Variance-reduced particle simulation of the Boltzmann transport equation in the relaxation-time approximation

Gregg A. Radtke* and Nicolas G. Hadjiconstantinou†
Department of Mechanical Engineering, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139, USA
(Received 7 February 2009; revised manuscript received 15 April 2009; published 26 May 2009)

We present an efficient variance-reduced particle simulation technique for solving the linearized Boltzmann transport equation in the relaxation-time approximation used for phonon, electron, and radiative transport, as well as for kinetic gas flows. The variance reduction is achieved by simulating only the deviation from equilibrium. We show that in the limit of small deviation from equilibrium of interest here, the proposed formulation achieves low relative statistical uncertainty that is also independent of the magnitude of the deviation from equilibrium, in stark contrast to standard particle simulation methods. Our results demonstrate that a space-dependent equilibrium distribution improves the variance reduction achieved, especially in the collision-dominated regime where local equilibrium conditions prevail. We also show that by exploiting the physics of relaxation to equilibrium inherent in the relaxation-time approximation, a very simple collision algorithm with a clear physical interpretation can be formulated.

DOI: 10.1103/PhysRevE.79.056711 PACS number(s): 05.10.Ln, 02.70.Ns, 51.10.+y

I. INTRODUCTION

Solution methods for the Boltzmann transport equation have received renewed attention in connection with the recent focus on small-scale devices, where particle-mediated (molecule, phonon, electron, or photon) transport occurs over length scales that are on the order of or smaller than the mean free path [1].

The relaxation-time approximation is a widely used simplification to the Boltzmann collision integral that provides a reasonably accurate description of a number of collisional phenomena. It is commonly used for rarefied gas flows [2–5]—where it is referred to as the Bhatnagar-Gross-Krook (BGK) model—as well as for electron [6,7], phonon [1,8,9], and radiative [5] transport. One of the disadvantages of the BGK model is that it yields a Prandtl number of unity, making the study of simultaneous heat and momentum transfer problems [2] difficult. To address this deficiency, a modified relaxation model with a variable Prandtl number, known as the ellipsoidal statistical model [2], has been developed.

Solving the Boltzmann equation, in general, remains a formidable task due to the high dimensionality of the distribution function and the complexities associated with the collision integral. Particle methods are the prevalent solution method for simulating the Boltzmann equation due to their natural treatment of the advection operator, ability to capture traveling discontinuities in the velocity distribution function, and straightforward application to complicated geometries. For dilute gas flows, the most widely used such method is the direct simulation Monte Carlo (DSMC) method [10].

When simulating systems close to equilibrium (e.g., low-Mach-number dilute gas flows, electron transport in the low-field limit, etc.), traditional particle methods (such as DSMC) exhibit an unacceptably low signal-to-noise ratio. For example, using equilibrium statistical mechanics [11], one can show that the relative statistical uncertainty in any flow velocity component is given by

$$\frac{\sigma_{u_i}}{u_0} = \frac{1}{\sqrt{MN_{\text{cell}}}} \frac{1}{Ma\sqrt{\gamma}}, \tag{1}$$

where σ_{u_i} is the standard deviation in component i of the flow velocity, u_0 is a reference velocity, M is the number of independent samples, N_{cell} is the number of particles per computational cell, Ma is the local Mach number (based on u_0), and γ is the ratio of specific heats. Thus, as Ma becomes small (typical of gas flows in microdevices), the computational cost becomes prohibitive.

Most of the previous approaches to reducing the statistical uncertainty associated with particle simulations (i.e., variance reduction) have focused on the original Boltzmann (hard-sphere-type) collision model. Baker and Hadjiconstantinou [12] originally proposed the concept of variance reduction by simulating only the deviation from equilibrium within the context of the original Boltzmann equation. Chun and Koch [13] proposed a variance-reduced particle scheme for the linearized hard-sphere model by introducing ghost particles and particle weights to simulate the deviation from a fixed equilibrium distribution; this approach required a particle cancellation scheme (leading to a discretized velocity space) in order to prevent an unbounded increase in the number of simulated particles. Baker and Hadjiconstantinou [14,15] proposed an alternative variance-reduced particle scheme in terms of deviational particles (particles representing the deviation from equilibrium) that may be positive or negative; this approach again required a particle cancellation scheme for stability.

Building on the ideas of Baker and Hadjiconstantinou [14,15], Homolle and Hadjiconstantinou [16,17] proposed a particle scheme which uses a special formulation of the hard-sphere collision operator to *achieve stability without particle cancellation*. This scheme is known as low-variance deviational simulation Monte Carlo (LVDSMC). Further theoretical analysis of deviational methods can be found in [18].

*radtke@mit.edu
†ngh@mit.edu

In this paper, we present a deviational method for treating the relaxation-time approximation which, as discussed above, is widely used within the applied physics community. This work takes advantage of the similarities between the BGK and the LVDSMC collision operators to develop a stable algorithm (not requiring a particle cancellation scheme). As explained in detail below, in the method proposed here we simulate the deviation from a space-dependent equilibrium distribution; this results in a substantially different algorithm compared to methods which simulate the deviation from a spatially homogeneous equilibrium distribution, such as the one recently proposed by Hadjiconstantinou *et al.* [19]. In particular, we show that under linearized conditions, the resulting collision algorithm is very simple and efficient. This algorithm is also in line with the physical interpretation of the collision integral; specifically, the action of the collision integral can be implemented as a relaxation of the underlying equilibrium distribution toward the local equilibrium distribution and removal of deviational particles from the simulation. In terms of variance reduction achieved, our results suggest that use of a spatially dependent equilibrium distribution improves variance reduction, especially at low Knudsen numbers where local equilibrium conditions prevail. In this paper we also demonstrate that deviational methods can be appropriately symmetrized to achieve second-order convergence in time; additionally, we introduce the use of the generalized ratio-of-uniforms method in deviational methods as a means of increasing the efficiency of sampling distributions.

Pareschi [20] and Dimarco and Pareschi [21] also studied the relaxation-time approximation and developed a discrete velocity approach which decomposes the distribution function into a convex combination of an equilibrium and a nonequilibrium distribution, with the latter represented by positive particles. These studies [20,21], which were primarily focused on the Euler hydrodynamic limit, studied the effect of a time-varying equilibrium distribution, which unfortunately becomes nonequilibrium upon action of the advection operator, thus requiring reconstruction.

A brief outline of the paper is as follows: background material for the relaxation-time approximation and variance reduction is introduced in Sec. II; our particle method for the linearized Boltzmann equation is developed in Sec. III and summarized in Sec. IV, where the order of time convergence is also discussed. Several validation test cases are considered, and the results are presented in Sec. V. We conclude in Sec. VI by discussing the features and differences of the present method to other methods. Finally, the application of the generalized ratio-of-uniforms method, an efficient technique for sampling distributions, is outlined in the Appendix.

Although our method is applicable to any physical system which can be modeled by the Boltzmann equation in the linearized relaxation-time approximation, for simplicity we will develop the theory based on a dilute gas flow model; extension to other systems of interest is straightforward.

II. BACKGROUND

The Boltzmann equation in the relaxation-time approximation is given by

$$\frac{\partial f}{\partial t} + \boldsymbol{c}\cdot\frac{\partial f}{\partial \boldsymbol{x}} + \boldsymbol{a}\cdot\frac{\partial f}{\partial \boldsymbol{c}} = -\frac{f - f^{\mathrm{loc}}}{t_{\mathrm{r}}}, \tag{2}$$

where $f(\boldsymbol{x},\boldsymbol{c},t)$ is the particle distribution function, t is the time, $\boldsymbol{x}$ is the spatial coordinate, $\boldsymbol{c}$ is the particle velocity, $\boldsymbol{a}$ is the external force field per unit mass, and t_{r} is the relaxation time. For the sake of simplicity, we will assume that $\boldsymbol{a}$ is zero; extension to $\boldsymbol{a}\neq 0$ directly follows. The local equilibrium particle density function is a Maxwell-Boltzmann distribution,

$$f^{\mathrm{loc}} = \frac{n}{(2\pi RT)^{3/2}}\exp\left(-\frac{\|\boldsymbol{c}-\boldsymbol{u}\|^2}{2RT}\right), \tag{3}$$

in terms of the local density n, velocity $\boldsymbol{u}$, and temperature T.

For convenience, this equation is made dimensionless based on the reference number density n_0, temperature T_0, most probable speed $c_0=\sqrt{2RT_0}$, and a characteristic length scale L. The dimensionless variables are therefore $\hat{f}=c_0^3 f/n_0$, $\hat{t}=c_0 t/L$, $\hat{\boldsymbol{c}}=\boldsymbol{c}/c_0$, and $\hat{\boldsymbol{x}}=\boldsymbol{x}/L$, yielding

$$\frac{\partial \hat{f}}{\partial \hat{t}} + \hat{\boldsymbol{c}}\cdot\frac{\partial \hat{f}}{\partial \hat{\boldsymbol{x}}} = -\frac{t_{\mathrm{r},0}}{t_{\mathrm{r}}}\frac{\hat{f}-\hat{f}^{\mathrm{loc}}}{k}. \tag{4}$$

Here, $t_{\mathrm{r},0}$ is the relaxation time at reference conditions and $k=c_0 t_{\mathrm{r},0}/L$ is a modified Knudsen number, which is related to the standard definition of Kn (ratio of the mean free path to L) by $k=\frac{\sqrt{\pi}}{2}Kn$. The local equilibrium distribution in dimensionless form is $\hat{f}^{\mathrm{loc}}=\hat{n}(\pi\hat{T})^{-3/2}e^{-\|\hat{\boldsymbol{c}}-\hat{\boldsymbol{u}}\|^2/\hat{T}}$, which is now defined in terms of the dimensionless number density $\hat{n}=n/n_0$, mean velocity $\hat{\boldsymbol{u}}=\boldsymbol{u}/c_0$, and temperature $\hat{T}=T/T_0$.

Variance reduction is achieved by simulating the particle distribution function $\hat{f}=\hat{f}^{\mathrm{MB}}+\hat{f}^{\mathrm{d}}$ in two parts: an equilibrium (Maxwell-Boltzmann) distribution $\hat{f}^{\mathrm{MB}}$ and the deviation from the equilibrium distribution $\hat{f}^{\mathrm{d}}$. The distribution $\hat{f}^{\mathrm{MB}}=\hat{n}_{\mathrm{MB}}(\pi\hat{T}_{\mathrm{MB}})^{-3/2}e^{-\|\hat{\boldsymbol{c}}-\hat{\boldsymbol{u}}_{\mathrm{MB}}\|^2/\hat{T}_{\mathrm{MB}}}$ is characterized by the parameters $(\hat{n}_{\mathrm{MB}},\hat{\boldsymbol{u}}_{\mathrm{MB}},\hat{T}_{\mathrm{MB}})$ which may be spatially variable; the distribution $\hat{f}^{\mathrm{d}}$, which represents a deviation from the equilibrium distribution, is characterized by signed particles. Evolution rules for $\hat{f}^{\mathrm{MB}}$ and $\hat{f}^{\mathrm{d}}$ are derived below.

III. PARTICLE METHOD FOR THE LINEARIZED BOLTZMANN EQUATION

The particle simulation can be initialized after selecting an initial equilibrium distribution $\hat{f}^{\mathrm{MB}}$. Although this choice is in principle arbitrary, an appropriately chosen initial $\hat{f}^{\mathrm{MB}}$ reduces the number of particles required to simulate the deviation therefrom and thus minimizes the computational cost. If the system is initially at equilibrium, it is typically wise to identify $\hat{f}^{\mathrm{MB}}$ with the initial equilibrium distribution; this requires no deviational particles to describe the system initial state (nonequilibrium then originates at the walls or source terms which generate deviational particles as explained in later sections). If the system initial state is nonequilibrium, then an appropriate initial $\hat{f}^{\mathrm{MB}}$ is a nearby equilibrium state;

the deviational distribution is then generated by sampling $\hat{f}^{\mathrm{d}}=\hat{f}-\hat{f}^{\mathrm{MB}}$.

As in DSMC, a parameter known as the *effective number*, N_{eff}, is used to relate the number of simulated particles to physical particles. However, in contrast to DSMC, where the number of simulated particles is solely determined by N_{eff} and the number of physical particles, in the present method the number of simulated particles depends on the departure from equilibrium since particles are used to simulate $\hat{f}^{\mathrm{d}}$, a fraction of the total distribution. Formally, we represent the deviational distribution as $\hat{f}^{\mathrm{d}}=N_{\mathrm{eff}}\Sigma_p s_p \delta(\hat{\boldsymbol{x}}-\hat{\boldsymbol{x}}_p)\delta(\hat{\boldsymbol{c}}-\hat{\boldsymbol{c}}_p)$, where $\hat{\boldsymbol{x}}_p$ are the particle positions, $\hat{\boldsymbol{c}}_p$ are the particle velocities, and $s_p(=\pm 1)$ are the particle signs.

Typical of particle methods, the simulation proceeds by using a splitting method, whereby integration in time with time step $\delta\hat{t}$ is split into a collisionless advection substep and a collision substep. During the collision substep, the effect of collisions is simulated in a spatially homogeneous fashion within spatial cells of characteristic linear dimension Δx. With this in mind, in our simulation we take the parameters $(\hat{n}^{\mathrm{MB}},\hat{\boldsymbol{u}}^{\mathrm{MB}},\hat{T}^{\mathrm{MB}})$ to be constant within each cell. This choice also impacts the advection substep as explained below.

The collision and advection substeps are now discussed in more detail.

A. Collision substep

Introducing $\hat{f}=\hat{f}^{\mathrm{MB}}+\hat{f}^{\mathrm{d}}$ into the right-hand side of Eq. (4) yields (after appropriate linearization)

$$\left[\frac{\partial\hat{f}}{\partial\hat{t}}\right]_{coll}=\frac{\hat{f}^{\mathrm{loc}}-\hat{f}^{\mathrm{MB}}}{k}-\frac{\hat{f}^{\mathrm{d}}}{k}. \tag{5}$$

As shown in a previous method by Hadjiconstantinou *et al.* [19], the collision step may be simulated by treating $\hat{f}^{\mathrm{MB}}$ as constant (e.g., absolute equilibrium) and treating the first and second terms on the right-hand side as a source and a sink term for deviational particles, respectively. Here, however, we proceed by allowing $\hat{f}^{\mathrm{MB}}$ to vary as a function of space [18] and be updated once every time step by introducing an arbitrary change $\delta\hat{f}^{\mathrm{MB}}$. In other words, we write

$$\left[\frac{\partial\hat{f}}{\partial\hat{t}}\right]_{coll}\delta\hat{t}=\left[\frac{\delta\hat{t}}{k}(\hat{f}^{\mathrm{loc}}-\hat{f}^{\mathrm{MB}})-\delta\hat{f}^{\mathrm{MB}}\right]+\delta\hat{f}^{\mathrm{MB}}-\frac{\delta\hat{t}}{k}\hat{f}^{d}. \tag{6}$$

In the linearized regime, it is possible to make the bracketed term in the right-hand side (above) equal to zero by choosing $\delta\hat{f}^{\mathrm{MB}}$ appropriately; this change in $\hat{f}^{\mathrm{MB}}$ is accomplished by updating the properties $(\hat{n}_{\mathrm{MB}},\hat{\boldsymbol{u}}_{\mathrm{MB}},\hat{T}_{\mathrm{MB}})$. The final term is implemented by deleting deviational particles with probability $\delta\hat{t}/k$.

We now proceed to find the appropriate form of $\delta\hat{f}^{\mathrm{MB}}$ that makes the first term on the right-hand side of Eq. (6) equal to zero by expanding $\hat{f}^{\mathrm{MB}}$ and $\hat{f}^{\mathrm{loc}}$ about a fixed equilibrium distribution $\hat{f}^0=\pi^{-3/2}e^{-\hat{c}^2}$, where $\hat{c}=\|\hat{\boldsymbol{c}}\|$; the result for $\hat{f}^{\mathrm{MB}}$ is

$$\hat{f}^{\mathrm{MB}}=\left[1+\omega_{\mathrm{MB}}+2\hat{\boldsymbol{c}}\cdot\hat{\boldsymbol{u}}_{\mathrm{MB}}+\left(\hat{c}^2-\tfrac{3}{2}\right)\tau_{\mathrm{MB}}\right]\hat{f}^0, \tag{7}$$

while a similar result is obtained for $\hat{f}^{\mathrm{loc}}$. Here, $\omega=\hat{n}-1$ and $\tau=\hat{T}-1$ are the perturbations from the equilibrium density and temperature, respectively. From Eq. (7) we obtain

$$\delta\hat{f}^{\mathrm{MB}}=\left[\delta\omega_{\mathrm{MB}}+2\hat{\boldsymbol{c}}\cdot\delta\hat{\boldsymbol{u}}_{\mathrm{MB}}+\left(\hat{c}^2-\tfrac{3}{2}\right)\delta\tau_{\mathrm{MB}}\right]\hat{f}^0, \tag{8}$$

which can be used to show that, to linear order, the first term on the right-hand side of Eq. (6) is identically zero when the Maxwell-Boltzmann properties are updated according to

$$\delta\omega_{\mathrm{MB}}=\frac{\delta\hat{t}}{k}(\omega-\omega_{\mathrm{MB}}), \tag{9}$$

$$\delta\hat{\boldsymbol{u}}_{\mathrm{MB}}=\frac{\delta\hat{t}}{k}(\hat{\boldsymbol{u}}-\hat{\boldsymbol{u}}_{\mathrm{MB}}), \tag{10}$$

$$\delta\tau_{\mathrm{MB}}=\frac{\delta\hat{t}}{k}(\tau-\tau_{\mathrm{MB}}). \tag{11}$$

Note that this step does not require $\hat{f}^{\mathrm{MB}}$ to be time dependent but is rather understood as a type of change of basis which is performed once every time step. For further discussion of this point the reader is referred to Ref. [18].

The two nontrivial tasks in performing the collision step—changing the Maxwell-Boltzmann properties and deleting deviational particles—mirror the physics inherent in the relaxation-time approximation; specifically, the equilibrium part of the simulated distribution $\hat{f}^{\mathrm{MB}}$ approaches the local equilibrium distribution $\hat{f}^{\mathrm{loc}}$, while the deviational part $\hat{f}^{\mathrm{d}}$ decays to zero. Solving the nonlinear Boltzmann equation (in the relaxation-time approximation) can be performed by generating additional particles at each time step [22].

B. Advection substep

In the advection substep, we numerically implement

$$\frac{\partial\hat{f}}{\partial\hat{t}}+\hat{\boldsymbol{c}}\cdot\frac{\partial\hat{f}}{\partial\hat{\boldsymbol{x}}}=0, \tag{12}$$

which for a spatially dependent $\hat{f}^{\mathrm{MB}}$ gives

$$\frac{\partial\hat{f}^{\mathrm{d}}}{\partial\hat{t}}+\hat{\boldsymbol{c}}\cdot\frac{\partial\hat{f}^{\mathrm{d}}}{\partial\hat{\boldsymbol{x}}}=-\hat{\boldsymbol{c}}\cdot\frac{\partial\hat{f}^{\mathrm{MB}}}{\partial\hat{\boldsymbol{x}}}. \tag{13}$$

Simple advection of deviational particles, as is performed in ordinary DSMC, provides a homogeneous solution of Eq. (13), while additional particles must be generated from the cell interfaces at every time step to satisfy the inhomogeneous term. According to Refs. [16–18], the additional particles are generated from a difference in f^{MB} fluxal distributions: i.e.,

$$\mathcal{F}_{\text{adv}}\delta\hat{t} = \left[\frac{\partial \hat{f}^{\text{d}}}{\partial \hat{t}}\right]_{\text{adv,gen}} S_n\delta\hat{t} = S_n\hat{c}_n\delta\hat{t}(\hat{f}_-^{\text{MB}} - \hat{f}_+^{\text{MB}}), \quad (14)$$

where S_n is the (dimensionless) surface area of the cell interface, $\hat{c}_n = \hat{\boldsymbol{c}}\cdot\boldsymbol{n}$ is the normal component of the particle velocity, $\hat{f}_\pm^{\text{MB}}$ are the Maxwell-Boltzmann distributions for adjacent cells, and $\boldsymbol{n}$ is the cell-interface surface-normal pointing from $\hat{f}_-^{\text{MB}}$ to $\hat{f}_+^{\text{MB}}$.

Using Eq. (7), distribution (14) can be written as

$$\mathcal{F}_{\text{adv}}\delta\hat{t} = S_n\hat{c}_n\delta\hat{t}\left[(\omega_{\text{MB}-} - \omega_{\text{MB}+}) + 2\hat{\boldsymbol{c}}\cdot(\hat{\boldsymbol{u}}_{\text{MB}-} - \hat{\boldsymbol{u}}_{\text{MB}+}) + \left(\hat{c}^2 - \frac{3}{2}\right)(\tau_{\text{MB}-} - \tau_{\text{MB}+})\right]\hat{f}^0, \quad (15)$$

Implementation consists of generating $\int_{\Re^3} d^3\hat{\boldsymbol{c}}\mathcal{F}_{\text{adv}}\delta\hat{t}/N_{\text{eff}}$ deviational particles sampled from the distribution $\mathcal{F}_{\text{adv}}$ and advecting them for a uniformly distributed random fraction of the time step away from the interface at which they are created. These particles are efficiently generated using the ratio-of-uniforms method [23] as outlined in the Appendix. For isothermal [24] flow parallel to the boundary surface, the generation term reduces to a particularly simple form: $\sim\hat{c}_x\hat{c}_z\hat{f}^0$ where $\hat{c}_x$ is normal to the boundary and $\hat{c}_z$ is in the flow direction, resulting in a product of a Gaussian and analytically invertible distributions.

C. Boundary conditions

The boundary conditions are enforced according to the Maxwell accommodation model [4,18],

$$\hat{f}(\hat{\boldsymbol{c}}) = (1-\alpha)\hat{f}(\hat{\boldsymbol{c}} - 2[(\hat{\boldsymbol{c}} - \hat{\boldsymbol{u}}_{\text{b}})\cdot\boldsymbol{n}]\boldsymbol{n}) + \alpha\hat{f}^{\text{b}}(\hat{\boldsymbol{c}}) \quad (16)$$

for $(\hat{\boldsymbol{c}} - \hat{\boldsymbol{u}}_{\text{b}})\cdot\boldsymbol{n} > 0$, where in the interest of simplicity, the explicit time dependence of the distribution function is suppressed. In Eq. (16), $\hat{f}$ is evaluated at the boundary; $\boldsymbol{n}$ is the normal vector pointing into the gas. The (linearized) boundary distribution is

$$\hat{f}^{\text{b}} = \left[1 + \omega_{\text{b}}^{(1)} + \omega_{\text{b}}^{(2)} + 2\hat{\boldsymbol{c}}\cdot\hat{\boldsymbol{u}}_{\text{b}} + \left(\hat{c}^2 - \tfrac{3}{2}\right)\tau_{\text{b}}\right]\hat{f}^0, \quad (17)$$

where $\hat{\boldsymbol{u}}_{\text{b}}$ is the boundary velocity and τ_{b} is the temperature perturbation of the boundary. The number density perturbation consists of two parts: $\omega_{\text{b}}^{(1)}$ and $\omega_{\text{b}}^{(2)}$; both are discussed below. For simplicity, we assume that the wall normal component of the boundary velocity is zero ($\hat{\boldsymbol{u}}_{\text{b}}\cdot\boldsymbol{n} = 0$). Mass conservation is enforced by balancing the inward and outward particle fluxes, i.e.,

$$\int_{\hat{c}_n>0} d^3\hat{\boldsymbol{c}}\hat{c}_n\hat{f}^{\text{b}} = -\int_{\hat{c}_n<0} d^3\hat{\boldsymbol{c}}\hat{c}_n\hat{f}, \quad (18)$$

which, upon substitution of $\hat{f} = \hat{f}^{\text{MB}} + \hat{f}^{\text{d}}$ and Eq. (17), may be written in two parts as

$$\int_{\hat{c}_n>0} d^3\hat{\boldsymbol{c}}\hat{c}_n\left[\omega_{\text{b}}^{(2)} + 2\hat{\boldsymbol{c}}\cdot\hat{\boldsymbol{u}}_{\text{b}} + \left(\hat{c}^2 - \frac{3}{2}\right)\tau_{\text{b}}\right]\hat{f}^0 = -\int_{\hat{c}_n<0} d^3\hat{\boldsymbol{c}}\hat{c}_n\left[\omega_{\text{MB}} + 2\hat{\boldsymbol{c}}\cdot\hat{\boldsymbol{u}}_{\text{MB}} + \left(\hat{c}^2 - \frac{3}{2}\right)\tau_{\text{MB}}\right]\hat{f}^0 \quad (19)$$

and

$$\int_{\hat{c}_n>0} d^3\hat{\boldsymbol{c}}\hat{c}_n\hat{f}^0(1 + \omega_{\text{b}}^{(1)}) = -\int_{\hat{c}_n<0} d^3\hat{\boldsymbol{c}}\hat{c}_n\hat{f}^d. \quad (20)$$

Introducing $\hat{f} = \hat{f}^{\text{MB}} + \hat{f}^{\text{d}}$ and Eq. (17) into Eq. (16) yields

$$\hat{f}^{\text{d}}(\hat{\boldsymbol{c}}) = (1-\alpha)\hat{f}^{\text{d}}(\hat{\boldsymbol{c}} - 2\hat{c}_n\boldsymbol{n}) + \alpha(1 + \omega_{\text{b}}^{(1)})\hat{f}^0(\hat{\boldsymbol{c}}) + \alpha\left[\omega_{\text{b}}^{(2)} - \omega_{\text{MB}} + 2\hat{\boldsymbol{c}}\cdot(\hat{\boldsymbol{u}}_{\text{b}} - \hat{\boldsymbol{u}}_{\text{MB}}) + \left(\hat{c}^2 - \tfrac{3}{2}\right)(\tau_{\text{b}} - \tau_{\text{MB}})\right]\hat{f}^0(\hat{\boldsymbol{c}}) \quad (21)$$

for $\hat{c}_n > 0$, where the properties ω_{MB}, $\hat{\boldsymbol{u}}_{\text{MB}}$, and τ_{MB} are from the cell adjacent to the boundary. The first and second terms on the right-hand side correspond to the ordinary DSMC procedures—i.e., deviational particles are specularly reflected with probability $1-\alpha$ and diffusely reflected with probability α—while the third term requires additional particles to be generated at the boundary. Specular reflections are accomplished by changing the sign of the normal component of the particle velocity according to $\hat{\boldsymbol{c}} \to \hat{\boldsymbol{c}} - 2\hat{c}_n\boldsymbol{n}$. Diffusely reflected particles are "absorbed" by the wall and redrawn from the fluxal wall distribution, where the constant $\omega_{\text{b}}^{(1)}$ is automatically determined [see Eq. (20)] by ensuring that the same net mass of particles that was absorbed is also emitted. Specifically, if N^+ positive and N^- negative particles are diffusely reflected from a boundary surface during a time step, we emit $|N^+ - N^-|$ particles drawn from $\hat{c}_n\hat{f}^0$ for $\hat{c}_n > 0$ with sign $\text{sgn}(N^+ - N^-)$, which leads to a reduced number of particles by canceling pairs of positive and negative particles striking the boundary [16,17].

The third term in Eq. (21) is implemented by generating additional particles in a similar manner to the advection routine, namely, particles are generated by sampling the distribution

$$\mathcal{F}_{\text{b}}\delta\hat{t} = S_n\hat{c}_n\delta\hat{t}\alpha\left[\omega_{\text{b}}^{(2)} - \omega_{\text{MB}} + \hat{\boldsymbol{c}}\cdot(\hat{\boldsymbol{u}}_{\text{b}} - \hat{\boldsymbol{u}}_{\text{MB}}) + \left(\hat{c}^2 - \tfrac{3}{2}\right)(\tau_{\text{b}} - \tau_{\text{MB}})\right]\hat{f}^0 \quad (22)$$

at each boundary surface. The density perturbation $\omega_{\text{b}}^{(2)}$ is determined from the mass conservation statement (19). Performing the integration results in $\omega_{\text{b}}^{(2)} - \omega_{\text{MB}} = \frac{1}{2}(\tau_{\text{MB}} - \tau_{\text{b}}) - \sqrt{\pi}\hat{\boldsymbol{u}}_{\text{MB}}\cdot\boldsymbol{n}$.

D. Property evaluation

The perturbed hydrodynamic properties are evaluated according to their usual definitions as moments of the distribution function (see, for example, Ref. [4]). By using the formal representation of the distribution function $\hat{f} = \hat{f}_{\text{MB}} + N_{\text{eff}}\Sigma_p s_p\delta(\hat{\boldsymbol{x}} - \hat{\boldsymbol{x}}_p)\delta(\hat{\boldsymbol{c}} - \hat{\boldsymbol{c}}_p)$ and integrating over a single

computational cell (with volume V_{cell}), the hydrodynamic properties can be reduced to the following expressions:

$$\omega = \omega_{\text{MB}} + \frac{N_{\text{eff}}}{V_{\text{cell}}} \sum_{p \in \text{cell}} s_p, \tag{23}$$

$$\hat{\boldsymbol{u}} = \hat{\boldsymbol{u}}_{\text{MB}} + \frac{N_{\text{eff}}}{V_{\text{cell}}} \sum_{p \in \text{cell}} s_p \hat{\boldsymbol{c}}_p, \tag{24}$$

$$\tau = \tau_{\text{MB}} - (\omega - \omega_{\text{MB}}) + \frac{2}{3}\frac{N_{\text{eff}}}{V_{\text{cell}}} \sum_{p \in \text{cell}} s_p \hat{c}_p^2, \tag{25}$$

$$P = \omega + \tau, \tag{26}$$

$$P_{ij} = \delta_{ij}(\omega_{\text{MB}} + \tau_{\text{MB}}) + 2\frac{N_{\text{eff}}}{V_{\text{cell}}} \sum_{p \in \text{cell}} s_p (\hat{c}_i)_p (\hat{c}_j)_p, \tag{27}$$

$$\hat{\boldsymbol{q}} = -\frac{5}{2}(\hat{\boldsymbol{u}} - \hat{\boldsymbol{u}}_{\text{MB}}) + \frac{N_{\text{eff}}}{V_{\text{cell}}} \sum_{p \in \text{cell}} s_p \hat{\boldsymbol{c}}_p \hat{c}_p^2. \tag{28}$$

Here, $P_{ij}=p_{ij}/p_0-\delta_{ij}$ is the perturbed pressure tensor, $P=p/p_0-1$ is the perturbed pressure, and $\hat{\boldsymbol{q}}=\boldsymbol{q}/(p_0 c_0)$ is the dimensionless heat flux vector.

E. Linearized body force

A standard technique for simulating pressure driven flow [25,26] and thermal creep [26,27] (or a combination of both effects) in the linear regime is to introduce an effective "body force," which results in an additional substep in the splitting algorithm for time integration, namely,

$$\left[\frac{\partial \hat{f}^{\text{d}}}{\partial \hat{t}}\right]_{\text{body}} \delta\hat{t} = \hat{c}_z \left[\kappa + \left(\frac{5}{2} - \hat{c}^2\right)\beta\right] \hat{f}^0 \delta\hat{t}. \tag{29}$$

Here, $\kappa=-\frac{L}{p_0}\frac{dp}{dz}$ and $\beta=\frac{L}{T_0}\frac{dT}{dz}$ denote the dimensionless pressure and temperature gradients, respectively, assumed to be in the z direction.

According to Allshouse and Hadjiconstantinou [28], the body force due to the pressure gradient, $\hat{c}_z \kappa \hat{f}^0 \delta\hat{t}$, can be applied by shifting the Maxwell-Boltzmann velocity by the amount $\delta\hat{u}_{\text{MB},z}=\frac{1}{2}\kappa\delta\hat{t}$ every time step. The thermal creep portion of the body force term is applied [29] by generating particles from the distribution,

$$\mathcal{F}_\beta \delta\hat{t} = V_{\text{cell}} \hat{c}_z \left(\tfrac{5}{2} - \hat{c}^2\right) \beta \hat{f}^0 \delta\hat{t}. \tag{30}$$

This distribution was efficiently sampled by the generalized ratio-of-uniforms method as outlined in the Appendix.

IV. IMPLEMENTATION

Based on previous discussions [30,31] indicating higher convergence rates for appropriately symmetrized algorithms, a symmetric version of the above-described algorithm was implemented. Here, the advection and body force substeps are split into half substeps on either side of the collision substep; half steps are performed as described above, only with a time step of $\frac{1}{2}\delta\hat{t}$. The algorithm is summarized below.

ALGORITHM: Symmetrized time splitting

1. half advection step for all cells
   a. advect deviational particles $\hat{\boldsymbol{x}} \to \hat{\boldsymbol{x}} + \frac{\delta\hat{t}}{2}\hat{\boldsymbol{c}}$
   b. enforce boundary conditions for all deviational particles which cross a boundary during step 1.a
      i. with probability $1-\alpha$, perform specular reflections $\hat{\boldsymbol{c}} \to \hat{\boldsymbol{c}} - 2\hat{c}_n \boldsymbol{n}$
      ii. otherwise (i.e., with probability α) absorb particle and count total sign $N^+ - N^-$ absorbed
      iii. emit $|N^+ - N^-|$ particles from the fluxal wall distribution $\hat{c}_n \hat{f}^0$ for $\hat{c}_n > 0$ with sign $\text{sgn}(N^+ - N^-)$
      iv. generate particles by sampling $\mathcal{F}_{\text{b}}\frac{\delta\hat{t}}{2}$
2. half linearized body force step for all cells
   a. shift velocity by $\delta\hat{u}_{\text{MB},z}=\frac{1}{2}\kappa\frac{\delta\hat{t}}{2}$
   b. generate particles by sampling $\mathcal{F}_\beta\frac{\delta\hat{t}}{2}$
3. full collision step for all cells
   a. backward deletion step: delete particles with probability $1-(1+\frac{\delta\hat{t}}{2k})^{-1}$
   b. update properties ω, $\hat{\boldsymbol{u}}$, and τ using Eqs. (23)–(25)
   c. change Maxwell-Boltzmann distribution using Eqs. (9)–(11)
   d. forward deletion step: delete particles with probability $\frac{\delta\hat{t}}{2k}$
4. half linearized body force step for all cells (repeat step 2)
5. half advection step (repeat step 1)
6. sample properties ω, $\hat{\boldsymbol{u}}$, τ, P_{ij}, P, and $\hat{\boldsymbol{q}}$ using Eqs. (23)–(28)
7. repeat steps 1–6 until final time is reached

For comparison purposes (see Sec. V), a nonsymmetrized algorithm was also implemented. This algorithm performs only steps 1, 2, 3.b–3.d, 6, and 7, where steps 1.a, 1.b.iv, 2, and 3.d are of duration $\delta\hat{t}$ rather than $\frac{\delta\hat{t}}{2}$.

V. RESULTS

We validate the method presented here by considering three archetypal problems—namely, pressure driven flow, thermal creep flow, and heat flux between plates held at different temperatures—within a one-dimensional parallel plate geometry with a wall separation of L. The heat flux calculation is also used to demonstrate second-order time convergence for the symmetrized algorithm, as well as to quantify the level of relative statistical uncertainty. All simulations, except where indicated, were performed using the symmetrized version of the algorithm.

A. Pressure driven flow and thermal creep

Using the linearized body force formulation of Sec. III E, we calculated the flow rate induced by a pressure gradient ($\kappa=\epsilon$, $\beta=0$) and temperature gradient ($\kappa=0$, $\beta=\epsilon$) in the direction parallel to the confining plates where $0<\epsilon\ll 1$. Both diffusely reflecting ($\alpha=1$) and $\alpha=0.8$ cases were considered. Our results are compared to those of Loyalka [26]

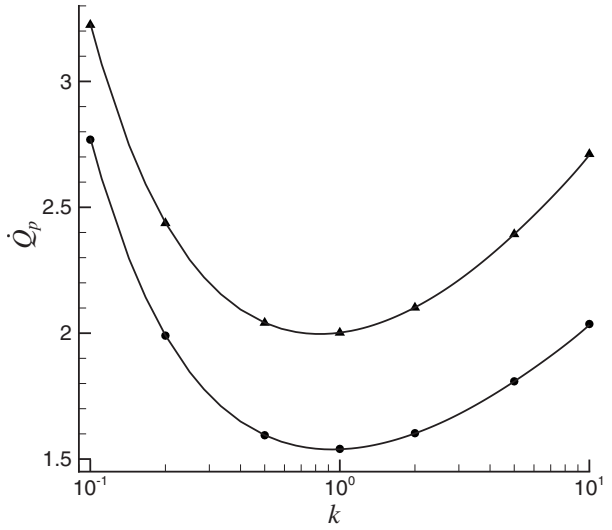

FIG. 1. Dimensionless pressure driven flow rate $\dot{Q}_p=2\bar{u}_z/(\kappa c_0)$ between parallel plates vs Knudsen number for diffusely reflecting boundaries (●) and $\alpha=0.8$ (▲) compared to the results of Loyalka [26] (solid lines).

(obtained by numerical solution of the linearized Boltzmann equation) and are shown in Figs. 1 and 2. In these figures, $\dot{Q}_p=2\bar{u}_z/(\kappa c_0)$ denotes the dimensionless flow rate for pressure driven flow, while $\dot{Q}_T=2\bar{u}_z/(\beta c_0)$ denotes the dimensionless flow rate due to thermal creep; the overbar denotes a spatial average across the system width (L). Excellent agreement is obtained in all cases.

B. Heat flux between parallel plates at different temperatures

Figure 3 shows a comparison between our simulation results for the heat flux between diffusely reflecting parallel

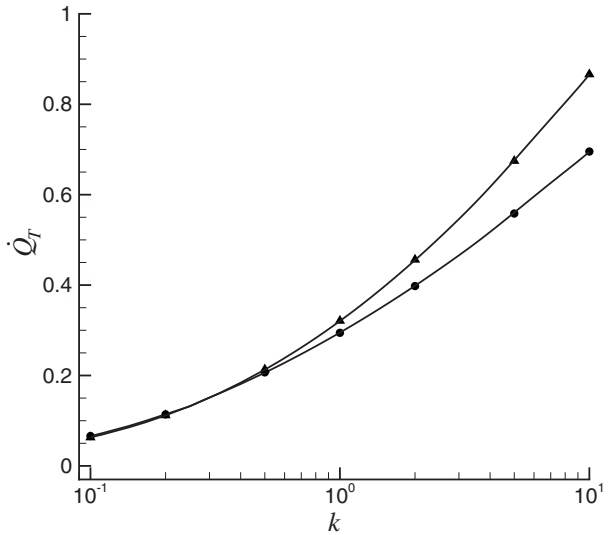

FIG. 2. Dimensionless thermal creep flow rate $\dot{Q}_T=2\bar{u}_z/(\beta c_0)$ between parallel plates vs Knudsen number for diffusely reflecting boundaries (●) and $\alpha=0.8$ (▲) compared to the results of Loyalka [26] (solid lines).

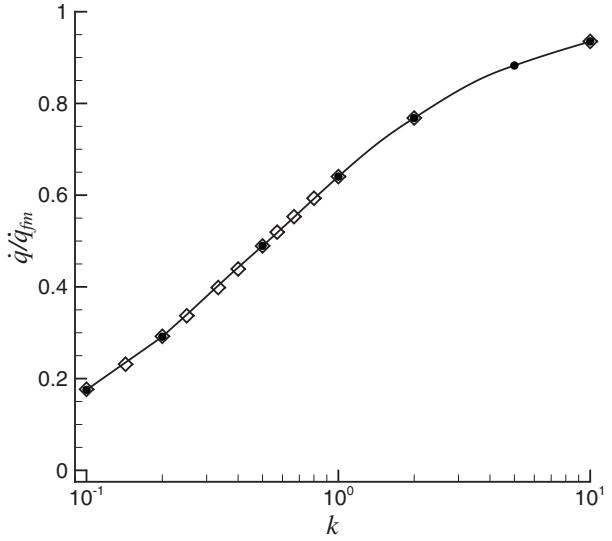

FIG. 3. Heat flux between fully accommodating parallel plates held at a temperature difference ΔT vs Knudsen number (●). The heat flux is normalized by the free molecular flow value $\dot{q}_{fm}=\frac{c_0 p_0}{\sqrt{\pi}}\frac{\Delta T}{T_0}$ and compared to the results from Bassanini *et al.* [32] (◇). The solid line is presented only to guide the eyes.

plates held at a temperature difference ΔT and the results of Bassanini *et al.* [32]; the heat flux is normalized by the free molecular flow ($k\to\infty$) value $\dot{q}_{fm}=\frac{c_0 p_0}{\sqrt{\pi}}\frac{\Delta T}{T_0}$. Excellent agreement between the two sets of results is observed.

We also validate our algorithm by simulating the evolution of the temperature field between parallel plates resulting from an impulsive increase in the wall temperature to $T_b(\hat{x}=0,\hat{t}>0)=T_0+\Delta T$ with $\Delta T/T_0=0.04$, which is near the limit where linearization is appropriate. We compare the results from our algorithm with those from a DSMC simulation (which uses the relaxation-time approximation) as shown in Fig. 4 for diffusely reflecting walls and Fig. 5 for $\alpha=0.7$. Both methods used $\mathcal{N}_{\text{cell}}=50$ cells, approximately 10^5 particles [33] and a time step of $\mathcal{N}_{\text{cell}}\delta\hat{t}=0.1$. The results were averaged over $M=400$ independent samples for the variance-reduced method and over $M=8000$ for DSMC. Although the variance-reduced simulation used fewer samples, it produced noise-free results in contrast to the relatively noisy DSMC results. Despite the differences in uncertainty level, the results show excellent agreement.

C. Validation of second-order convergence in time

In order to demonstrate second-order convergence of the symmetrized version of the algorithm, we simulated the steady-state heat flux between parallel plates for $k=0.1$ using both the symmetrized and nonsymmetrized versions of the algorithm. These computations were performed using $\mathcal{N}_{\text{cell}}=200$ computational cells and approximately 5×10^5 particles in the simulation in order to achieve high spatial resolution and low noise; the solutions were further averaged over $M=1000$ independent samples once steady state was reached. The accuracy was evaluated by comparison with a reference solution $\dot{q}_{\text{ref}}$ obtained using the symmetrized algo-

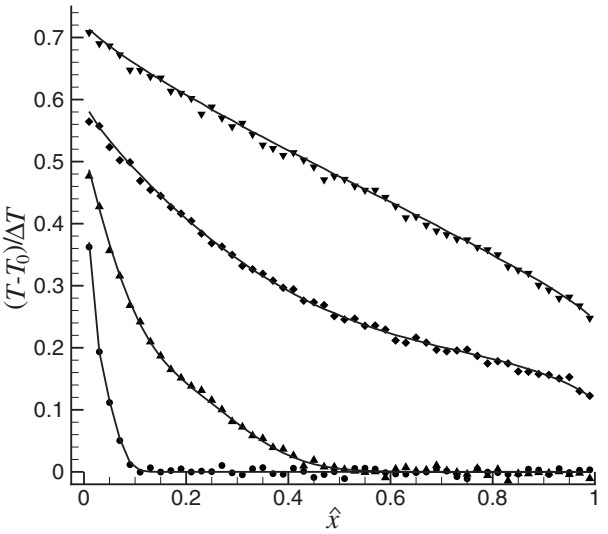

FIG. 4. Transient temperature field resulting from an impulsive change in boundary temperature of magnitude $\Delta T/T_0=0.04$ for $k=1$ and $\alpha=1$. The variance-reduced results (solid lines) are compared to DSMC results (symbols) for $\hat{t}=0.04$ (●), $\hat{t}=0.2$ (▲), $\hat{t}=1$ (♦), and $\hat{t}=4$ (▼).

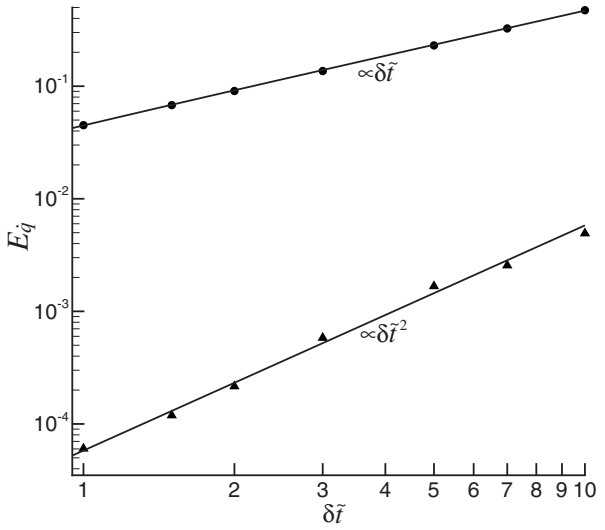

FIG. 6. Numerical error in the heat flux $E_{\dot{q}}=|\dot{q}_{\mathrm{ref}}-\dot{q}|/\dot{q}_{\mathrm{ref}}$ between fully accommodating parallel plates at $k=0.1$ vs time step $\delta\tilde{t}=\mathcal{N}_{\mathrm{cell}}\delta\hat{t}$, comparing the nonsymmetrized (●) and symmetrized (▲) algorithms. The solid lines are proportional to $\delta\tilde{t}$ and $\delta\tilde{t}^2$, respectively.

rithm with time step $\mathcal{N}_{\mathrm{cell}}\delta\hat{t}=0.01$. The difference between $\dot{q}_{\mathrm{ref}}$ and the true solution is due to finite cell size and finite number of samples used [31]. Shown in Fig. 6 is the fractional difference $E_{\dot{q}}=|\dot{q}-\dot{q}_{\mathrm{ref}}|/\dot{q}_{\mathrm{ref}}$ in the heat flux for a range of time steps $\delta\tilde{t}=\mathcal{N}_{\mathrm{cell}}\delta\hat{t}$. Clearly, the symmetrized algorithm achieves second-order time convergence with substantially reduced error compared to the nonsymmetrized algorithm.

D. Variance reduction

In this section we discuss some measures of the variance reduction achieved by the proposed method. Using the steady heat transfer problem described above as a prototype problem, we measured the relative statistical uncertainty (inverse of the signal-to-noise ratio) in the temperature, quantified here by $\sigma_T/\Delta T$, where σ_T is the temperature standard deviation. Figure 7 shows the results obtained for $k=0.1,1,10$ for a computational cell in the center of the simulation domain with $N_{\mathrm{cell}}\approx 950$ particles. This is compared to a nonvariance-reduced particle method, here represented by DSMC; the relative statistical uncertainty for DSMC was estimated using the theory developed in Ref. [11]. We observe that, as expected, the variance-reduced method exhibits

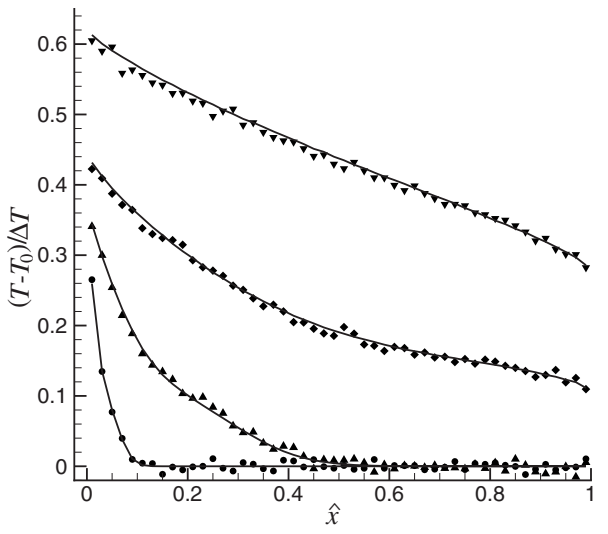

FIG. 5. Transient temperature field resulting from an impulsive change in boundary temperature of magnitude $\Delta T/T_0=0.04$ for $k=1$ and $\alpha=0.7$. The variance-reduced results (solid lines) are compared to DSMC results (symbols) for $\hat{t}=0.04$ (●), $\hat{t}=0.2$ (▲), $\hat{t}=1$ (♦), and $\hat{t}=4$ (▼).

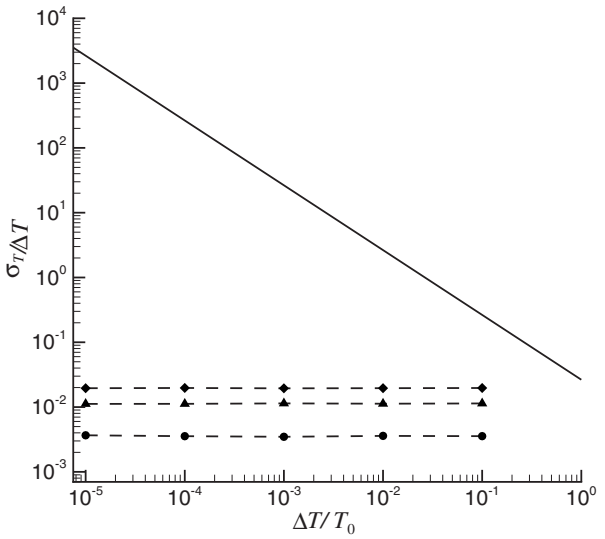

FIG. 7. Normalized statistical error in temperature versus temperature difference for heat transfer between parallel plates at $k=0.1$ (●), $k=1$ (▲), and $k=10$ (♦); these results are compared to a standard particle simulation method represented by the DSMC (solid line). The results for DSMC were estimated using the theory developed in Ref. [11].

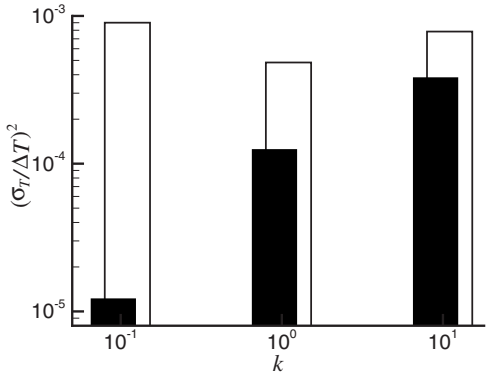

FIG. 8. Normalized variance in temperature vs Knudsen number for heat transfer between parallel plates. Variance observed using the present method (black) is significantly reduced compared to the method of Hadjiconstantinou *et al.* [19] (white), especially for $k \ll 1$.

a constant signal-to-noise ratio in the limit of small deviation from equilibrium, as opposed to DSMC which exhibits a signal-to-noise ratio that is significantly smaller and which continues to decrease as ΔT decreases. Although the exact speedup will be implementation dependent, the deviational algorithm described here is sufficiently similar to DSMC that Fig. 7 may be used to provide an order-of-magnitude estimate of the speedup achieved (which scales as the square of $\sigma_T/\Delta T$). We also observe that the variance reduction achieved has a Knudsen number dependence, with the variance reduction increasing with decreasing Knudsen number; this result can be explained by recalling that local equilibrium conditions set in as $k \to 0$.

Using the same prototype problem, we compare the variance reduction achieved by the present approach with that achieved by the approach of Hadjiconstantinou *et al.* [19] (which simulates the deviation from a spatially homogeneous equilibrium distribution). A comparison between the two methods as a function of k is shown in Fig. 8; this comparison was performed at one value of ΔT since as shown above, in deviational methods, $\sigma_T/\Delta T$ is independent of $\Delta T/T_0$ for small $\Delta T/T_0$. Here, the wall temperatures were taken to be T_0 and $T_0+\Delta T$, while the temperature parameter of the spatially homogeneous equilibrium distribution in the method of Hadjiconstantinou *et al.* [19] was taken to be T_0 to represent the general case when the local value of the hydrodynamic properties is different from the equilibrium value. Figure 8 shows that, for the problem studied here, the present method provides a measurable improvement in variance reduction in the whole range of Knudsen numbers of interest ($0.1 < k < 10$); the difference becomes large for $k < 1$ and $k \ll 1$ in particular, where as a result of the local equilibrium conditions the difference in variance reduction observed is more than a factor of 50.

VI. CONCLUSION

We have presented an efficient variance-reduced particle simulation method for simulating the linearized Boltzmann equation in the relaxation-time approximation. The proposed method exploits the physical interpretation of the relaxation-time approximation to formulate a particularly simple collision algorithm. Numerical results show that, in addition to the constant signal-to-noise ratio expected of variance reduction methods, compared to algorithms simulating the deviation from a spatially homogeneous distribution, the proposed algorithm leads to a reduced number of particles for the same level of variance reduction. This difference is most significant in the limit $k \ll 1$, where local equilibrium conditions allow the present algorithm to be significantly more efficient. These features enable the simulation of near-equilibrium phenomena with significantly reduced statistical uncertainty compared to nondeviational methods such as DSMC.

ACKNOWLEDGMENTS

The authors gratefully acknowledge support from the Singapore-MIT Alliance.

APPENDIX: THE RATIO-OF-UNIFORMS SAMPLING METHOD

For computational efficiency, the samples from distributions (15), (22), and (30) were generated using the generalized multivariate ratio-of-uniforms method [23]. The ratio-of-uniforms method for sampling a function $\mathcal{F}(\hat{\boldsymbol{c}})$ is based on a variable transformation: $\mathcal{F}=H^{3r+1}$, $\hat{\boldsymbol{c}}=\boldsymbol{\xi}/H^r$. Thus, the function can be rewritten as

$$\mathcal{F}d^3\hat{\boldsymbol{c}} = \int_0^{\mathcal{F}} d\mathcal{F}d^3\hat{\boldsymbol{c}} = (3r+1)\int_0^{\mathcal{F}^{1/(3r+1)}} dHd^3\boldsymbol{\xi}. \quad \text{(A1)}$$

The distribution is sampled by choosing values uniformly distributed in $(H,\boldsymbol{\xi})$ which are accepted when $H \le \mathcal{F}^{1/(3r+1)}$, while the parameter r can be chosen to maximize the acceptance probability. According to Wakefield *et al.* [23], $r=\frac{1}{2}$ results in an optimal method for sampling a Gaussian distribution, and this is the value adopted.

The key difficulty for application is finding reasonable bounds for H and $\boldsymbol{\xi}$. The approach adopted here is to break the desired distribution into a linear combination of functions for which tight analytic bounds are known; from these, the bounds for the overall distribution are derived. Specifically, we are interested in the following distributions:

$$\mathcal{F}_x = \hat{c}_x\hat{f}^0, \quad \text{(A2)}$$

$$\mathcal{F}_{xx} = \hat{c}_x^2\hat{f}^0, \quad \text{(A3)}$$

$$\mathcal{F}_{xy} = \hat{c}_x\hat{c}_y\hat{f}^0, \quad \text{(A4)}$$

$$\mathcal{F}_{xcc} = \hat{c}_x\hat{c}^2\hat{f}^0. \quad \text{(A5)}$$

Note that $\mathcal{F}_x$ and $\mathcal{F}_{xy}$ take on both positive and negative values: for sampling such distributions, we sample $|\mathcal{F}|$ and determine the particle sign according to $s=\mathrm{sgn}(\mathcal{F})$. The desired bounds were evaluated (assuming $r=\frac{1}{2}$) according to the formulas in Ref. [23],

TABLE I. Ratio-of-uniforms bounds a and b_i for the functions $\mathcal{F}_Z$ defined in Eqs. (A2)–(A5).

	$\mathcal{F}_x$	$\mathcal{F}_{xx}$	$\mathcal{F}_{xy}$	$\mathcal{F}_{xcc}$
$\pi^{3/2}a^{5/2}$	$1/\sqrt{2e}$	$1/e$	$1/(2e)$	$[3/(2e)]^{3/2}$
$\pi^{3/2}b_x^5$	$(3/e)^3$	$[7/(2e)]^{7/2}$	$27e^{-7/2}/\sqrt{2}$	$(4/e)^4$
$\pi^{3/2}b_y^5$	$5^{5/2}/(2e)^3$	$(5/2)^{5/2}e^{-7/2}$	$27e^{-7/2}/\sqrt{2}$	$20^{5/2}/(27e^4)$
$\pi^{3/2}b_z^5$	$5^{5/2}/(2e)^3$	$(5/2)^{5/2}e^{-7/2}$	$5^{5/2}/(2e)^{7/2}$	$20^{5/2}/(27e^4)$

$$|H_Z| \le a_Z = \sup_{\hat{c}\in\Re^3} |\mathcal{F}_Z|^{2/5}, \tag{A6}$$

$$|\xi_i| \le b_{i,Z} = \sup_{\hat{c}\in\Re^3} |\hat{c}_i||\mathcal{F}_Z|^{1/5}, \tag{A7}$$

where $Z \in \{x,y,z,xy,yz,zx,xx,yy,zz,xcc,ycc,zcc\}$.

The bounds for Eqs. (A2)–(A5) were analytically evaluated and the results are presented in Table I. When the real distribution function is a linear combination of the above distributions, i.e., $\mathcal{F}=\Sigma_Z A_Z \mathcal{F}_Z$, it is straightforward to show that the following overall bounds hold:

$$|H| \le a^* = \left(\sum_Z |A_Z| a_Z^{5/2}\right)^{2/5}, \tag{A8}$$

$$|\xi_i| \le b_i^* = \left(\sum_Z |A_Z| b_{i,Z}^5\right)^{1/5}. \tag{A9}$$

Here, a^* and b_i^* are the bounds actually used in computations.

Note that it is essential to include the Jacobian of transformation $(3r+1)$ in the determination of the number of trial samples to generate when integrating a distribution. For example, when sampling $\mathcal{F}$ on $\hat{\boldsymbol{c}} \in \Re^3$, it is necessary to generate $(3r+1)a^*(2b^*)^3/N_{\text{eff}}=20a^*(b^*)^3/N_{\text{eff}}$ trial samples.

[1] G. Chen, *Nanoscale Energy Transfer and Conversion* (Oxford University Press, New York, 2005).
[2] C. Cercignani, *The Boltzmann Equation and its Applications* (Springer, New York, 1988).
[3] C. Cercignani, *Slow Rarefied Flows: Theory and Application to Micro-Electro-Mechanical Systems* (Birkhauser-Verlag, Basel, 2006).
[4] Y. Sone, *Molecular Gas Dynamics: Theory, Techniques, and Applications* (Birkhauser, Boston, 2007).
[5] W. G. Vincenti and C. H. Kruger, *Introduction to Physical Gas Dynamics* (Wiley, New York, 1965).
[6] W. A. Harrison, *Solid State Theory* (Dover, New York, 1980).
[7] M. Lundstrom, *Fundamentals of Carrier Transport*, 2nd ed. (Cambridge University Press, Cambridge, 2000).
[8] G. Chen, ASME J. Heat Transfer **124**, 320 (2002).
[9] A. A. Joshi and A. Majumdar, J. Appl. Phys. **74**, 31 (1993).
[10] G. A. Bird, *Molecular Gas Dynamics and the Direct Simulation of Gas Flows* (Clarendon, Oxford, 1994).
[11] N. G. Hadjiconstantinou, A. L. Garcia, M. Z. Bazant, and G. He, J. Comput. Phys. **187**, 274 (2003).
[12] L. L. Baker and N. G. Hadjiconstantinou, Phys. Fluids **17**, 051703 (2005).
[13] J. Chun and D. L. Koch, Phys. Fluids **17**, 107107 (2005).
[14] L. L. Baker and N. G. Hadjiconstantinou, in *Proceedings of the 4th International Conference Nanochannels, Microchannels and Minichannels* (Limerick, Ireland, 2006), pp. 377–383.
[15] L. L. Baker and N. G. Hadjiconstantinou, J. Comput. Theor. Nanosci. **5**, 165 (2008).
[16] T. M. M. Homolle and N. G. Hadjiconstantinou, Phys. Fluids **19**, 041701 (2007).
[17] T. M. M. Homolle and N. G. Hadjiconstantinou, J. Comput. Phys. **226**, 2341 (2007).
[18] W. Wagner, Monte Carlo Meth. Appl. **14**, 191 (2008).
[19] N. G. Hadjiconstantinou, G. A. Radtke, and L. L. Baker, e-print arXiv:0905.2218v1 (2009).
[20] L. Pareschi, in *Proceedings of the ESAIM*, edited by T. Goudon, E. Sonnendrucker, and D. Talay (EDP Sciences, Les Ulis Cedex, 2005), Vol. 15, pp. 87–120.
[21] G. Dimarco and L. Pareschi, Multiscale Model. Simul. **6**, 1169 (2008).
[22] This extension of the present method to solutions of the nonlinear Boltzmann equation (which is a significantly different approach to Ref. [19]) is the subject of current work.
[23] J. C. Wakefield, A. E. Gelfand, and A. F. M. Smith, Stat. Comput. **1**, 129 (1991).
[24] In the linearized regime, all flows with $\tau=0$ for the initial and boundary conditions are isothermal.
[25] C. Cercignani and A. Daneri, J. Appl. Phys. **34**, 3509 (1963).
[26] S. K. Loyalka, J. Chem. Phys. **63**, 4054 (1975).
[27] S. K. Loyalka and J. W. J. Cipolla, Phys. Fluids **14**, 1656 (1971).
[28] M. R. Allshouse and N. G. Hadjiconstantinou, AIP Conf. Proc. **1084**, 1015 (2009).
[29] While it is tempting to generalize the technique in Ref. [28] by shifting the velocity by $\delta\hat{u}_{\text{MB},z}=\frac{1}{2}(\kappa+\frac{5}{2}\beta)\delta\hat{t}$ and generating particles from the comparatively simpler distribution $-V_{\text{cell}}\hat{c}_z\hat{c}^2\beta\hat{f}^0\delta\hat{t}$, this approach results in a less efficient algorithm.
[30] N. G. Hadjiconstantinou, Phys. Fluids **18**, 111301 (2006).
[31] D. J. Rader, M. A. Gallis, J. R. Torczynski, and W. Wagner, Phys. Fluids **18**, 077102 (2006).
[32] P. Bassanini, C. Cercignani, and C. D. Pagani, Int. J. Heat Mass Transfer **10**, 447 (1967).
[33] For this problem, the number of deviational particles in the variance-reduced method increases from zero and reaches a nearly constant value at steady state. In our variance-reduced simulation, there were approximately 10^5 particles at steady state, while the DSMC maintained 10^5 particles for the entire simulation time.