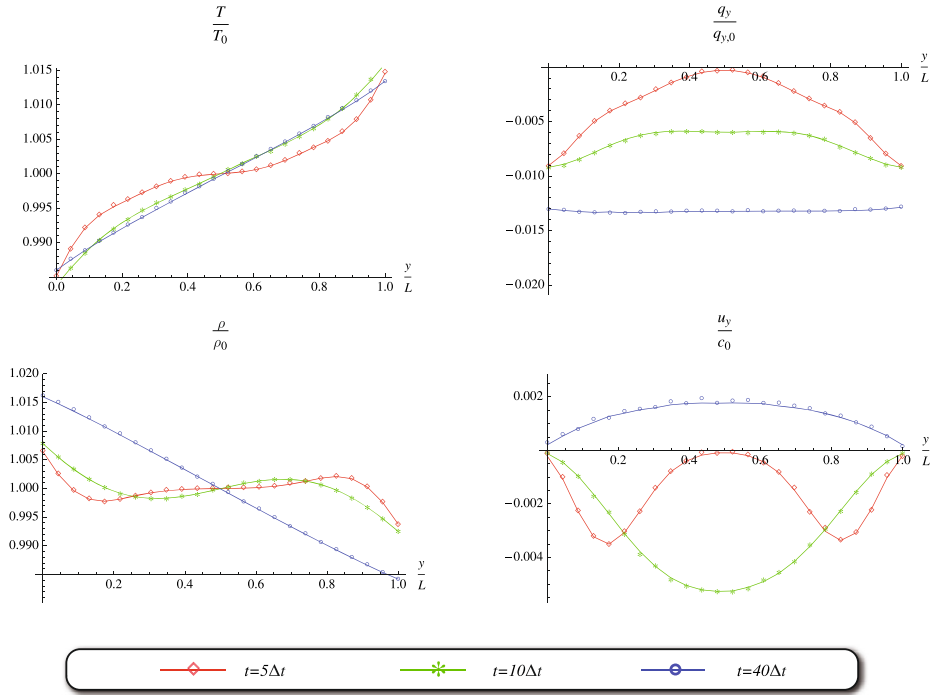


FIGURE 6. Transient results for an impulsive boundary temperature change problem for $Kn = 1.0$. Solid lines denote VRDSMC results with $\|S_i\| = 3$ and $N_{\text{cell}} = 1500$; DSMC results are shown in dots. The snapshots shown correspond to $t = \{5, 10, 40\} * \Delta t$ where $\Delta t = \frac{1}{24} \sqrt{\pi} \lambda / (2c_0)$.

were to be used in DSMC, then VRDSMC provides the full benefit shown in Figure 7. It should also be noted that our numerical experiments have shown that the numerical error associated with VRDSMC is not strongly affected by other discretization parameters (*e.g.* Δy) and thus provided ε remains small, much like DSMC [1,13,15], accurate solutions can be obtained with fairly coarse grids and thus total numbers of particles that are not excessively large. Furthermore, this moderate increase in N_{cell} is only practically limiting in applications characterized by $Kn \ll 1$, while most transition regime flows can be described to engineering accuracy with $N_{\text{cell}} = O(100)$, which is not substantially higher than the number of particles per cell required by DSMC.

5. SUMMARY AND CONCLUSIONS

We have presented a variance reduction method for direct Monte Carlo simulations of low-signal flows whose main ingredients are importance weights and Kernel Density Estimation. The method was validated using DSMC results for a number of different flows for $0.1 \leq Kn \leq 10$. Our numerical results show that, provided a sufficient number of particles is used, significant variance reduction is achieved with little additional discretization error.

More specifically, for $Kn \lesssim 3$ a basic tradeoff between accuracy (requiring a small ε) and stability (requiring a large $\|S_i\|$) exists. Consequently, a larger number of particles (compared to DSMC) may be needed to *accurately* simulate small Knudsen number flows. As we have shown above, provided a sufficient number of particles is used, all problems of interest ($0.1 \leq Kn \leq 10$) can be practically simulated with small numerical error. Overall, the extra computational cost of VRDSMC (due to the larger number of particles and the additional computation




FIGURE 7. Comparison of the relative statistical uncertainty (σ/U_{wall}) of the DSMC (squares) and VRDSMC (circles) methods. Results are for steady Couette flow at $Kn = 1.0$ for 500 particles per cell. The (theoretical) speedup scales with the square of the ratio of statistical uncertainties.

required per particle) is quite moderate [2], making the method very competitive with DSMC not only for low-signal flows, but also moderate signal flows (*e.g.* $Ma \sim 0.1$).

The major advantage of the proposed VRDSMC method compared to previous approaches is that it requires essentially no modification of the DSMC algorithm and introduces relatively little additional complexity. Moreover, it can be easily extended to other collision models and processes (*e.g.* chemical reactions). In fact, the variance reduction formulation presented is sufficiently general that extensions to other particle simulation methods may be possible.

Acknowledgements. This work has been funded, in part, by Saudi Aramco and, in part, by the Singapore-MIT Alliance. The authors would also like to thank Faisal Kashif, Colin Landon and Gregg Radtke for helpful discussions and Professor Tomasz Rychlik for helpful comments and suggestions.

REFERENCES

- [1] F.J. Alexander, A.L. Garcia and B.J. Alder, Cell size dependence of transport coefficients in stochastic particle algorithms. *Phys. Fluids* **10** (1998) 1540–1542.
- [2] H.A. Al-Mohssen, *An Excursion with the Boltzmann Equation at Low Speeds: Variance-Reduced DSMC*. Ph.D. Thesis, Massachusetts Institute of Technology, Dept. of Mechanical Engineering, Cambridge (2010).
- [3] H.A. Al-Mohssen and N.G. Hadjiconstantinou, Yet Another Variance Reduction Method for Direct Monte Carlo Simulations of Low-Signal Flows, in *26th International Symposium on Rarefied Gas Dynamics*, T. Abe Ed., AIP, Kyoto (2008) 257–262.
- [4] L.L. Baker and N.G. Hadjiconstantinou, Variance reduction for Monte Carlo solutions of the Boltzmann equation. *Phys. Fluids* **17** (2005) 051703.
- [5] L.L. Baker and N.G. Hadjiconstantinou, Variance-reduced particle methods for solving the Boltzmann equation. *J. Comput. Theor. Nanosci.* **5** (2008) 165–174.
- [6] L.L. Baker and N.G. Hadjiconstantinou, Variance-reduced Monte Carlo solutions of the Boltzmann equation for low-speed gas flows: A discontinuous Galerkin formulation. *Int. J. Numer. Methods Fluids* **58** (2008) 381–402.
- [7] G.A. Bird, *Molecular Gas Dynamics and the Direct Simulation of Gas Flows*. Clarendon Press (1994).
- [8] C. Cercignani, *The Boltzmann equation and its applications*. Springer-Verlag (1988).

- [9] C. Cercignani, *Mathematical Methods in Kinetic Theory*. Plenum Press (1990).
- [10] C. Cercignani, *Slow Rarefied Flows: Theory and Application to Micro-Electro-Mechanical Systems*. Springer (2006).
- [11] J. Chun and D.L. Koch, A direct simulation Monte Carlo method for rarefied gas flows in the limit of small Mach number. *Phys. Fluids* **17** (2005) 107107.
- [12] A. Doucet and X. Wang, Monte Carlo methods for signal processing: a review in the statistical signal processing context. *IEEE Signal Process. Mag.* **22** (2005) 152–170.
- [13] A.L. Garcia and W. Wagner, Time step truncation error in direct simulation Monte Carlo. *Phys. Fluids* **12** (2000) 2621–2633.
- [14] P. Glasserman, *Monte Carlo Methods in Financial Engineering*. Springer (2004).
- [15] N.G. Hadjiconstantinou, Analysis of discretization in the direct simulation Monte Carlo. *Phys. Fluids* **12** (2000) 2634–2638.
- [16] N.G. Hadjiconstantinou, The limits of Navier-Stokes theory and kinetic extensions for describing small-scale gaseous hydrodynamics. *Phys. Fluids* **18** (2006) 111301.
- [17] N.G. Hadjiconstantinou, A.L. Garcia, M.Z. Bazant and G. He, Statistical error in particle simulations of hydrodynamic phenomena. *J. Comput. Phys.* **187** (2003) 274–297.
- [18] T.M.M. Homolle and N.G. Hadjiconstantinou, Low-variance deviational simulation Monte Carlo. *Phys. Fluids* **19** (2007) 041701.
- [19] T.M.M. Homolle and N.G. Hadjiconstantinou, A low-variance deviational simulation Monte Carlo for the Boltzmann equation. *J. Comput. Phys.* **226** (2007) 2341–2358.
- [20] C.D. Landon, *Weighted Particle Variance Reduction of Direct Simulation Monte Carlo for the Bhatnagar-Gross-Krook Collision Operator*. M.S. Thesis, Massachusetts Institute of Technology, Dept. of Mechanical Engineering, Cambridge (2010).
- [21] H.C. Ottinger, B.H.A.A. van den Brule and M.A. Hulsen, Brownian configuration fields and variance reduced CONNFESSIT. *J. Non-Newton. Fluid Mech.* **70** (1997) 255–261.
- [22] W.H. Press, S.A. Teukolsky, W.T. Vetterling and B.P. Flannery, *Numerical Recipes*. Cambridge University Press (2007).
- [23] G.A. Radtke and N.G. Hadjiconstantinou, Variance-reduced particle simulation of the Boltzmann transport equation in the relaxation-time approximation. *Phys. Rev. E* **79** (2009) 056711.
- [24] R.Y. Rubinstein, *Simulation and the Monte Carlo Method*. Wiley (1981).
- [25] D.W. Scott, *Multivariate Density Estimation*. John Wiley & Sons (1992).
- [26] Y. Sone, *Kinetic Theory and Fluid Dynamics*. Birkhauser (2002).
- [27] W. Wagner, A convergence proof for Bird’s direct simulation Monte Carlo method for the Boltzmann equation. *J. Stat. Phys.* **66** (1992) 1011–1044.
- [28] W. Wagner, Deviational Particle Monte Carlo for the Boltzmann Equation. *Monte Carlo Methods Appl.* **14** (2008) 191–268.