RESEARCH ARTICLE | OCTOBER 09 2015

Frequency-domain deviational Monte Carlo method for linear oscillatory gas flows

Daniel R. Ladiges; John E. Sader

Physics of Fluids 27, 102002 (2015)
https://doi.org/10.1063/1.4932108

PHYSICS OF FLUIDS **27**, 102002 (2015)

Frequency-domain deviational Monte Carlo method for linear oscillatory gas flows

Daniel R. Ladiges and John E. Sader[a)]
School of Mathematics and Statistics, The University of Melbourne, Victoria 3010, Australia

(Received 5 May 2015; accepted 16 September 2015; published online 9 October 2015)

Oscillatory non-continuum low Mach number gas flows are often generated by nanomechanical devices in ambient conditions. These flows can be simulated using a range of particle based Monte Carlo techniques, which in their original form operate exclusively in the time-domain. Recently, a frequency-domain *weight-based* Monte Carlo method was proposed [D. R. Ladiges and J. E. Sader, "Frequency-domain Monte Carlo method for linear oscillatory gas flows," J. Comput. Phys. **284**, 351–366 (2015)] that exhibits superior statistical convergence when simulating oscillatory flows. This previous method used the Bhatnagar-Gross-Krook (BGK) kinetic model and contains a "virtual-time" variable to maintain the inherent time-marching nature of existing Monte Carlo algorithms. Here, we propose an alternative frequency-domain *deviational* Monte Carlo method that facilitates the use of a wider range of molecular models and more efficient collision/relaxation operators. We demonstrate this method with oscillatory Couette flow and the flow generated by an oscillating sphere, utilizing both the BGK kinetic model and hard sphere particles. We also discuss how oscillatory motion of arbitrary time-dependence can be simulated using computationally efficient parallelization. As in the weight-based method, this deviational frequency-domain Monte Carlo method is shown to offer improved computational speed compared to the equivalent time-domain technique. © *2015 AIP Publishing LLC.* [http://dx.doi.org/10.1063/1.4932108]

I. INTRODUCTION

Recent advances in nanoscale fabrication have driven strong activity in exploring the use of nanomechanical devices in a range of applications, including atomic-scale mass measurement,[1–3] imaging with atomic resolution,[4,5] micro- and nanoscale measurement of fluid properties,[6–8] and ultrasensitive environmental sensing.[9,10] Ability to model gas flows generated by small mechanical devices in ambient conditions (in air) is particularly important given the ubiquity of this environment.

Three key features control the nature of the gas flows generated: (i) the device dimensions are often comparable to the gas mean free path, (ii) these devices are commonly driven at resonance and thus generate an oscillatory flow whose frequency can be similar to the gas molecular collision frequency, and (iii) the gas flow speeds are in the low Mach number regime.

The first of these characteristics is quantified by the Knudsen number, $\mathrm{Kn} \equiv \lambda/L$, where λ is the gas mean free path and L is a dominant length scale of the flow. The second by the scaled oscillation frequency, $\theta \equiv \omega/\nu$, where ω is the device frequency and ν is the molecular collision frequency. Together, these parameters indicate the degree of rarefaction (non-continuum nature) of the flow. In the continuum limit, $\mathrm{Kn} \to 0$ and $\theta \to 0$, gas flows can be accurately modeled by the Navier-Stokes equations which can be solved using standard numerical or analytical approaches in many cases of practical interest. However, such solutions do not apply if the continuum hypothesis is violated, i.e., either Kn or θ is not infinitesimally small. While approaches exist to extend the validity of the

a)Author to whom correspondence should be addressed. Electronic mail: jsader@unimelb.edu.au

Navier-Stokes equations to weakly non-continuum flows through use of slip models,[11–13] solutions for arbitrary Kn and θ require alternative methods.

The Boltzmann equation describes the time evolution of the distribution of gas particles in both the particle velocity and spatial domains; such high dimensionality complicates its solution using standard numerical techniques. This issue is circumvented through the use of particle based approaches, foremost of which is the direct simulation Monte Carlo (DSMC) method, originally proposed by Bird.[14] DSMC employs a statistical approach in which simulator particles are used to represent a large number of real particles in a gas flow. Time evolution of these simulator particles is specified by the Boltzmann equation, with collisions being determined probabilistically. Despite its broad applicability, use of standard DSMC is problematic in modeling gas flows generated by nanomechanical devices due to their operation at low Mach number.

Low Mach number gas flows exhibit very small deviations from equilibrium. Use of a stochastic method such as DSMC will tend to obscure these deviations with statistical noise associated with this equilibrium. Recently, a range of alternative particle-based Monte Carlo techniques has been proposed to address this issue. Utilizing the method of control variates,[15] rather than solving for the complete distribution function of the gas, these techniques simulate only the deviation from a specified equilibrium. They may be broadly grouped into two classes:[16] (i) *deviational* and (ii) *weight-based* methods. In the first class, velocities of the simulator particles directly represent the deviation of the distribution function from equilibrium. In contrast, weight-based methods account for deviations from equilibrium through the allocation of weights as a function of molecular velocity. In the case of weight-based methods, the chosen equilibrium is a spatially constant *global* equilibrium. However, deviational methods may use either a global equilibrium or a spatially varying *local* equilibrium. Since the local equilibrium is specified using the local density, bulk velocity, and temperature at a given position, this tends to be "closer" to the actual distribution function of the flow at that position — often resulting in fewer simulator particles being required to represent the deviation.

The deviational approach has been implemented using the Bhatnagar-Gross-Krook (BGK) relaxation model,[17–19] hard sphere (HS) interactions,[20,21] and variable hard sphere (VHS) collisions.[16,18] The weight-based approach has also been implemented using the BGK relaxation model;[22,23] however, attempts at more realistic collision models[24–26] lead to difficulties with algorithmic stability. This effect is alleviated by periodically "reconstructing" the distribution function, at the price of introducing implicit velocity discretization, and increasing the computational cost. Other differences exist between these two classes in terms of computational efficiency and ease of implementation. For example, deviational methods concentrate simulator particles where the deviation from equilibrium is large, whilst weight-based methods use a spatially uniform distribution of particles. This lends an advantage to the deviational approach in unbounded problems, which often contains large regions of negligible flow. On the other hand, weight-based methods tend to be algorithmically simpler and easier to apply.

Each of these Monte Carlo methods can efficiently simulate the nanoscale gas flows described above, at low Mach number and arbitrary Knudsen number. Importantly, gas flows generated by many nanoscale devices are time-dependent and oscillatory in nature. While the above-mentioned Monte Carlo methods can also effectively simulate oscillatory flows at arbitrary frequency, they are inherently time-domain techniques. A solution in the frequency-domain is therefore useful and desirable because this would provide direct access to the amplitude and phase of the flow field (information that is often required) and potentially lead to computational savings.

Recently, a frequency-domain lattice Boltzmann (LB) method was proposed.[27] Based on the linearized Boltzmann-BGK equation,[28–31] this approach makes use of a "virtual-time" variable to maintain the inherent time-marching nature of the LB algorithm. In this method, steady solutions in virtual-time correspond to actual solutions in the frequency-domain. These, in turn, are directly related to oscillatory solutions in the time-domain. This approach (like all LB methods) is formulated primarily for continuum flows,[32] and its applicability outside of this regime is limited to weakly non-continuum flows, i.e., small Kn and θ.[33] In Ref. 34, we presented an approach that combines this virtual-time concept with a low Mach number Monte Carlo method, enabling direct simulation in the frequency-domain for arbitrary Kn and θ. In addition to providing direct access

to amplitude and phase information, this frequency-domain approach was shown to exhibit superior statistical convergence with a corresponding improvement in computational speed, compared to the equivalent time-domain Monte Carlo simulation. This previous frequency-domain method[34] applies the *weight-based* approach described above, using the BGK relaxation operator.[23] While extension to more realistic collision models is possible in principle, difficulties identical to those of the time-domain weight-based methods are encountered.[24,26]

In this article, we develop the frequency-domain Monte Carlo concept using a *deviational* low Mach number approach. This allows the method to be efficiently applied using a wider range of relaxation/collision models — we demonstrate the BGK and HS models here. In Section II, we provide an overview of the linearized Boltzmann equation. The algorithm for the frequency-domain Monte Carlo method using a deviational approach is then presented. This method is validated in Section III for two canonical flow problems: (i) oscillatory plane Couette flow and (ii) the flow generated by an oscillating sphere. Comparisons are made with existing accurate numerical BGK solutions,[35] deviational time-domain Monte Carlo solutions, and previous results using the weighted frequency-domain method.[34] In Section IV, a quantitative assessment of the convergence properties and computational cost of the deviational frequency-domain method, relative to the corresponding time-domain method, are provided. Finally, an approach to efficiently parallelize the algorithm for oscillatory flows of general time-dependence is given in Section V. Concluding comments are given in Section VI with additional comments and calculations relegated to the Appendices.

II. DEVIATIONAL FREQUENCY-DOMAIN MONTE CARLO METHOD

The deviational frequency-domain Monte Carlo method is now described. We begin by summarizing the linearized Boltzmann equation with both the BGK relaxation model and HS collision operator, and its form in the frequency-domain. The virtual-time concept[27] is then applied to this frequency-domain equation, which is solved using a modification to the deviational time-domain algorithms of Homolle, Radtke, and Hadjiconstantinou that have been implemented for both of these molecular models.[17,20] Throughout, the subscripts b, c, and f are used to indicate a boundary, cell, and cell interface, respectively.

A. Frequency-domain Boltzmann equation

The Boltzmann equation is

$$\frac{\partial f}{\partial t} + \mathbf{c} \cdot \frac{\partial f}{\partial \mathbf{x}} = Q[f], \tag{1}$$

where f is the number density distribution function of gas particles, t is the time, $\mathbf{x}$ is the position, and $\mathbf{c}$ is the particle velocity. Throughout, the term "distribution function" shall refer to a number density distribution function. The operator Q specifies the collision/relaxation process. As discussed above, we focus on both HS collisions and the BGK relaxation process. The HS collision operator is given by

$$Q_{\mathrm{HS}}[f(\mathbf{c})] = \int_{-\infty}^{\infty} \int_{0}^{4\pi} \left[f(\mathbf{c}') f(\mathbf{c}_1') - f(\mathbf{c}) f(\mathbf{c}_1) \right] |\mathbf{c} - \mathbf{c}_1| \eta \, d\Omega \, d\mathbf{c}_1, \tag{2}$$

where $\mathbf{c}'$ and $\mathbf{c}_1'$ indicate post collision velocities, η is the collisional cross section of the particles, and Ω is the solid angle. The BGK relaxation operator is

$$Q_{\mathrm{BGK}}[f] = \nu (f^{\mathrm{eq}} - f), \tag{3}$$

where ν is the molecular collision frequency, and f^{eq} is the local equilibrium distribution function, given by the Maxwell-Boltzmann distribution

$$f^{\mathrm{eq}} = \frac{n(\mathbf{x},t)}{[2\pi R T(\mathbf{x},t)]^{3/2}} \exp\left(-\frac{|\mathbf{c} - \mathbf{u}(\mathbf{x},t)|^2}{2 R T(\mathbf{x},t)} \right), \tag{4}$$

where R is the specific gas constant.

The bulk flow properties of the gas are obtained from the moments of the distribution function, f, with the local number density, flow velocity, and temperature specified by

$$n = \int_{-\infty}^{\infty} f d\mathbf{c}, \qquad \mathbf{u} = \frac{1}{n}\int_{-\infty}^{\infty} \mathbf{c} f d\mathbf{c}, \quad T = \frac{1}{3nR}\int_{-\infty}^{\infty} |\mathbf{c}-\mathbf{u}|^2 f d\mathbf{c}, \tag{5}$$

respectively.

We assume that all gas particles impacting a solid boundary are re-emitted diffusely in equilibrium with the boundary velocity and temperature, i.e.,

$$f_b = \frac{n_b(\mathbf{x},t)}{[2\pi R T_b(\mathbf{x},t)]^{3/2}} \exp\left(-\frac{|\mathbf{c}-\mathbf{u}_b(\mathbf{x},t)|^2}{2RT_b(\mathbf{x},t)}\right), \tag{6}$$

where $\mathbf{u}_b$ and T_b are the velocity and temperature of the boundary; the subscript b refers to the boundary. Here, n_b is the number density of particles emitted and is determined by mass conservation, with

$$\int_{\mathbf{n}\cdot(\mathbf{c}-\mathbf{u}_b)<0} f\mathbf{n}\cdot(\mathbf{c}-\mathbf{u}_b)d\mathbf{c} = -\int_{\mathbf{n}\cdot(\mathbf{c}-\mathbf{u})>0} f_b\mathbf{n}\cdot(\mathbf{c}-\mathbf{u})d\mathbf{c}, \tag{7}$$

where $\mathbf{n}$ is the outward normal.

We are interested in flows that deviate only slightly from a known global equilibrium distribution, given by

$$\bar{f}^{\mathrm{eq}} = \frac{n_0}{[2\pi R T_0]^{3/2}} \exp\left(-\frac{c^2}{2RT_0}\right), \tag{8}$$

where n_0 and T_0 represent the number density and temperature of the gas at rest, and $c = |\mathbf{c}|$. We therefore express the distribution function as

$$f = \bar{f}^{\mathrm{eq}} + g, \tag{9}$$

where g defines the perturbation from the global equilibrium. The local number density and temperature may be expressed as

$$n = n_0(1+\sigma), \quad T = T_0(1+\kappa), \tag{10}$$

with the local number density, velocity, and temperature perturbations given by

$$\begin{aligned} \sigma &= \int_{-\infty}^{\infty} \frac{g}{n_0}\, d\mathbf{c}, \\ \mathbf{u} &= \int_{-\infty}^{\infty} \mathbf{c}\frac{g}{n_0}\, d\mathbf{c}, \\ \kappa &= \frac{2}{3}\int_{-\infty}^{\infty} \left(\frac{c^2}{2RT_0} - \frac{3}{2}\right)\frac{g}{n_0}\, d\mathbf{c}. \end{aligned} \tag{11}$$

We note that g, σ, $\mathbf{u}$, and κ may all be time-dependent. The local equilibrium distribution function, Eq. (4), is expanded to linear order, giving

$$f^{\mathrm{eq}} = \bar{f}^{\mathrm{eq}} + g^{\mathrm{eq}}, \tag{12}$$

where

$$g^{\mathrm{eq}} = \bar{f}^{\mathrm{eq}}\left\{\sigma + \frac{\mathbf{c}\cdot\mathbf{u}}{RT_0} + \left(\frac{c^2}{2RT_0} - \frac{3}{2}\right)\kappa\right\} \tag{13}$$

is the perturbation to the global equilibrium distribution function, $\bar{f}^{\mathrm{eq}}$. We designate g^{eq} the "equilibrium perturbation."

Substituting Eq. (9) into Eq. (1) yields the required linearized Boltzmann equation

$$\frac{\partial g}{\partial t} + \mathbf{c}\cdot\frac{\partial g}{\partial \mathbf{x}} = L[g], \tag{14}$$

where L is the linearized collision/relaxation operator,[17,20]

$$L_{\mathrm{HS}}[g(\mathbf{c})] = \int_{-\infty}^{\infty} K(\mathbf{c},\mathbf{c}_1) g(\mathbf{c}_1) d\mathbf{c}_1 - \hat{\nu}(\mathbf{c}) g(\mathbf{c}), \tag{15a}$$

$$L_{\mathrm{BGK}}[g] = \nu(g^{\mathrm{eq}} - g). \tag{15b}$$

The functions K and $\hat{\nu}$ are derived in Ref. 36; further details are given in Appendix A. The diffuse boundary condition for g is

$$g_b = \bar{f}^{\mathrm{eq}} \left\{ \sigma_b + \frac{\mathbf{c}\cdot\mathbf{u}_b}{RT_0} + \left(\frac{c^2}{2RT_0} - \frac{3}{2} \right) \kappa_b \right\}, \tag{16}$$

where

$$\sigma_b = \mathbf{u}_b \cdot \mathbf{n} \sqrt{\frac{\pi}{2RT_0}} - \frac{1}{2}\kappa_b - 2\sqrt{\pi} \int_{\mathbf{n}\cdot\mathbf{c}<0} \mathbf{n}\cdot\mathbf{c} \frac{g}{n_0} d\mathbf{c}. \tag{17}$$

Because we are interested in oscillatory (time-dependent) flows, the perturbation functions g and g^{eq} are expressed in terms of the following time dependence:

$$g(t,\mathbf{c},\mathbf{x}|\omega) = \tilde{g}(\mathbf{c},\mathbf{x}|\omega) e^{-\mathrm{i}\omega t}, \quad g^{\mathrm{eq}}(t,\mathbf{c},\mathbf{x}|\omega) = \tilde{g}^{\mathrm{eq}}(\mathbf{c},\mathbf{x}|\omega) e^{-\mathrm{i}\omega t}, \tag{18}$$

where ω is the angular oscillation frequency and i is the imaginary unit. Equation (14) then becomes

$$\mathbf{c}\cdot\frac{\partial \tilde{g}}{\partial \mathbf{x}} = L[\tilde{g}] + \mathrm{i}\omega\tilde{g}, \tag{19}$$

which is the linearized Boltzmann equation in the frequency-domain. The tilde "~" indicates the spatial and particle velocity component of the distribution function, with the time-dependence removed; see Eq. (18). This separation of the time-dependence is applied to all bulk flow variables, leading to the following expressions for the number density, bulk velocity, and temperature perturbations:

$$\begin{aligned} \tilde{\sigma} &= \int_{-\infty}^{\infty} \frac{\tilde{g}}{n_0} d\mathbf{c}, \\ \tilde{\mathbf{u}} &= \int_{-\infty}^{\infty} \mathbf{c}\frac{\tilde{g}}{n_0} d\mathbf{c}, \\ \tilde{\kappa} &= \frac{2}{3}\int_{-\infty}^{\infty} \left(\frac{c^2}{2RT_0} - \frac{3}{2} \right) \frac{\tilde{g}}{n_0} d\mathbf{c}, \end{aligned} \tag{20}$$

and the equilibrium perturbation

$$\tilde{g}^{\mathrm{eq}} = \bar{f}^{\mathrm{eq}} \left\{ \tilde{\sigma} + \frac{\mathbf{c}\cdot\tilde{\mathbf{u}}}{RT_0} + \left(\frac{\mathbf{c}^2}{2RT_0} - \frac{3}{2} \right) \tilde{\kappa} \right\}. \tag{21}$$

The boundary conditions in Eqs. (16) and (17) are specified accordingly.

As in Ref. 34, we introduce the virtual-time variable, τ,[27] and thus replace the left hand side of Eq. (19) with the material derivative of $\tilde{g}$,

$$\frac{D\tilde{g}}{D\tau} = \frac{\partial \tilde{g}}{\partial \tau} + \mathbf{c}\cdot\frac{\partial \tilde{g}}{\partial \mathbf{x}} = L[\tilde{g}] + \mathrm{i}\omega\tilde{g}. \tag{22}$$

Clearly, the steady solution to Eq. (22) corresponds to the solution of Eq. (19). This simple modification allows us to retain the basic structure of a steady time-domain Monte Carlo method, whilst explicitly solving in the frequency-domain. In Sec. II B, we employ the low Mach number time-domain deviational algorithms of Homolle, Radtke, and Hadjiconstantinou to solve Eq. (22).

B. Deviational frequency-domain Monte Carlo method

We begin by splitting Eq. (22) into advection, collision/relaxation, and oscillatory operations,

$$\left.\frac{\partial \tilde{g}}{\partial \tau}\right|_{\mathrm{adv}} = -\mathbf{c}\cdot\frac{\partial \tilde{g}}{\partial \mathbf{x}}, \tag{23a}$$

$$\left.\frac{\partial \tilde{g}}{\partial \tau}\right|_{\text{col}} = L[\tilde{g}], \tag{23b}$$

$$\left.\frac{\partial \tilde{g}}{\partial \tau}\right|_{\text{osc}} = \mathrm{i}\omega\tilde{g}, \tag{23c}$$

respectively. We note two major differences between the frequency-domain Monte Carlo and the corresponding time-domain methods: (i) the perturbation, g, to the distribution function is now complex-valued and (ii) the addition of new oscillatory operation, Eq. (23c).

In a *time-domain deviational* Monte Carlo method, the perturbation g is represented by the distribution of simulator particles in the spatial and particle velocity domains. Because g may be negative, each particle carries a "value" of $\mathcal{W} = \pm 1$ to indicate whether it represents a positive or negative part of the distribution. This value is distinct from the weight used in *weight-based* methods, in that it is restricted to two discrete values only. The number of physical particles represented by a simulator particle is $\mathcal{W}N_p$, where N_p is a constant. When $\mathcal{W}N_p$ is positive, this represents an addition of N_p particles to the equilibrium state at the simulator particle's location and velocity. When it is negative, a subtraction of N_p particles is indicated.

In the *frequency-domain deviational* method, we extend this approach to use complex particle values, $\tilde{\mathcal{W}}$. In this case, the real and imaginary parts of $\tilde{\mathcal{W}}N_p$ indicate how many physical particles are added or subtracted to both the real and imaginary parts of the equilibrium, respectively. Whilst $\mathcal{W}$ is restricted to a discrete set of values in the time-domain method, this restriction applies to the value of $|\tilde{\mathcal{W}}|$ in the frequency-domain method, with $|\tilde{\mathcal{W}}| = 1$ or $\sqrt{2}$, as we discuss in Secs. II B 1 and II B 2.

As for the time-domain deviational approach, the actions of collisions and boundary interactions on $\tilde{g}$ are represented by changes to the distribution of simulator particles in velocity space, while advection is represented by changes to the spatial distribution of the simulator particles. The oscillatory operation in Eq. (23c) is applied using updates to $\tilde{\mathcal{W}}$ for each particle, which are detailed below. Following the standard DSMC algorithm, the spatial domain is divided into discrete cells. Simulations are progressed by solving and applying the operations in Eq. (23) in each cell, over a small virtual-time step $\Delta\tau$.

As mentioned above, simulator particles are used to indicate the deviation from an equilibrium state. This may be either the spatially uniform *global* equilibrium or the spatially varying *local* equilibrium. Following the approach of Refs. 17 and 20, we split the perturbation from the global equilibrium distribution into two parts,

$$\tilde{g} = \tilde{g}^{\text{eq}} + \tilde{g}^{\text{d}}, \tag{24}$$

where $\tilde{g}^{\text{eq}}$ is the spatially varying equilibrium perturbation; see Eq. (21). The distribution $\tilde{g}^{\text{d}}$ now represents the deviation from *local* equilibrium; see Eq. (9). When using simulator particles to represent $\tilde{g}^{\text{d}}$, the number density, bulk velocity, and temperature perturbations are stored for each cell. We may then avoid the use of particles to represent $\tilde{g}^{\text{eq}}$ by calculating it directly using Eq. (21). When $\tilde{g}^{\text{eq}}$ is given as a spatially piecewise-constant function, this approach has been demonstrated[17,21] to reduce the number of particles required in simulations of simple one-dimensional flows. Discussion of the simulation of higher dimensional cases using a continuous function for $\tilde{g}^{\text{eq}}$ is contained in Refs. 37 and 38.

Depending on the flow being simulated, here, we either use the piecewise-constant method to specify the *local* equilibrium or simply employ particles to represent the whole of $\tilde{g}$, i.e., the deviation from the *global* equilibrium. The effect of this choice is discussed in Section III. We now give a separate description of each approach.

1. Using particles to represent the deviation from global equilibrium

a. Oscillatory operation. We first consider Eq. (23c), which we solve explicitly over the interval $\Delta\tau$. This yields the following relationship connecting the solution for $\tilde{g}$ at the current

virtual-time τ to a future virtual-time $\tau + \Delta\tau$:

$$\tilde{g}(\tau + \Delta\tau) = e^{i\omega\Delta\tau}\tilde{g}(\tau). \tag{25}$$

Equation (25) is applied by updating $\tilde{\mathcal{W}}$ for every particle using

$$\begin{bmatrix} \tilde{\mathcal{W}}_{\text{re,new}} \\ \tilde{\mathcal{W}}_{\text{im,new}} \end{bmatrix} = \begin{bmatrix} \cos(\Delta\tau\omega) & -\sin(\Delta\tau\omega) \\ \sin(\Delta\tau\omega) & \cos(\Delta\tau\omega) \end{bmatrix} \begin{bmatrix} \tilde{\mathcal{W}}_{\text{re}} \\ \tilde{\mathcal{W}}_{\text{im}} \end{bmatrix}, \tag{26}$$

where the subscripts "re" and "im" refer to the real and imaginary components, respectively. This operation preserves the value of $|\tilde{\mathcal{W}}|$. We next consider Eq. (23b) and solve for the BGK model and HS particles separately.

b. Collision operation (BGK model). The BGK relaxation operation is solved analytically over the interval $\Delta\tau$, giving

$$\tilde{g}(\tau + \Delta\tau, \mathbf{c}) = e^{-\nu\Delta\tau}\tilde{g}(\tau, \mathbf{c}) + (1 - e^{-\nu\Delta\tau})\tilde{g}^{\text{eq}}(\tau, \mathbf{c}). \tag{27}$$

To implement this solution in each cell, we evaluate the moments of $\tilde{g}(\tau, \mathbf{c})$ using Monte Carlo integration,

$$\begin{aligned} \tilde{\sigma} &= \frac{N_p}{n_0 V_c} \sum_i^{N_c} \tilde{\mathcal{W}}_i, \\ \tilde{\mathbf{u}} &= \frac{N_p}{n_0 V_c} \sum_i^{N_c} \tilde{\mathcal{W}}_i \mathbf{c}_i, \\ \tilde{\kappa} &= \frac{N_p}{n_0 V_c} \frac{2}{3} \sum_i^{N_c} \tilde{\mathcal{W}}_i \left(\frac{c_i^2}{2RT_0} - \frac{3}{2} \right), \end{aligned} \tag{28}$$

where V_c is the volume of the cell, and the sums are taken over the N_c particles in the cell, with the index i indicating the ith particle. Next, $N_c(1 - e^{-\nu\Delta\tau})$ particles are deleted from the cell, retaining the fraction $e^{-\nu\Delta\tau}$ as per Eq. (27). Finally, following the deviational approach, particles are generated from the distribution

$$\tilde{g}^{\text{gen}}(\tau, \mathbf{c}) = V_c(1 - e^{-\nu\Delta\tau})\tilde{g}^{\text{eq}}(\tau, \mathbf{c}), \tag{29}$$

where $\tilde{g}^{\text{eq}}$ is found by inserting the moments calculated using Eq. (28) into Eq. (21). The superscript "gen" is used to indicate distributions that are used to generate particles.

There are a range of methods available[17,39] to sample the distribution $\tilde{g}^{\text{gen}}(\tau, \mathbf{c})$. For simplicity, we describe a single general method for distribution sampling for use throughout this article, the *acceptance-rejection* method,[36] whilst noting that considerable performance gains are available by tailoring the approach to individual distributions.[17,36] We first select N_g candidate particles from the distribution $p(\mathbf{c})$, where

$$p(\mathbf{c}) > \max\{|\tilde{g}^{\text{gen}}_{\text{re}}(\tau, \mathbf{c})|, |\tilde{g}^{\text{gen}}_{\text{im}}(\tau, \mathbf{c})|\} \tag{30}$$

and

$$N_g = \frac{1}{N_p} \int_{-\infty}^{\infty} p(\mathbf{c}) d\mathbf{c}. \tag{31}$$

A particle is accepted if

$$\mathcal{R}p(\mathbf{c}) < \max\{|\tilde{g}^{\text{gen}}_{\text{re}}(\tau, \mathbf{c})|, |\tilde{g}^{\text{gen}}_{\text{im}}(\tau, \mathbf{c})|\}, \tag{32}$$

where $\mathcal{R}$ is a uniformly distributed random variable with $\mathcal{R} \in [0, 1]$; if the condition in Eq. (32) is not met, the particle is discarded. In the case,

$$\mathcal{R}p(\mathbf{c}) < |\tilde{g}^{\text{gen}}_{\text{re}}(\tau, \mathbf{c})|, \tag{33}$$

the real particle value is set to

$$\tilde{\mathcal{W}}_{\text{re}} = \text{sgn}[\tilde{g}^{\text{gen}}_{\text{re}}(\tau, \mathbf{c})], \tag{34}$$

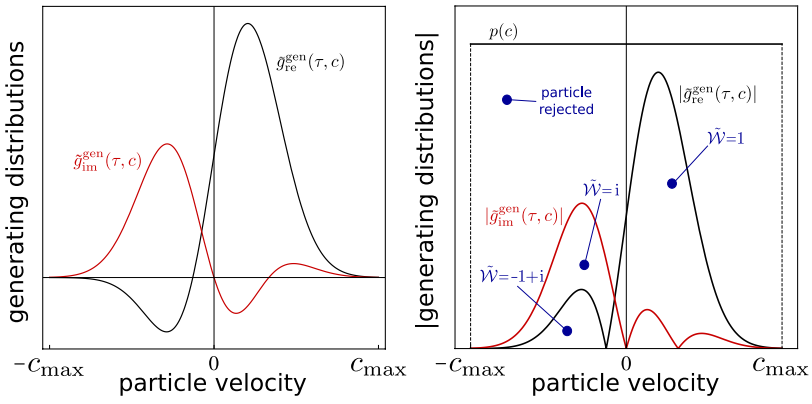

FIG. 1. Examples of particle generation for any generating distribution. Left panel: real and imaginary parts of an example generating distribution, $\tilde{g}^{\rm gen}$. Right panel: illustration showing the particle selection and assignment of $\tilde{\mathcal{W}}$. Magnitudes of the real and imaginary parts of $\tilde{\mathcal{W}}$ (either 0 or 1) are dictated by the distributions $|\tilde{g}^{\rm gen}_{\rm re}|$ and $|\tilde{g}^{\rm gen}_{\rm im}|$ that enclose the chosen particle velocity; their signs are given by those of the original distributions $\tilde{g}^{\rm gen}_{\rm re}$ and $\tilde{g}^{\rm gen}_{\rm im}$. Examples: points (solid dots) are selected randomly beneath the bounding distribution $p(c)$. The upper left point lies above both $|\tilde{g}^{\rm gen}_{\rm re}|$ and $|\tilde{g}^{\rm gen}_{\rm im}|$; it is therefore rejected. In the lower left cases, one point lies below $|\tilde{g}^{\rm gen}_{\rm im}|$ only and is assigned the value $\tilde{\mathcal{W}} = \mathrm{i}$; because $\tilde{g}^{\rm gen}_{\rm im} > 0$. The other point lies below both $|\tilde{g}^{\rm gen}_{\rm re}|$ and $|\tilde{g}^{\rm gen}_{\rm im}|$ and is assigned the value $\tilde{\mathcal{W}} = -1 + \mathrm{i}$; because $\tilde{g}^{\rm gen}_{\rm re} < 0$ and $\tilde{g}^{\rm gen}_{\rm im} > 0$. The right-most point is correspondingly assigned the value $\tilde{\mathcal{W}} = 1$.

otherwise $\tilde{\mathcal{W}}_{\rm re} = 0$. Similarly, where the condition

$$\mathcal{R}p(\mathbf{c}) < |\tilde{g}^{\rm gen}_{\rm im}(\tau,\mathbf{c})| \tag{35}$$

is met, the imaginary particle value is given by

$$\tilde{\mathcal{W}}_{\rm im} = \mathrm{sgn}[\tilde{g}^{\rm gen}_{\rm im}(\tau,\mathbf{c})], \tag{36}$$

with $\tilde{\mathcal{W}}_{\rm im} = 0$ otherwise. If only one of the two conditions in Eqs. (33) and (35) is met, a particle with $|\tilde{\mathcal{W}}| = 1$ results. If both conditions are met, a particle with $|\tilde{\mathcal{W}}| = \sqrt{2}$ is created. This particle generation procedure is illustrated in Fig. 1.

The accepted particles are placed randomly within the cell. The simplest choice of $p(\mathbf{c})$ is a uniform sphere given by $|\mathbf{c}| < c_{\max}$, where $c_{\max}$ is appropriately large. Alternatively, performance may be improved by using importance sampling.[36]

c. Collision operation (HS model). Unlike the BGK model, the HS collision operation is not analytically tractable. We therefore use the approximate finite difference solution,[20,21,36]

$$\tilde{g}(\tau+\Delta\tau,\mathbf{c}) = [1-\Delta\tau\hat{\nu}(\mathbf{c})]\tilde{g}(\tau,\mathbf{c}) + \Delta\tau\int_{-\infty}^{\infty} K(\mathbf{c},\mathbf{c}_1)\tilde{g}(\tau,\mathbf{c}_1)d\mathbf{c}_1, \tag{37}$$

where the collision kernel, K, and the collision frequency term, $\hat{\nu}$, are detailed in Appendix A. To implement Eq. (37), we first mark each particle in a given cell for deletion with probability $\Delta\tau\hat{\nu}(\mathbf{c})$. Note that, in contrast with the BGK deletion process, this probability is velocity dependent. These particles are not immediately removed, because they are now required for the particle generation step. The distribution from which the particles are generated is

$$\tilde{g}^{\rm gen}(\tau,\mathbf{c}) = \Delta\tau\int_{-\infty}^{\infty} K(\mathbf{c},\mathbf{c}_1)\tilde{g}(\tau,\mathbf{c}_1)d\mathbf{c}_1. \tag{38}$$

As with the BGK model, we again apply the acceptance-rejection method to sample this distribution. Monte Carlo integration is used to evaluate Eq. (38), and this integral is performed for every candidate particle. Computational time may be saved, at the expense of increased variance, by using only a subset of the particles in the cell for each integration. Evaluating the real and imaginary parts separately gives

$$g^{\rm gen}_{\rm re}(\tau,\mathbf{c}) = \Delta\tau\frac{N_c}{N_s}\sum_{i}^{N_s} K(\mathbf{c},\mathbf{c}_i)\tilde{\mathcal{W}}_{{\rm re},i}, \tag{39a}$$

$$g_{\text{im}}^{\text{gen}}(\tau,\mathbf{c}) = \Delta\tau \frac{N_c}{N_s} \sum_{i}^{N_s} K(\mathbf{c},\mathbf{c}_i) \tilde{\mathcal{W}}_{\text{im},i}, \tag{39b}$$

where N_s is the size of the subset of particles in each cell and the index i again refers to the ith particle.

Equations (39a) and (39b) are calculated for N_g candidate particles (Eq. (31)), which are then accepted or rejected according to Eq. (32). If accepted, the real and imaginary particle values are then set according to Eqs. (33)-(36). Upon completion of the particle generation step, the particles earlier marked for deletion are removed from the simulation.

d. Advection operation. Advection operation, Eq. (23a), is solved deterministically by moving each particle in a straight line, and updating the location accordingly,

$$\mathbf{x}(t+\Delta\tau) = \mathbf{x}(\tau) + \Delta\tau\mathbf{c}. \tag{40}$$

Boundary interactions are handled during the advection stage. First, particles crossing a boundary are removed from the simulation.

Open boundaries are then modeled using a reservoir of gas which is in equilibrium with specified density, bulk velocity, and temperature perturbations, given by $\tilde{\sigma}_b$, $\tilde{\mathbf{u}}_b$, and $\tilde{\kappa}_b$, respectively; the subscript b again indicates the boundary. The generating distribution at the boundary is therefore

$$\tilde{g}_b^{\text{gen}}(\tau,\mathbf{c}) = A_b \Delta\tau \bar{f}^{\text{eq}}(\mathbf{c})\, \mathbf{c}\cdot\mathbf{n} \left\{ \tilde{\sigma}_b + \frac{\mathbf{c}\cdot\tilde{\mathbf{u}}_b}{RT_0} + \left(\frac{c^2}{2RT_0} - \frac{3}{2} \right) \tilde{\kappa}_b \right\}, \quad \mathbf{c}\cdot\mathbf{n} > 0, \tag{41}$$

where A_b is the area of the boundary in question. Particles are generated using this distribution in a similar fashion to the collision step, i.e., Eqs. (30)-(36). In addition to the procedures detailed above, we restrict candidate particles to velocities with $\mathbf{c}\cdot\mathbf{n} > 0$.

For *solid boundaries*, the density perturbation is given by

$$\tilde{\sigma}_b = \tilde{\mathbf{u}}_b \cdot \mathbf{n} \sqrt{\frac{\pi}{2RT_0}} - \frac{1}{2}\tilde{\kappa}_b + \tilde{\sigma}_r. \tag{42}$$

The net mass of particles crossing a solid boundary over the virtual-time step is $\tilde{m}_r$, and $\tilde{\sigma}_r$ is chosen to ensure that the same mass of particles is re-emitted, i.e.,

$$\tilde{\sigma}_r = \frac{\tilde{m}_r N_p}{A_b \Delta\tau\, n_0} \sqrt{\frac{2\pi}{RT}}. \tag{43}$$

Finally, all particles generated on the (open or closed) boundary are advected for a random fraction of the virtual-time step, $\mathcal{R}\Delta\tau$.

2. Using particles to represent the deviation from local equilibrium

To represent the deviation from a local equilibrium, the effect of the operations in Eq. (23) in each cell must be considered separately for the equilibrium perturbation $\tilde{g}^{\text{eq}}$ and the particles representing $\tilde{g}^{\text{p}}$; see Eq. (24). Recall that in each cell, we first calculate the local density, bulk velocity, and temperature moments which are then used to construct $\tilde{g}^{\text{eq}}$ from Eq. (21). We denote these local moments $\tilde{\sigma}_c$, $\tilde{\mathbf{u}}_c$, and $\tilde{\kappa}_c$, respectively, where the subscript c is used to indicate cell specific values.

a. Oscillatory operation. We once again begin with oscillatory operation, Eq. (23c). As before, the particle values for $\tilde{\mathcal{W}}$ are updated using Eq. (26), and the moments for each cell are evaluated in the same fashion,

$$\begin{bmatrix} \sigma_{c,\text{re,new}} \\ \sigma_{c,\text{im,new}} \end{bmatrix} = \begin{bmatrix} \cos(\Delta\tau\omega) & -\sin(\Delta\tau\omega) \\ \sin(\Delta\tau\omega) & \cos(\Delta\tau\omega) \end{bmatrix} \begin{bmatrix} \sigma_{c,\text{re}} \\ \sigma_{c,\text{im}} \end{bmatrix}, \tag{44}$$

with $\tilde{\mathbf{u}}_c$ and $\tilde{\kappa}_c$ treated in exactly the same manner.

b. Collision operation (BGK model). When considering BGK relaxation process, Eq. (27), we note that the particle generation term in Eq. (29) is a Maxwell-Boltzmann distribution. Because we know the analytical form of this function, we may apply Eq. (29) directly to the stored value of $\tilde{g}^{\text{eq}}$,

which is evaluated (from Eq. (21)) using the stored moments of the cell,

$$\begin{aligned}
\tilde{\sigma}_{c,\text{new}} &= \tilde{\sigma}_c + (1 - e^{-\nu\Delta\tau})\frac{N_p}{n_0 V_c}\sum_i^{N_c} \mathcal{W}\tilde{W}_i, \\
\tilde{\mathbf{u}}_{c,\text{new}} &= \tilde{\mathbf{u}}_c + (1 - e^{-\nu\Delta\tau})\frac{N_p}{n_0 V_c}\sum_i^{N_c} \mathcal{W}\tilde{W}_i\mathbf{c}_i, \\
\tilde{\kappa}_{c,\text{new}} &= \tilde{\kappa}_c + (1 - e^{-\nu\Delta\tau})\frac{N_p}{n_0 V_c}\frac{2}{3}\sum_i^{N_c} \mathcal{W}\tilde{W}_i\left(\frac{c_i^2}{2RT_0} - \frac{3}{2}\right).
\end{aligned} \tag{45}$$

Equations (21) and (45) yield analytical equations for the component of $\tilde{g}$ represented by $\tilde{g}^{\text{eq}}$, eliminating the need to produce simulator particles as part of the collision process. As in Sec. II B 1, $N_c(1 - e^{-\nu\Delta\tau})$ particles are deleted from each cell.

c. Collision operation (HS model). In contrast to the BGK relaxation process, HS generating distribution, Eq. (38), does not exactly correspond to a Maxwell-Boltzmann distribution and therefore cannot be represented exactly by changes to $\tilde{g}^{\text{eq}}$. Nonetheless, we proceed, as above, by updating the stored cell moments using the moments of the generating distribution. Details regarding the moments of the generating distribution are given in Appendix A, Eq. (A5). Evaluating these moments using Monte Carlo integration, and adding them to the stored values in the cells, gives

$$\begin{aligned}
\tilde{\sigma}_{c,\text{new}} &= \tilde{\sigma}_c + \Delta\tau\frac{N_p}{n_0 V_c}\sum_i^{N_c} \hat{\nu}(\mathbf{c}_i)\mathcal{W}\tilde{W}_i = \tilde{\sigma}_c + \Delta\tilde{\sigma}_c, \\
\tilde{\mathbf{u}}_{c,\text{new}} &= \tilde{\mathbf{u}}_c + \Delta\tau\frac{N_p}{n_0 V_c}\sum_i^{N_c} \hat{\nu}(\mathbf{c}_i)\mathcal{W}\tilde{W}_i\mathbf{c}_i = \tilde{\mathbf{u}}_c + \Delta\tilde{\mathbf{u}}_c, \\
\tilde{\kappa}_{c,\text{new}} &= \tilde{\kappa}_c + \Delta\tau\frac{N_p}{n_0 V_c}\frac{2}{3}\sum_i^{N_c} \hat{\nu}(\mathbf{c}_i)\mathcal{W}\tilde{W}_i\left(\frac{c_i^2}{2RT_0} - \frac{3}{2}\right) = \tilde{\kappa}_c + \Delta\tilde{\kappa}_c.
\end{aligned} \tag{46}$$

Particles are then generated to account for the remainder of the change to $\tilde{g}$, by subtracting the change in $\tilde{g}^{\text{eq}}$ from Eq. (38),

$$\tilde{g}^{\text{gen}}(\tau, \mathbf{c}) = \Delta\tau\int_{-\infty}^{\infty} K(\mathbf{c}, \mathbf{c}_1)\tilde{g}(\tau, \mathbf{c}_1)d\mathbf{c}_1 - V_c\Delta\tilde{g}^{\text{eq}}, \tag{47}$$

where

$$\Delta\tilde{g}^{\text{eq}} = \bar{f}^{\text{eq}}\left\{\Delta\tilde{\sigma}_c + \frac{\mathbf{c}\cdot\Delta\tilde{\mathbf{u}}_c}{RT_0} + \left(\frac{c^2}{2RT_0} - \frac{3}{2}\right)\Delta\tilde{\kappa}_c\right\} \tag{48}$$

is obtained by substituting the change in the moments, $\Delta\tilde{\sigma}_c$, $\Delta\tilde{\mathbf{u}}_c$, and $\Delta\tilde{\kappa}_c$ (see Eq. (46)), into Eq. (21). Equation (48) yields the Maxwell-Boltzmann distribution "closest" to Eq. (38), minimizing the number of particles generated by Eq. (47); discussion of this approach is given in Ref. 36.

At the beginning of the collision step, particles are again marked for deletion with probability $\Delta\tau\hat{\nu}(\mathbf{c})$ and removed at the end of the operation.

d. Advection operation. The advection operation is applied by inserting Eq. (23b) into Eq. (24) and considering each part separately,

$$\left.\frac{\partial\tilde{g}^{\text{d}}}{\partial\tau}\right|_{\text{adv}} = -\mathbf{c}\cdot\frac{\partial\tilde{g}^{\text{d}}}{\partial\mathbf{x}}, \tag{49a}$$

$$\left.\frac{\partial\tilde{g}^{\text{eq}}}{\partial\tau}\right|_{\text{adv}} = -\mathbf{c}\cdot\frac{\partial\tilde{g}^{\text{eq}}}{\partial\mathbf{x}}. \tag{49b}$$

To implement Eq. (49a), the particles representing $\tilde{g}^{\text{d}}$ are advected in the usual manner; see Eq. (40). Equation (49b) represents advection due to differing values of the equilibrium perturbation in adjacent cells, i.e., $\tilde{g}_+^{\text{eq}}$ and $\tilde{g}_-^{\text{eq}}$. Its solution is given by generating particles at the cell boundaries, with

the distribution given by

$$\tilde{g}_f^{\text{gen}}(\tau,\mathbf{c}) = A_f \Delta\tau\, \mathbf{c}\cdot\mathbf{n}\left[\tilde{g}_-^{\text{eq}}(\tau,\mathbf{c}) - \tilde{g}_+^{\text{eq}}(\tau,\mathbf{c})\right], \tag{50}$$

where A_f indicates the area of the cell interface, and the normal, $\mathbf{n}$, is oriented from minus to plus; the subscript f refers to the cell interface. Particles are generated from the distribution, $\tilde{g}_f^{\text{gen}}(\tau,\mathbf{c})$, placed randomly on the cell interface and advected for a random fraction of the virtual-time step, $\mathcal{R}\Delta\tau$.

As before, particles crossing boundaries are removed from the simulation. The generating distribution of the boundaries is now given by

$$\tilde{g}_b^{\text{gen}}(\tau,\mathbf{c}) = A_b \Delta\tau\, \bar{f}^{\text{eq}}(\mathbf{c})\, \mathbf{c}\cdot\mathbf{n}\left\{\tilde{\sigma}_m + \frac{\mathbf{c}\cdot(\tilde{\mathbf{u}}_b - \tilde{\mathbf{u}}_c)}{RT_0} + \left(\frac{c^2}{2RT_0} - \frac{3}{2}\right)(\tilde{\kappa}_b - \tilde{\kappa}_c)\right\}, \quad \mathbf{c}\cdot\mathbf{n} > 0, \tag{51}$$

where subscript c is used here to refer to the cell adjacent to the boundary. For *open* boundaries,

$$\tilde{\sigma}_m = \tilde{\sigma}_b - \tilde{\sigma}_c, \tag{52}$$

and for *closed* boundaries,

$$\tilde{\sigma}_m = (\tilde{\mathbf{u}}_b - \tilde{\mathbf{u}}_c)\cdot\mathbf{n}\sqrt{\frac{\pi}{2RT_0}} - \frac{1}{2}(\tilde{\kappa}_b - \tilde{\kappa}_c) + \tilde{\sigma}_r. \tag{53}$$

Once again, each particle generated from Eq. (51) is placed randomly on the boundary and advected for a random fraction of the virtual-time step, $\mathcal{R}\Delta\tau$.

C. Algorithm summary

We now summarize each step of the deviational frequency-domain algorithm as follows.

1. *Particle movement*: Particle positions are updated according to Eq. (23a). If a local equilibrium perturbation is used, as per Eq. (24), particles are generated on the cell interfaces with Eq. (50) and advected for a random fraction of the virtual-time step; no such operation exists when using the global equilibrium.
2. *Boundary interaction*: Particles intersecting a boundary during the virtual-time step are removed from the simulation. For solid boundaries, the net mass of these particles is recorded; this operation is not performed for open boundaries. Particles are then generated on the boundary, using Eq. (41) for the deviation from global equilibrium, or Eq. (51) for the deviation from local equilibrium. In each case, the recorded mass of simulator particles crossing the boundary is used to calculate $\tilde{\sigma}_r$. The generated particles are advected for a random fraction of the virtual-time step.
3. (a) *Particle relaxation* (BGK): In each cell, the local density, bulk velocity, and temperature moments due to the particles are calculated using Eq. (28). The fraction $(1 - e^{-\nu\Delta\tau})$ of these particles is then removed from the simulation. When generating the deviation from a global equilibrium, the moments are used to construct Eq. (29), from which new particles are drawn. From a local equilibrium, these values are used to update the stored moments of the cell, as in Eq. (28).
   (b) *Particle collisions* (HS): Particles are marked for deletion with probability $\Delta\tau\hat{\nu}(\mathbf{c})$ in each cell. When simulating the deviation from global equilibrium, Eq. (38) is used to generate new particles. For a local equilibrium, the local density, bulk velocity, and temperature moments due to the particles are calculated. These values are used to (i) update the stored moments of the cell using Eq. (46) and (ii) construct Eq. (47) from which new particles are generated. Finally, the particles marked for deletion are removed.
4. *Particle oscillation*: The value of every particle is updated according to Eq. (26). The stored moments for each cell are also updated when using a local equilibrium, as per Eq. (44); this step does not exist when using a global equilibrium.

For simplicity, an asymmetric time/virtual-time splitting algorithm is used. We note that convergence may be improved through the use of symmetrized splitting;[17,40,41] its incorporation should be straightforward, although this has not been investigated.

Importantly, setting ω to zero recovers the original time-domain algorithms[17,20]—these can be used to simulate time-dependent (transient) flows with the addition of time-dependent boundary conditions.

III. TEST CASES

To assess the deviational frequency-domain Monte Carlo method, two standard oscillatory flows are considered: (i) one-dimensional Couette flow between two flat surfaces and (ii) the three-dimensional flow generated by a sphere oscillating in an unbounded gas.

For Couette flow, simulations are performed using both of the approaches detailed in Sec. II, i.e., deviation from local and global equilibria. For low frequency flows, we simulate the deviation from a local equilibrium, whereas for high frequency flows, deviation from the global equilibrium is used, as discussed below. For each of these complementary approaches, we compare the convergence properties of the time and frequency-domain methods as a function of both cell size and time/virtual-time step size. Frequency-domain BGK solutions are then validated against existing accurate numerical solutions of the Boltzmann-BGK equation.[35] Frequency-domain HS solutions are validated by comparison to the corresponding time-domain simulations (obtained by setting $\omega = 0$ in the frequency-domain method; see above).

For the 3D flow generated by an oscillating sphere, simulations are performed of the deviation from the global equilibrium only. We make a single comparison of the frequency-domain BGK and HS solutions to their corresponding time-domain solutions. Additionally, the deviational frequency-domain BGK solution is compared to that obtained using the weighted frequency-domain method, results of which are presented in Ref. 34.

The relaxation frequency ν is a free parameter in the BGK model. In all BGK simulations, we relate this value to the mean free path via the most probable molecular speed, i.e., $\nu = \sqrt{2RT_0}/\lambda$. Henceforth, we make use of scaled variables, which are denoted with "*." Time and virtual-time are scaled by the relaxation frequency, with $t^* = t\nu$ and $\tau^* = \tau\nu$, respectively. Spatial position is scaled by the geometric length scale of the flow, $\mathbf{x}^* = \mathbf{x}/L$, in both the time and frequency-domain methods. For Couette flow, L is the distance between the two plates, whereas for the oscillating sphere, L is its radius.

When performing HS simulations, a modified Knudsen number $\hat{\mathrm{Kn}} = \mathrm{Kn}\sqrt{\pi}/2$ is used.[42] This matches the HS viscosity to the BGK viscosity. However, given the differing Prandtl numbers (Pr) of the two models, Pr = 1 (BGK) and Pr = 2/3 (HS),[43] we cannot simultaneously match their thermal diffusivities.

An extensive review of BGK and HS solutions for one-dimensional steady isothermal flows is given in Ref. 44, where good quantitative agreement between the two collision models is found. Our results confirm that this agreement holds for oscillatory Couette flow (in Fig. 7). We also examine the consequences of the above-mentioned Prandtl number difference for the oscillating solid sphere (in Fig. 11), which inherently generates a non-isothermal flow.

Simulation results are given in the form of time/virtual-time series of the instantaneous bulk flow properties, such as density, bulk velocity, stress, and temperature. Because Monte Carlo simulations inherently contain stochastic noise, these series are analyzed using a least-squares fit. Specifically, series for the frequency and time-domain methods are fit to the functions

$$\tilde{\Gamma}(\tau^*, \gamma, \zeta) = \gamma + \zeta \mathrm{i}, \tag{54a}$$

$$\Gamma(t^*, \xi, \mu) = \xi \cos(\theta t^*) + \mu \sin(\theta t^*), \tag{54b}$$

respectively, where γ and ζ are the real and imaginary fit parameters in the frequency-domain, while ξ and μ are the corresponding in-phase and out-of-phase fit parameters in the time-domain.

Simulations are run until a steady state is achieved, following which the above fit functions are applied. Each simulation is then time/virtual-time stepped, and the fit reevaluated, until the variation in the above parameters is less than 1%.

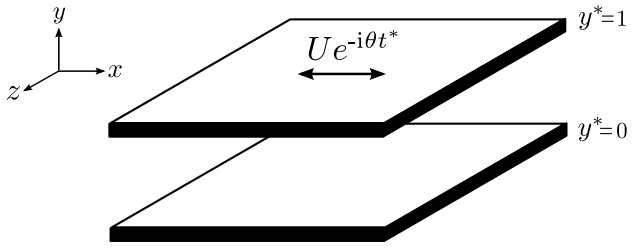

FIG. 2. Schematic of oscillatory Couette flow geometry showing the coordinate system and the time-dependent velocity of the upper plate. The bottom plate is stationary. Reprinted with permission from D. R. Ladiges and J. E. Sader, J. Comput. Phys. **284**, 351 (2015). Copyright 2015 Elsevier.

A. Oscillatory plane Couette flow

We begin with the oscillatory Couette flow problem. The flow geometry is bounded in the y-direction by two infinite parallel plates at $y^* = 0$ and $y^* = 1$; see Fig. 2. The domain is unbounded in the x- and z-directions and discretized into cells of height $\Delta y^* \ll \min(\mathrm{Kn}, 1)$. The upper plate oscillates in the x-direction with velocity amplitude U and scaled angular frequency θ, whilst the lower plate is stationary. The solid plate boundaries are assumed to be diffusely reflecting with temperature T_0, as described in Section II.

Figures 3 and 4 show the percentage relative error/difference in the simulated shear stress on the oscillating plate, obtained using the deviational frequency-domain and time-domain Monte

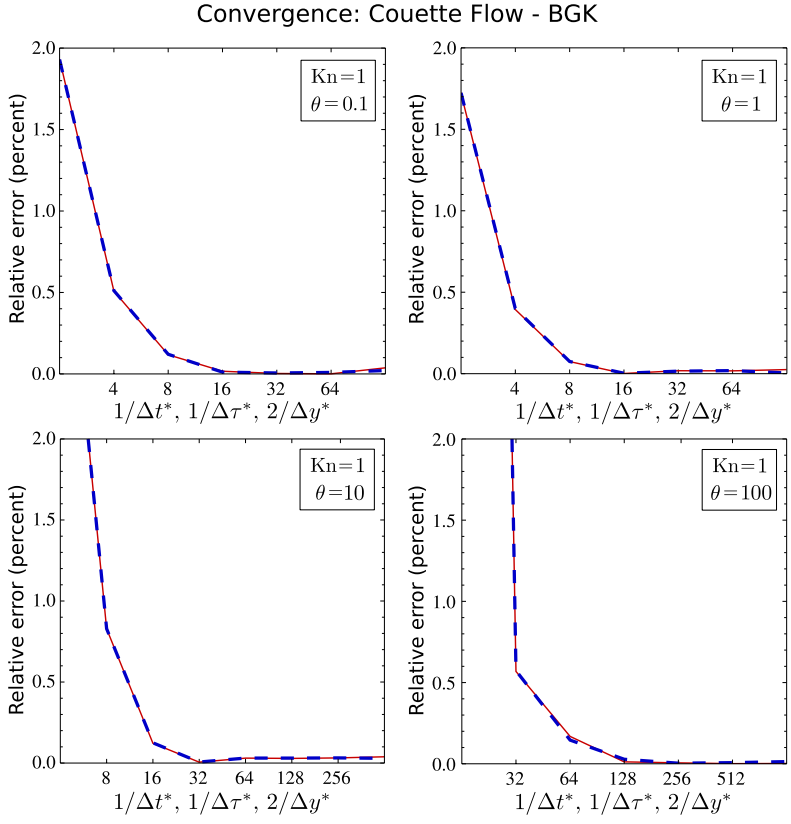

FIG. 3. [BGK model, oscillatory Couette flow]: relative percentage error in shear stress at oscillating plate as a function of step and cell size (which are varied simultaneously); benchmark in Ref. 35. Red solid line: time-domain solution, in-phase shear stress. Blue dashed line: frequency-domain solution, real shear stress. In each simulation, Δy^* has been selected such that the number of cells per mean free path is double the number of steps per relaxation period. For $\theta = 0.1$ and $\theta = 1$, the deviation from local equilibrium is simulated. For $\theta = 10$ and $\theta = 100$, deviation from global equilibrium is used. All simulations are run with $N_p = 5 \times 10^9$.

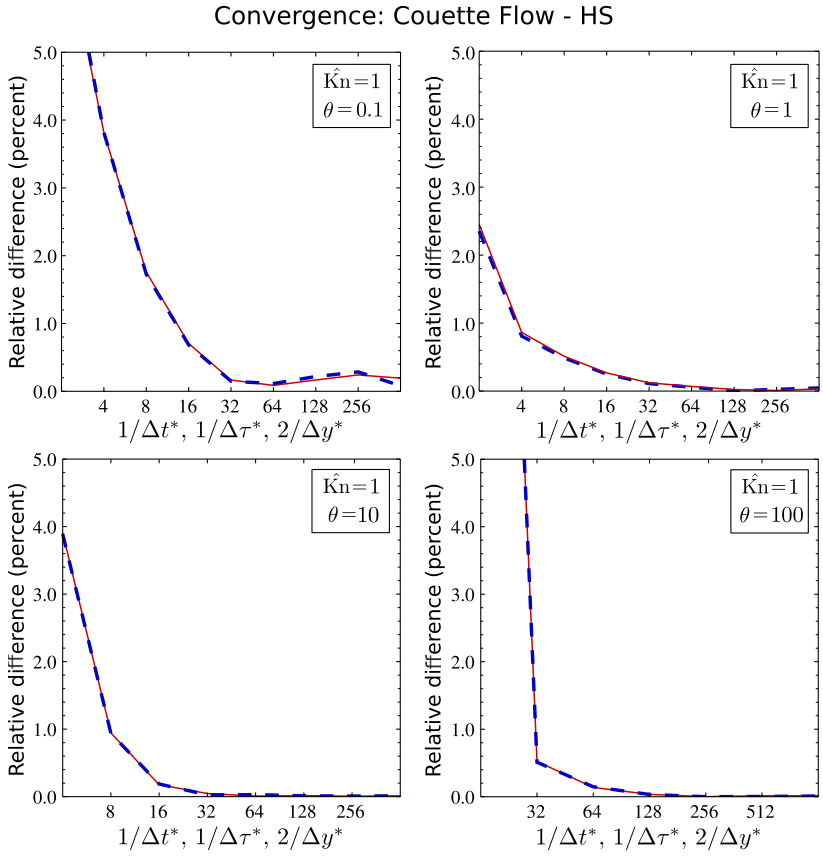

FIG. 4. [HS model, oscillatory Couette flow]: relative percentage difference in shear stress at oscillating plate as a function of step and cell size; relative to the average result for the smallest two step/cell sizes, which are varied simultaneously. See caption of Fig. 3 for simulation parameters and description.

Carlo methods with the BGK and HS operators, respectively. The error is shown as a function of step and cell size, which are varied simultaneously; see the captions of Figs. 3 and 4. The high accuracy numerical results in Ref. 35 are used as the benchmark for the BGK simulations, whilst the specified difference in the HS simulations is given relative to the average result for the smallest two step/cell sizes. For the frequency-domain simulations, the real component of the shear stress is reported. Equivalently, the in-phase component of the time-domain simulations is shown. For smaller normalized frequencies, i.e., $\theta = 0.1$ and $\theta = 1$, the deviation from a local equilibrium is simulated, which is more efficient than simulating deviations from the global equilibrium. For larger normalized frequencies, simulating the deviation from global equilibrium is found to be more efficient, requiring a smaller number of particles. Note that the time and frequency-domain methods exhibit nearly identical convergence with respect to step and cell size, regardless of the normalized frequency or type of deviation used.

Next, we examine the bulk velocity profiles for a range of Knudsen numbers and scaled oscillation frequencies. Figure 5 compares the deviational frequency-domain BGK results to the high accuracy literature solutions,[35] whereas Fig. 6 provides a comparison of the deviational frequency-domain HS results and the corresponding time-domain solutions. In each case, the frequency-domain method is found to accurately capture the bulk velocity profile over the entire flow domain.

It is also interesting to compare the results of the BGK and HS models, to explore the quantitative differences in their predictions. Figure 7 provides this comparison, showing that the flow profiles are very similar — they virtually overlap. This comparison demonstrates that the BGK model is capable of accurately modeling more realistic oscillatory gas flows.

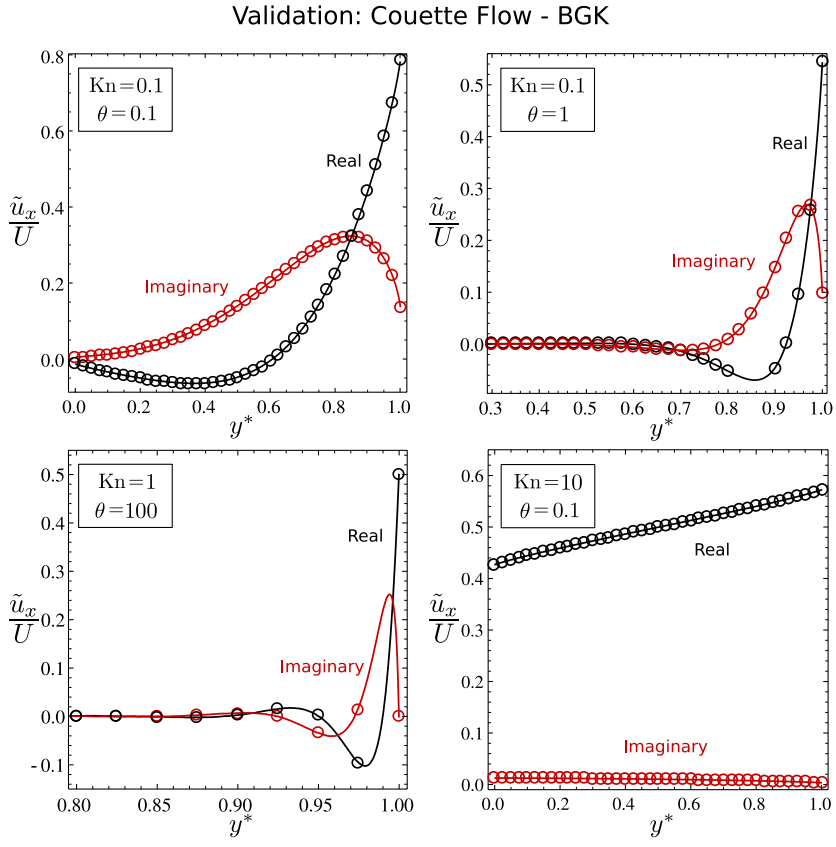

FIG. 5. [BGK model, oscillatory Couette flow]: streamwise bulk velocity profiles. Solid lines: frequency-domain results. Open circles: published high-accuracy numerical results.[35] For the $\theta = 0.1$ and $\theta = 1$ cases, $1/\Delta\tau^* , 1/\Delta t^* = 32$, $\mathrm{Kn}/\Delta y^* = 60$, $N_p = 5\times10^9$, and the deviation from local equilibrium is simulated. For the $\theta = 100$ case, $1/\Delta\tau^* , 1/\Delta t^* = 1024$, $\mathrm{Kn}/\Delta y^* = 2048$, $N_p = 2.5\times10^9$, and the deviation from global equilibrium is simulated.

Collectively, these results show that the deviational frequency-domain method provides a robust method to accurately simulate oscillatory unidirectional flows. This provides a complementary approach to the previously described weight-based frequency-domain method,[34] whilst more easily accommodating a broader range of collision models. The relative computational requirements of the time and frequency-domain methods are discussed in Section IV. An example extending the use of the frequency-domain method beyond strictly sinusoidal flows is given in Section V.

B. 3D oscillating sphere

Next, we consider a sphere of radius L, with a diffusely reflecting surface of fixed temperature T_0, centered at the origin of a spherical domain of radius R_b; see Fig. 8. The sphere executes rectilinear oscillations along the x-axis, with velocity amplitude U and scaled frequency θ. The outer boundary is open, as described in Section II, with density, bulk velocity, and temperature set at the global equilibrium values. The $z = 0$ plane is discretized radially into segments of length Δr. These segments are further divided into cells subtending an angle $\Delta\alpha$. For each segment, this angle is selected such that $r\Delta\alpha \approx \Delta r$ whilst fitting an integer number of cells at the given radius. The flow in question is symmetric around the x-axis. Consequently, the collision/relaxation process is independent of ϕ, and we are therefore able to rotate these cells through the entire coordinate, $\phi \in [0, 2\pi)$.

The collision/relaxation process is performed in a spherical coordinate frame. Because the advection process is implemented in the usual Cartesian frame, particle velocities must be converted into spherical coordinates to evaluate the integrals given by Eqs. (28), (39), (45), and (46). This is

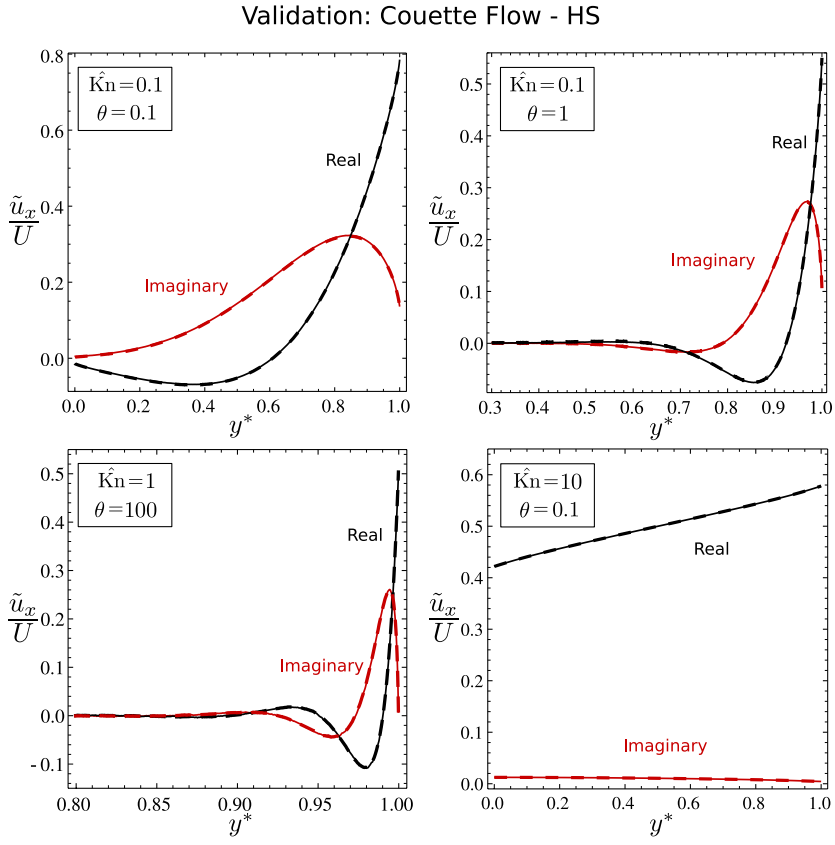

FIG. 6. [HS model, oscillatory Couette flow]: streamwise bulk velocity profiles. Solid lines: time-domain results. Dashed lines: frequency-domain results. For the $\theta = 0.1$ and $\theta = 1$ cases, $1/\Delta\tau^*, 1/\Delta t^* = 128$, $\mathrm{Kn}/\Delta y^* = 250$, $N_p = 5\times10^9$, and the deviation from local equilibrium is simulated. For the $\theta = 100$ case, $1/\Delta\tau^*, 1/\Delta t^* = 4096$, $\mathrm{Kn}/\Delta y^* = 8192$, $N_p = 0.6\times10^9$, and the deviation from global equilibrium is simulated.

achieved using

$$\begin{bmatrix} c_r \\ c_\alpha \\ c_\phi \end{bmatrix} = \begin{bmatrix} \cos\alpha & \sin\alpha & 0 \\ -\sin\alpha & \cos\alpha & \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} 1 & 0 & 0 \\ 0 & \cos\phi & \sin\phi \\ 0 & -\sin\phi & \cos\phi \end{bmatrix} \begin{bmatrix} c_x \\ c_y \\ c_z \end{bmatrix}. \tag{55}$$

Due to the high computational cost of three-dimensional simulations, we restrict our analysis to a comparison of the time and frequency-domain methods at Kn = 1 (BGK), $\hat{\mathrm{Kn}}$ = 1 (HS), and θ = 1—this represents a strongly non-continuum flow. The method of simulating the deviation from global equilibrium is used, because the approach of using a piecewise-constant function to represent the local equilibrium is not efficient for 3-dimensional flows; see Refs. 37 and 38 for discussion of more advanced approaches that may be applied in this case. The radius of the computational domain is incremented in steps of $5L$, until the relative changes in both the real (frequency-domain) and in-phase (time-domain) drags are less than 1%; a value of $R_b^* = 20$ fulfills this requirement.

Figures 9 and 10 show the BGK and HS solutions, respectively. The temperature, density, and bulk velocity along the axis of rectilinear oscillation are shown, along with the angular slip velocity on the surface of the sphere. The time and frequency-domain results obtained using the deviational method are given, along with frequency-domain results using the weight-based method described in Ref. 34. In all cases, we observe close agreement, demonstrating the validity of the deviational frequency-domain method for higher-dimensional flows.

Figure 11 compares the frequency-domain BGK and HS results using the deviational frequency-domain method. This comparison is particularly interesting because the generated flow

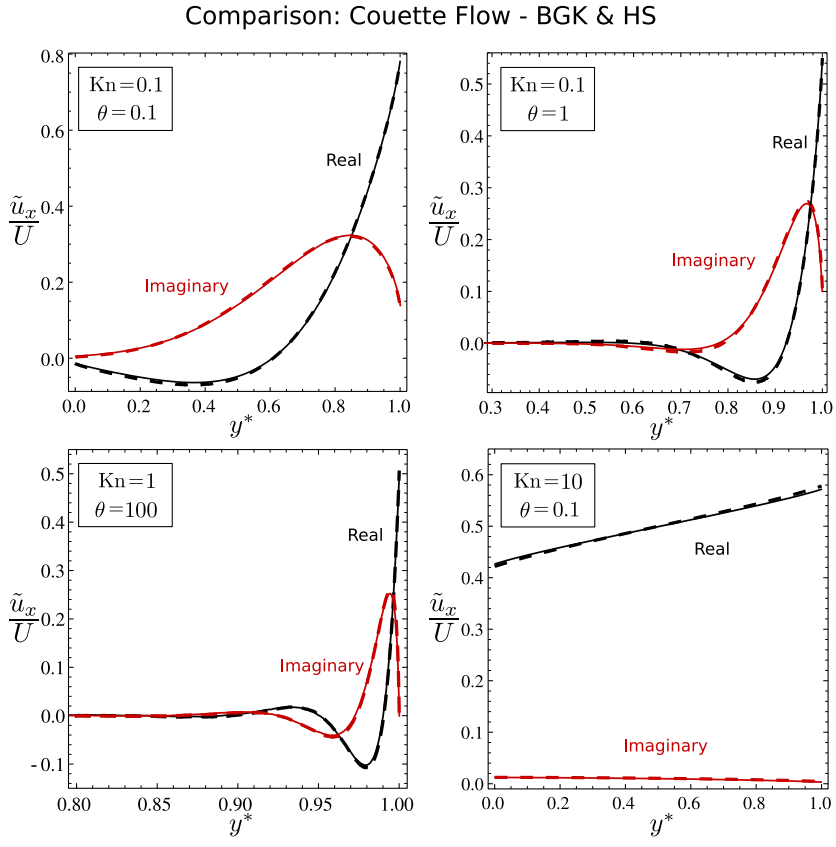

FIG. 7. [BGK and HS model comparison, oscillatory Couette flow]: streamwise bulk velocity profiles. Solid lines: frequency-domain BGK results. Dashed lines: frequency-domain HS results. See captions of Figs. 5 and 6 for simulation parameters.

is non-isothermal, due to the presence of a temperature jump at the solid sphere surface — this arises because the surface orientation is not parallel to the oscillation direction everywhere.[12] As discussed above, the HS and BGK models exhibit different Prandtl numbers: for the BGK model, we have Pr = 1, whilst Pr = 2/3 for the HS model.[43] Since the Prandtl number controls the relative diffusion of vorticity and heat, a significant difference in the predictions of these two models may be expected. Even so, excellent agreement between the predictions of the two models is observed for the bulk velocity and density profiles. Further, even the temperature profiles agree well, albeit

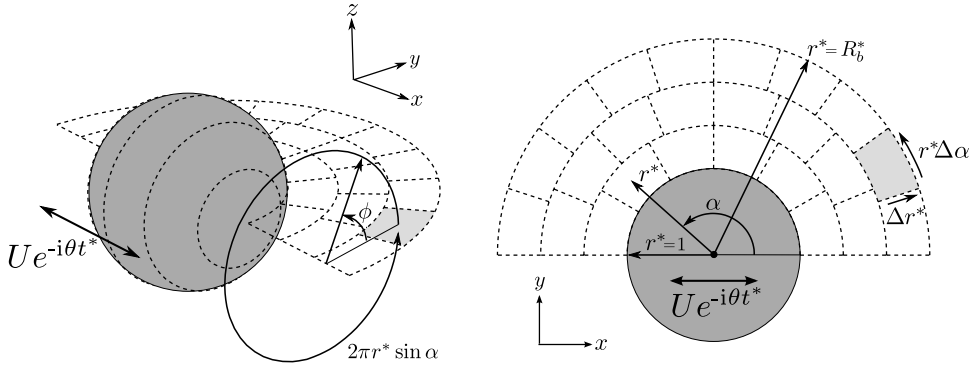

FIG. 8. Schematic of oscillating sphere geometry showing coordinate system. Left: 3D perspective. Right: $z = 0$ cross section (x-y plane). Reprinted with permission from D. R. Ladiges and J. E. Sader, J. Comput. Phys. **284**, 351 (2015). Copyright 2015 Elsevier.

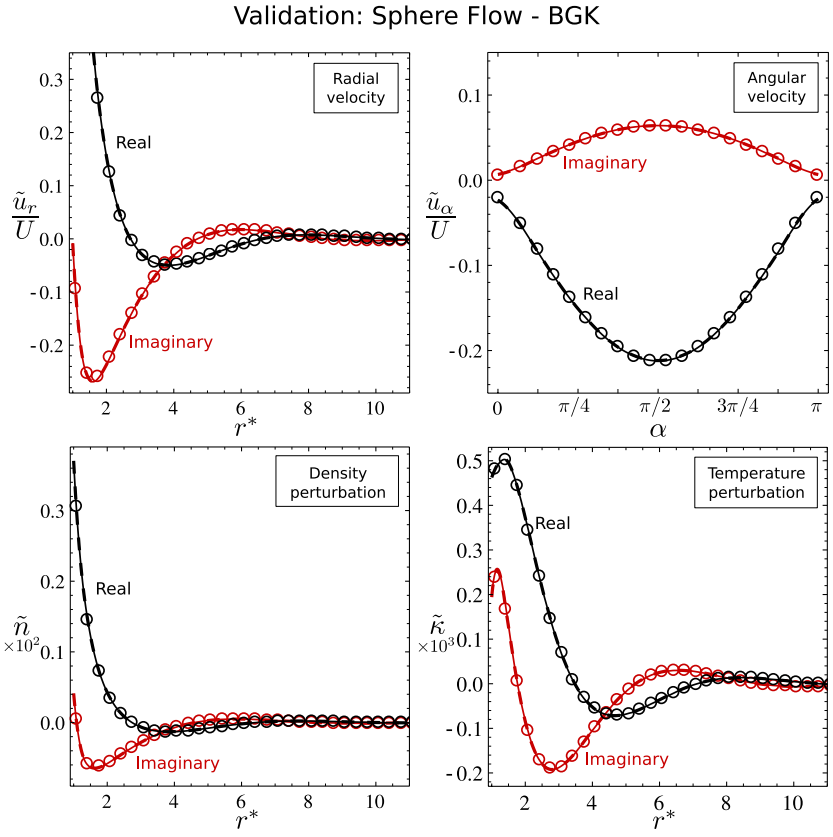

FIG. 9. [BGK model, oscillating sphere]: density, bulk velocity, and temperature profiles. Solid lines: time-domain results. Dashed lines: frequency-domain results. Open circles: results obtained using weighted frequency-domain method.[34] Simulations run using $N_p = 32\times10^9$, $1/\Delta r^* = 6$, and $1/\Delta\tau^*, 1/\Delta t^* = 64$.

with slightly greater differences occurring. This highlights the utility of the BGK model for oscillatory flows, which can provide good quantitative predictions for higher-dimensional non-continuum flows, that are generally non-isothermal due to the inevitable presence of a temperature jump at a solid wall.[12]

IV. COMPUTATIONAL PERFORMANCE

As discussed in Section III, bulk flow properties (such as the density, velocity, and temperature) are obtained using a least-squares fit to the time/virtual-time series of instantaneous values; see Eq. (54). Because these series inherently contain stochastic noise, the observed variances in all bulk flow properties decrease with an increasing number of time/virtual-time steps, i.e., as a larger sample is obtained. We refer to this reduction in variance as the *statistical convergence*.

The total computational time of a simulation depends on two factors: (i) the statistical convergence and (ii) the time required to perform each time/virtual time step. In this section, we compare these factors in the deviational time and frequency-domain methods, to determine the relative computational performance of the methods.

A. Statistical convergence

Formulas describing the relative statistical convergence of the weight-based time and frequency-domain Monte Carlo methods were previously derived in Ref. 34, using the theoretical framework of Ref. 45. These formulas also apply to the deviational method proposed in this

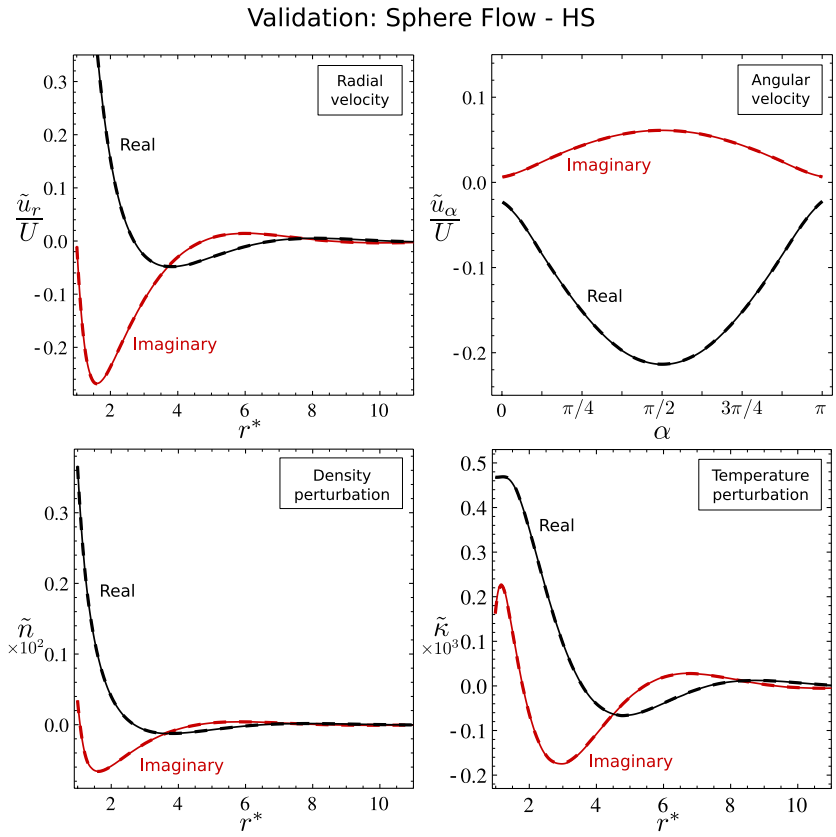

FIG. 10. [HS model, oscillating sphere]: density, bulk velocity, and temperature profiles. Solid lines: time-domain results. Dashed lines: frequency-domain results. Simulations run using $N_p = 32\times10^9$, $1/\Delta r^* = 6$, and $1/\Delta\tau^*, 1/\Delta t^* = 64$.

article and give the variance of the fit parameters in Eq. (54), as a function of the number of time/virtual-time steps. In Ref. 34, it was shown that the weight-based frequency-domain method exhibits superior statistical convergence relative to the corresponding time-domain method. Here, we investigate the relative statistical convergence of the proposed deviational frequency-domain method.

The formulas in Ref. 34 show that the relative statistical convergence of the deviational frequency and time-domain methods may be described in two asymptotic limits: $\beta M \ll 1$ and $\beta M \gg 1$. Here, M is the number of time/virtual-time steps in the simulation, and $\beta = \theta\Delta t^* = \theta\Delta\tau^*$ is proportional to the ratio of the time (or virtual-time) step to the oscillation period. When $\beta \ll 2\pi$, many time (or virtual-time) steps occur in an oscillation period, and $\beta \gg 2\pi$ indicates that a single step encompasses many oscillation periods.

In the small βM limit, the time-domain fit parameters exhibit a variance of $O\left((\beta M)^{-2}M^{-1}\right)$, whereas the variance of the frequency-domain fit parameters is of $O(M^{-1})$. Statistical convergence of the frequency-domain method (i.e., with increasing M) is therefore superior in this region. In the large βM limit, both methods display $O(M^{-1})$ behavior. Importantly in this limit, the variance of each fit parameter (ξ, μ) in the time-domain is given by a constant multiple of the corresponding variance in the frequency-domain,

$$\mathrm{Var}(\xi) = \frac{3}{2}\left(1 + \frac{1}{3}a\right)\mathrm{Var}(\gamma), \tag{56a}$$

$$\mathrm{Var}(\mu) = \frac{3}{2}\left(1 + \frac{1}{3}a^{-1}\right)\mathrm{Var}(\zeta), \tag{56b}$$

where a is a positive constant dependent on the properties of the simulations.[34] It follows that each time-domain fit parameter (ξ, μ) must have a variance more than 1.5 times greater than the

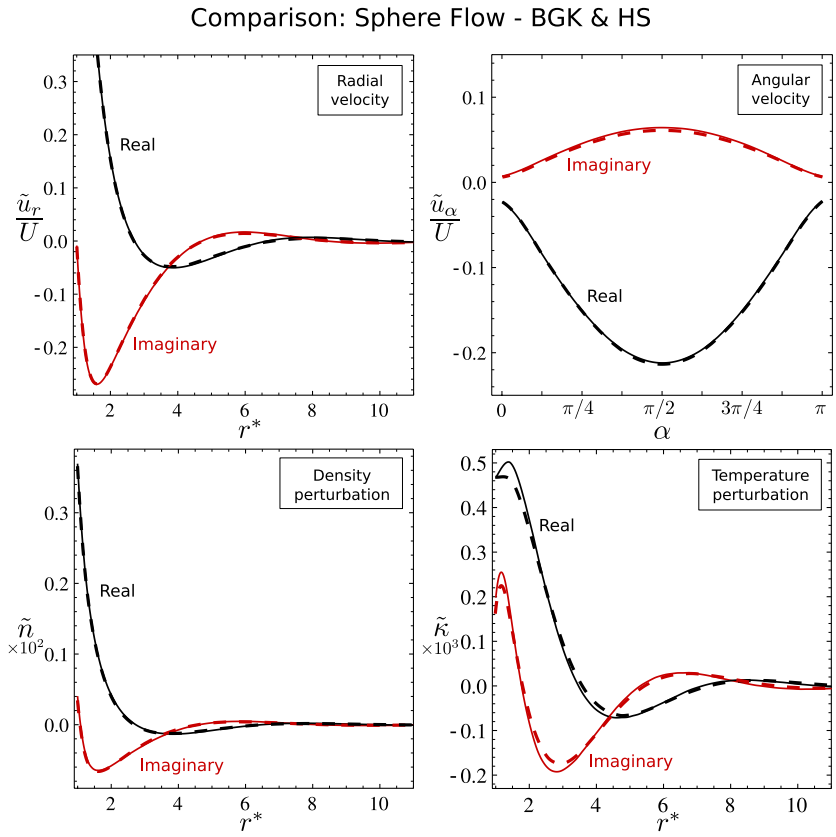

FIG. 11. [BGK and HS comparison, oscillating sphere]: density, bulk velocity, and temperature profiles. Solid lines: frequency-domain BGK results. Dashed lines: frequency-domain HS results. See captions of Figs. 9 and 10 for simulation parameters.

corresponding frequency-domain parameter (γ, ζ). Additionally, if one time-domain parameter has variance close to this minimum value, the variance of the other parameter will tend to be larger.[46] Next, we assess the validity of these theoretical predictions by direct comparison to simulation results using both time and frequency-domain methods.

Figure 12 presents results for the two asymptotic limits described above, using two of the oscillatory Couette flow cases from Section III A. These results focus on the in-phase component of the time-domain simulations and the corresponding real component of the frequency-domain solution. Results for a small value of β (and βM) are given in the left panel of Fig. 12 as a function of the number of steps, M. These results show that variance (standard deviation is plotted) of the time-domain fit parameter is orders of magnitude larger than that of the corresponding frequency-domain result. Results for large βM are shown in the right panel of Fig. 12, corresponding to the region where variance in each method is related by a constant factor; in this case, the factor is 1.6, see Eq. (56). In all cases, excellent agreement is achieved between the above-mentioned formulas and the observed statistical convergence in the simulation results; further examples are given in Appendix B.

Figure 13 shows comparisons for both the real/in-phase and imaginary/out-of-phase parameters, for the oscillating sphere problem of Section III B. Again, excellent agreement is found between the formulas and simulations. In the region $\beta M \gg 1$, ratio of the real and in-phase variance is 1.6, close to the minimum value. As expected from the above discussion, the ratio of the imaginary and out-of-phase variances is much larger; a value of 8.1 is observed.

The results in Figs. 12 and 13, together with those in Appendix B, show that the deviational frequency-domain method exhibits superior statistical convergence relative to the corresponding time-domain method.

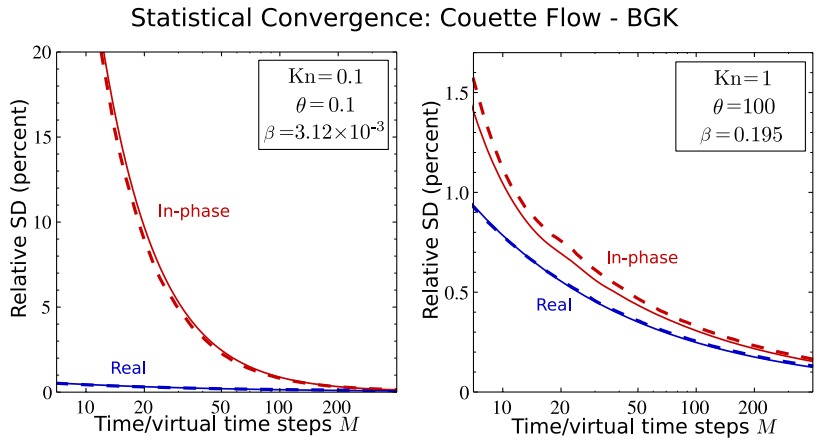

FIG. 12. [BGK model, oscillatory Couette flow]: percentage relative standard deviation in the shear stress at the oscillating wall. Red dashed lines: time-domain simulations, in-phase component. Blue dashed lines: frequency-domain simulations, real component. Solid lines: analytical predictions of equations given in Ref. 34. Ratio of the standard deviations of the time and frequency-domain results is theoretically predicted to be greater than 1.22 for a large number of steps. After 400 steps, data in the left panel ($\beta M = 1.25$) give a ratio of 1.64. Data in the right panel ($\beta M = 78$) give a ratio of 1.26.

B. Computational expense per step

We begin by noting that the computational expense on a *per particle* basis is similar for the time and frequency-domain deviational Monte Carlo methods. Some operations are duplicated between the real and imaginary parts in a frequency-domain simulation, whilst additional expense is incurred generating simulator particles from time varying distributions in the time-domain. Even so, the majority of operations are identical between the two methods.

Time and frequency-domain simulations will, in general, not contain the same number of simulator particles. As a consequence, the relative computational efficiency of the two methods cannot be inferred entirely from the statistical convergence as a function of the number of time/virtual-time steps. The computational time taken to complete each step must also be considered. This contrasts to a weight-based Monte Carlo approach, where time and frequency-domain simulations employ an equal (and constant) number of particles.

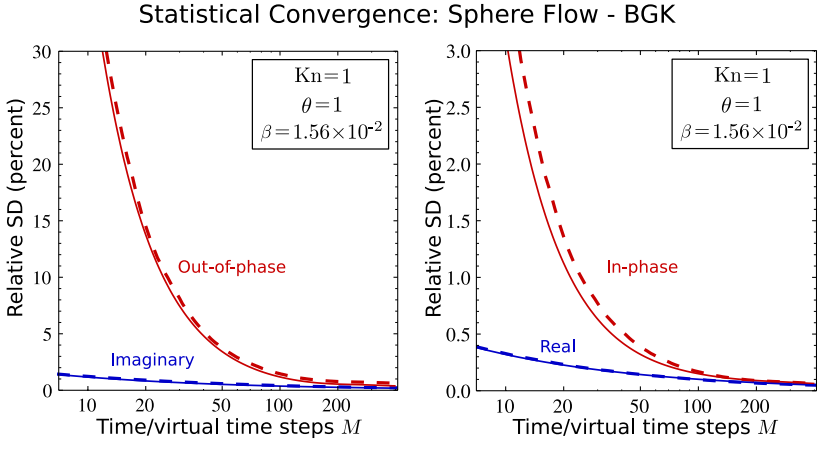

FIG. 13. [BGK model, oscillating sphere]: percentage relative standard deviation in the drag. Red dashed lines: time-domain simulations. Blue dashed lines: frequency-domain simulations. Solid lines: analytical predictions of equations given in Ref. 34. Left: imaginary and out-of-phase components. Right: real and in-phase components. Ratio of the imaginary and out-of-phase standard deviations after 400 steps is 2.85, whilst the ratio of the real and in-phase standard deviations is 1.28. Each of these values agrees with the theoretically predicted minimum of 1.22 when $\beta M \gg 1$.

The ratio of the number of particles in a frequency-domain simulation to that in the corresponding time-domain simulation, denoted N_{ratio}, cannot be predicted without prior knowledge of the flow field. However, bounds on N_{ratio} can be estimated by examining the components of the algorithm where particles are generated: the collision/relaxation process, boundary interactions, and flux between cells. In each of these components, (i) particles are created by sampling from a distribution which we refer to as a "generating distribution," and (ii) the number of particles created is proportional to the magnitude of this distribution. In the frequency-domain method, these generating distributions are given by Eqs. (29), (38), (41), (47), (50), and (51). Here, we focus on boundary generating distribution, Eq. (41), noting that the same analysis holds for the other distributions.

In a time-domain simulation, Eq. (41) is real-valued and varies sinusoidally with time,

$$g_b^{\text{gen}}(t,\mathbf{c}) = A_b \Delta\tau \bar{f}^{\text{eq}}(\mathbf{c})\, \mathbf{c}\cdot\mathbf{n} \left\{ \sigma_b + \frac{\mathbf{c}\cdot\mathbf{u}_b}{RT_0} + \left(\frac{c^2}{2RT_0} - \frac{3}{2} \right) \kappa_b \right\} \cos(\omega t), \quad \mathbf{c}\cdot\mathbf{n} > 0. \tag{57}$$

We designate

$$\left\langle g_b^{\text{gen}}(t) \right\rangle = Y \int_{\mathbf{c}\cdot\mathbf{n}>0} |g_b^{\text{gen}}(t,\mathbf{c})|\, d\mathbf{c}, \tag{58}$$

where Y is a normalization constant chosen such that $\left\langle g_b^{\text{gen}}(t) \right\rangle$ has an amplitude of 1, giving a time-averaged value of $2/\pi$, i.e., the average magnitude of a sinusoid with unit amplitude; see Fig. 14. The number of particles generated from Eq. (57), in a time-domain simulation, is proportional to $\left\langle g_b^{\text{gen}}(t) \right\rangle$.

Similarly, in a frequency-domain simulation, we define

$$\left\langle \tilde{g}_b^{\text{gen}}(\tau) \right\rangle = Y \int_{\mathbf{c}\cdot\mathbf{n}>0} \max\{|g_{\text{re}}^{\text{gen}}(\tau,\mathbf{c})|, |g_{\text{im}}^{\text{gen}}(\tau,\mathbf{c})|\}\, d\mathbf{c}, \tag{59}$$

where we use the same normalization constant Y, as chosen for Eq. (58), to enable comparison of the two functions, $\left\langle g_b^{\text{gen}}(t) \right\rangle$ and $\left\langle \tilde{g}_b^{\text{gen}}(\tau) \right\rangle$. The number of particles generated in a frequency-domain simulation using Eq. (41) is proportional to $\left\langle \tilde{g}_b^{\text{gen}}(\tau) \right\rangle$. Depending on the phase of $\tilde{g}^{\text{gen}}$, the value of $\left\langle \tilde{g}_b^{\text{gen}}(\tau) \right\rangle$ may vary from $1/\sqrt{2}$, when the magnitude of the real and imaginary parts of $\tilde{g}^{\text{gen}}$ is equal, to 1 when the real part of $\tilde{g}^{\text{gen}}$ attains its maximum amplitude and the imaginary part is zero (or vice versa). We label these two states "mixed-phase" and "in-phase," respectively; see Fig. 14 (left panel). Therefore, for the in-phase case, the boundary will generate $1/(2/\pi) \approx 1.57$ times more particles in a frequency-domain simulation than the average in a time-domain simulation. For mixed-phase, this ratio will be $(1/\sqrt{2})/(2/\pi) \approx 1.11$.

The same bounds hold for all other particle generating processes throughout the flow domain. The overall maximum N_{ratio} (see definition above) value is therefore obtained if all generating functions are in the in-phase state. The lowest N_{ratio} value occurs if all generating functions are in

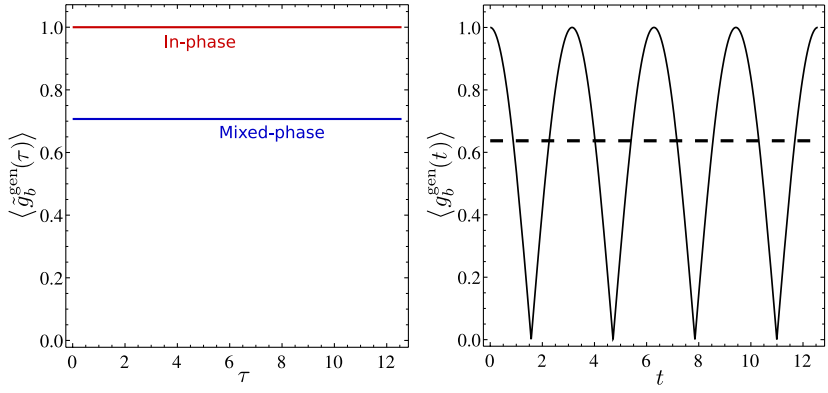

FIG. 14. Relative magnitude of particles generated at boundary. Left: frequency-domain, in-phase case shown in blue and mixed-phase case shown in red. Right: time-domain, solid line indicates time dependent value and dashed line is time averaged value. The ratio of the in-phase case to the time average is 1.57, and the ratio of the mixed-phase case to the time average is 1.11.

TABLE I. Ratio of the number of particles in the frequency-domain to the average number in the time-domain, denoted N_{ratio}, as a function of the boundary velocity phase, φ. BGK simulation conditions in Fig. 5 are used, with a range of boundary phases: $\tilde{\mathbf{u}}_b = (U[1+0\text{i}], 0, 0)$, $\tilde{\mathbf{u}}_b = (U[0.924+0.383\text{i}], 0, 0)$, and $\tilde{\mathbf{u}}_b = (U[0.707+0.707\text{i}], 0, 0)$. Phase shift $\varphi = 0$ represents the "in-phase" condition discussed above. A phase shift of $\varphi = \pi/4$ represents the "mixed-phase" condition. Note that HS simulations produce a similar outcome. The value of $N_{\text{ratio}} = 1.58$ attained for $\theta = 100$ approximately agrees with the predicted maximum value of 1.57.

		N_{ratio}		
Kn	θ	$\varphi = 0$	$\varphi = \pi/8$	$\varphi = \pi/4$
0.1	0.1	1.40	1.39	1.42
0.1	1	1.44	1.30	1.44
1	100	1.58	1.48	1.17
10	0.1	1.52	1.43	1.27

the mixed-phase state. In practice, the phase varies throughout the flow field, and most generating functions will yield a ratio of particles between the frequency and time-domain methods that lie somewhere between these two extremes. The overall ratio of particles in the frequency and time-domains, N_{ratio}, will therefore fall between 1.11 and 1.57.

Table I gives values for N_{ratio} corresponding to the simulations shown in Fig. 5. A range of cases is given where the phase of the boundary velocity, φ, is varied; whilst this phase can be set arbitrarily (because there is no external reference), its specification affects the number of particles used; see Eq. (59). When $\varphi = 0$, the boundary velocity is real and the maximum number of boundary particles is generated; this boundary condition is used in Section III and corresponds to the "in-phase" state discussed above. For $\varphi = \pi/4$, the real and imaginary parts of the boundary velocity have equal magnitude, corresponding to the "mixed-phase" condition which minimizes the number of particles generated at the boundary.

For high frequency flows, such as $\theta = 100$, the magnitude of $\tilde{g}$ decreases rapidly away from the wall, and the majority of the simulator particles are observed to be in-phase with the boundary. In this case, changing the boundary phase dramatically affects the number of particles in the simulation, as seen in row 3 of Table I. For high Knudsen number flows, e.g., Kn = 10, a relatively low spatial variation in the phase of $\tilde{g}$ occurs. Once again, changing the phase of the boundary strongly impacts the total number of particles. This is illustrated in row 4 of Table I. Away from each of these extreme limits, the choice of boundary phase exerts a smaller effect on the overall number of particles; see rows 1 and 2 of Table I.

For the BGK model, computational expense of the algorithm scales linearly with the number of particles. The relative computational expense per step of the time and frequency-domain methods is therefore specified by the relative number of particles, N_{ratio}. For HS simulations, most of the algorithm again scales linearly with the number of particles, with the exception of particle generation during the collision operation in Eqs. (38) and (39); these terms exhibit $O(N_c^2)$ scaling behaviour, i.e., quadratic dependence on the number of particles per cell. However, the relative effect of this nonlinear dependence is minor in all of the examples considered above, i.e., the relative time taken to complete a step in each domain is found to be proportional to the relative number of particles to a very good approximation.

C. Relative computational performance

Here, we combine the findings regarding the relative statistical convergence and the relative computational expense per step, to determine the *relative computational performance* of the time and frequency-domain methods. As the relative statistical convergences of the two methods have

been theoretically calculated in the small and large βM limits, $\beta M \ll 1$ and $\beta M \gg 1$ (see above), we now describe the overall computational performance in terms of these two limits.

1. Small βM

Based on the minimum and maximum values of N_{ratio} given in Section IV B, a single step in the frequency-domain method takes 1.11 to 1.57 times longer to complete than the average step in the time-domain method. However, as discussed in Section IV A and Appendix B, the frequency-domain method exhibits dramatically superior statistical convergence when $\beta M \ll 1$; it far exceeds the above maximum factor of 1.57. The frequency-domain method is therefore much more efficient than the time-domain method in this parameter region.

2. Large βM

For large βM, the lowest variance achievable in a single fit parameter using the time-domain method is 1.5 times greater than that of the corresponding parameter using the frequency-domain method; see Eq. (56). However, as noted in the discussion pertaining to Eq. (56), there is a reciprocal relationship between the variances of the two fit parameters in the time-domain method: decreasing the variance of the in-phase parameter increases the variance of the out-of-phase parameter, and vice versa.

If the frequency-domain simulations contain the maximum relative number of particles, $N_{\text{ratio}} = 1.57$, and the in-phase fit parameter in the time-domain simulations displays the minimum relative variance of 1.5, the two effects largely cancel, i.e., both methods have similar computational performance *in determining the in-phase and real fit parameters.* However, due to the reciprocal relationship described above, the frequency-domain method displays superior computational performance when comparing the out-of-phase and imaginary parameters. That is, the frequency-domain method displays better computational performance in both fit parameters collectively, even when the most favorable conditions are selected for the time-domain method.

3. Arbitrary βM

The above discussion shows that the frequency-domain method exhibits superior computational efficiency relative to the time-domain method, in the asymptotic limits of small and large βM. Since operation at intermediate βM exhibits features of both these limits, an identical conclusion is expected for all βM and is borne out in simulations (data not shown).

D. Time for flow to be fully developed

In the preceding comparison of the computational performance of the time and frequency-domain methods, the flow is fully developed, i.e., all transients in the simulation have decayed. The time taken for the flow to fully develop increases with decreasing Kn and θ and for larger domain size.

A fully developed flow can always be obtained by starting the simulation with less computationally intensive parameters, such as larger values for N_p, Δt, or $\Delta\tau$. These values are systematically reduced over the course of flow development, thus refining the solution. This approach can be used to ensure that the flow development time decreases to a negligible level relative to the overall computational time.

V. OSCILLATORY FLOWS OF ARBITRARY TIME-DEPENDENCE

In Section III, time and frequency-domain deviational methods were used to simulate flows generated by boundary conditions with a sinusoidal time-dependence; these methods can also be applied to flows driven by forces that vary sinusoidally. In Section IV, it was shown that, due

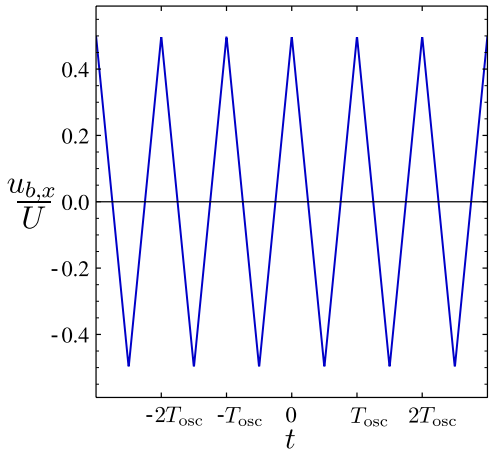

FIG. 15. Triangular time-dependent velocity boundary condition used in non-sinusoidal oscillatory Couette flow.

to superior statistical convergence, frequency-domain simulations provide improved computational efficiency. Here, we describe the extension of the deviational frequency-domain method for the efficient simulation of non-sinusoidally varying flows. While we focus on the deviational method, a similar approach may be used for the weight-based method of Ref. 34 (results not shown).

To apply the deviational frequency-domain method to such flows, we first express the time-domain boundary conditions (or driving forces) of the flow in their Fourier series. Each sinusoidal term in these Fourier series is then simulated individually using the frequency-domain method and combined to obtain the required result. This process naturally affords computational parallelization, which when combined with the superior computation efficiency of the frequency-domain method, and enables highly efficient simulations of oscillatory flows with arbitrary time-dependencies.

We illustrate this approach by way of example, using an unsteady Couette flow whose velocity boundary condition at its upper solid wall is a "triangular wave" of period $T_{\rm osc}$,

$$\mathbf{u}_b(t) = \left(U\left(\frac{t}{T_{\rm osc}} + \frac{1}{2} - \left\lfloor \frac{t}{T_{\rm osc}} + 1 \right\rfloor \right) (-1)^{\lfloor t/T_{\rm osc}+1 \rfloor}, 0, 0 \right), \tag{60}$$

where $\lfloor\ \rfloor$ indicates the floor (greatest integer) function and U is the velocity amplitude, see Fig. 15; as before, the lower wall is stationary. This boundary condition may be expressed in its Fourier series of the x-component,

$$u_{b,x}(t) = \sum_{j=1}^{\infty} \frac{4U}{\pi^2(2j-1)^2} \cos(2[2j-1]\omega_{\rm osc} t), \tag{61}$$

where $\omega_{\rm osc} = 2\pi/T_{\rm osc}$, and the subscript x denotes the x-component of the velocity. The time-domain boundary condition in Eq. (60) may then be approximated by a finite number of individual frequency-domain simulations, corresponding to the lowest order terms in Eq. (61) with velocity boundary conditions,

$$u_x^{(j)} = \frac{4U}{\pi^2(2j-1)^2} \tag{62}$$

and angular frequencies

$$\omega_j = 2j\omega_{\rm osc}. \tag{63}$$

The frequency components in Eq. (62) are independent and can therefore be treated separately. However, dramatic gains in computational efficiency are obtained by exploiting any "overlap" in the distribution functions, $\tilde{g}_j$, i.e., where a candidate particle is enclosed by more than one $\tilde{g}_j$, where each j corresponds to a single (and different) frequency component.

As noted previously, processes such as collisions/relaxation and boundary interactions involve the creation of simulator particles through the sampling of "generating distributions." This procedure is defined in Eqs. (30)-(36) and illustrated in Fig. 1. We may extend this approach to the overlapping parts of the generating functions for different j, representing different frequencies of the Fourier series in Eq. (61). This procedure is encapsulated in the following mathematical description.

For J individual simulations, representing different frequencies of the Fourier series, we rewrite Eq. (30) as

$$p(\mathbf{c}) > \max\{|g^{\text{gen}}_{1,\text{re}}(\tau,\mathbf{c})|, |g^{\text{gen}}_{1,\text{im}}(\tau,\mathbf{c})|, \ldots, |g^{\text{gen}}_{J,\text{re}}(\tau,\mathbf{c})|, |g^{\text{gen}}_{J,\text{im}}(\tau,\mathbf{c})|\}. \tag{64}$$

As such, a particle is accepted when

$$\mathcal{R}p(\mathbf{c}) < \max\{|g^{\text{gen}}_{1,\text{re}}(\tau,\mathbf{c})|, |g^{\text{gen}}_{1,\text{im}}(\tau,\mathbf{c})|, \ldots, |g^{\text{gen}}_{J,\text{re}}(\tau,\mathbf{c})|, |g^{\text{gen}}_{J,\text{im}}(\tau,\mathbf{c})|\}. \tag{65}$$

Each simulator particle now carries a value $\tilde{\mathcal{W}}_j$ for each of the J simulations. For the jth simulation, in the case

$$\mathcal{R}p(\mathbf{c}) < |g^{\text{gen}}_{j,\text{re}}(\tau,\mathbf{c})|, \tag{66}$$

the real part of the jth value is set to

$$\tilde{\mathcal{W}}_{j,\text{re}} = \text{sgn}[g^{\text{gen}}_{j,\text{re}}(\tau,\mathbf{c})]. \tag{67}$$

Correspondingly, when

$$\mathcal{R}p(\mathbf{c}) < |g^{\text{gen}}_{j,\text{im}}(\tau,\mathbf{c})|, \tag{68}$$

the imaginary part of the jth value is set to

$$\tilde{\mathcal{W}}_{j,\text{im}} = \text{sgn}[g^{\text{gen}}_{j,\text{im}}(\tau,\mathbf{c})]. \tag{69}$$

This particle generation procedure is illustrated in Fig. 16.

Calculation of the moments (or any other integration) of each $\tilde{g}_j(\tau,\mathbf{c})$ is performed separately, using the corresponding particle values. For example, Eq. (28) becomes

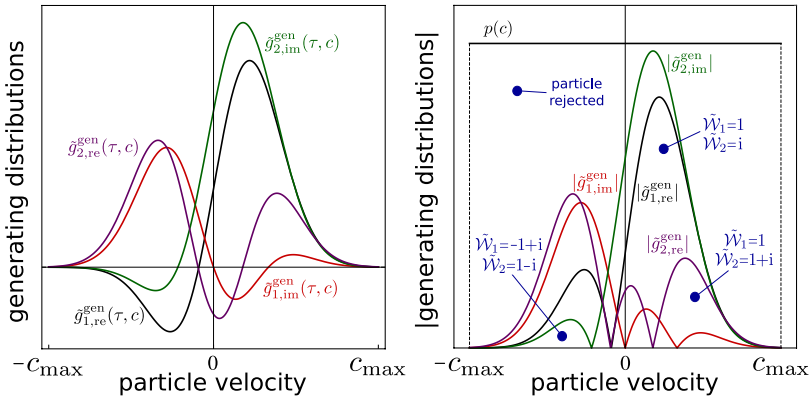

FIG. 16. Particle generation examples using more than one frequency component. Left panel: two example generating distributions, $\tilde{g}^{\text{gen}}_1$ and $\tilde{g}^{\text{gen}}_2$. Right panel: extension of the particle generation procedure given in Section II B to the case of multiple distributions. Overlapping parts of separate distributions are treated in the same manner as the real and imaginary parts of a single distribution; see Fig. 1. In the lower left example, the randomly selected point lies beneath all of $|\tilde{g}^{\text{gen}}_{1,\text{re}}|$, $|\tilde{g}^{\text{gen}}_{1,\text{im}}|$, $|\tilde{g}^{\text{gen}}_{2,\text{re}}|$, and $|\tilde{g}^{\text{gen}}_{2,\text{im}}|$, and the particle is therefore assigned real and imaginary values for both distributions; namely, because $\tilde{g}^{\text{gen}}_{1,\text{re}} < 0$ and $\tilde{g}^{\text{gen}}_{1,\text{im}} > 0$, we assign $\tilde{\mathcal{W}}_1 = -1+\text{i}$, and as $\tilde{g}^{\text{gen}}_{2,\text{re}} > 0$ and $\tilde{g}^{\text{gen}}_{2,\text{im}} < 0$, we assign $\tilde{\mathcal{W}}_2 = 1-\text{i}$. Other particles are assigned in the same manner, as indicated.

$$\begin{aligned}\tilde{\sigma}_j &= \frac{N_p}{n_0 V_c}\sum_i^{N_c}\tilde{\mathcal{W}}_{j,i},\\ \tilde{\mathbf{u}}_j &= \frac{N_p}{n_0 V_c}\sum_i^{N_c}\tilde{\mathcal{W}}_{j,i}\mathbf{c}_i,\\ \tilde{\kappa}_j &= \frac{N_p}{n_0 V_c}\frac{2}{3}\sum_i^{N_c}\tilde{\mathcal{W}}_{j,i}\left(\frac{c_i^2}{2RT_0}-\frac{3}{2}\right),\end{aligned} \tag{70}$$

where again the index i refers to the ith particle. Note that the factors multiplying $\tilde{\mathcal{W}}_{j,i}$ in Eq. (70) are independent of the frequency index j and are not recalculated for each simulation. The particle oscillation operation is performed separately on each set of cell moments and particle values, for example, Eq. (26) becomes

$$\begin{bmatrix}\tilde{\mathcal{W}}_{j,\mathrm{re,new}}\\ \tilde{\mathcal{W}}_{j,\mathrm{im,new}}\end{bmatrix} = \begin{bmatrix}\cos(\Delta\tau\omega_j) & -\sin(\Delta\tau\omega_j)\\ \sin(\Delta\tau\omega_j) & \cos(\Delta\tau\omega_j)\end{bmatrix}\begin{bmatrix}\tilde{\mathcal{W}}_{j,\mathrm{re}}\\ \tilde{\mathcal{W}}_{j,\mathrm{im}}\end{bmatrix}. \tag{71}$$

The rest of the frequency-domain algorithm, i.e., particle deletion and advection, is unchanged. Because these latter operations are independent of the frequency index j, they apply to multiple frequency components simultaneously. For example, if a simulator particle carries non-zero values for $\tilde{\mathcal{W}}_1$ and $\tilde{\mathcal{W}}_2$, i.e., it represents both $\tilde{g}_1$ and $\tilde{g}_2$ (for two frequency components), the advection of this particle applies simultaneously to both frequency components 1 and 2. This results in a considerable improvement in computational efficiency compared to advecting two particles in separate simulations of each frequency component.

Increasing the number of individual (and concurrent) simulations, J, results in an increase in the total computational cost. However, the additional cost for each simulation (at each frequency

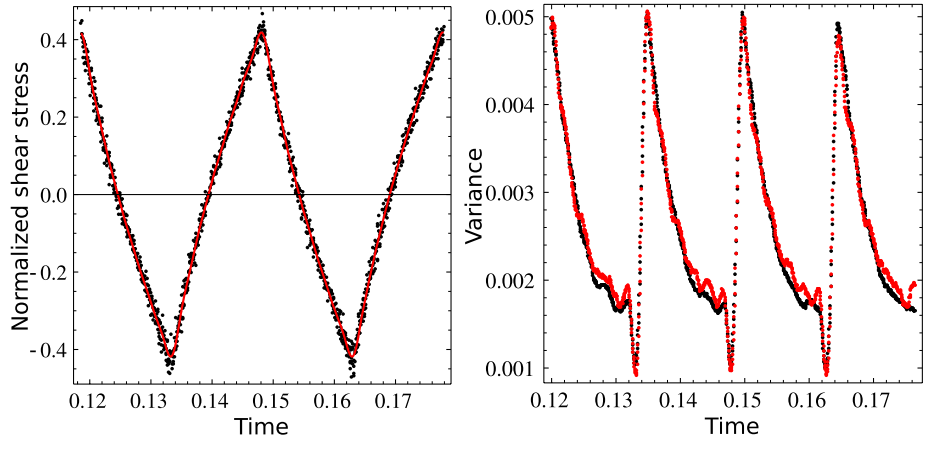

FIG. 17. [BGK model, triangular time-dependent Couette flow]: shear stress at the oscillating wall, comparison of time, and frequency-domain deviational results when simulating a non-sinusoidal "triangle" flow. Left: the black dots are the instantaneous output in the direct time-domain simulation of the triangular boundary condition. The red line is the six term Fourier series constructed from the parallelized frequency-domain method. For each frequency-component, a fit is made to the frequency-domain simulation output to determine the coefficient in the Fourier series. Right: black curve is the variance of the direct time-domain simulation (black dots in left panel), i.e., the variance of the instantaneous output, measured using a moving window of 20 points and averaged over multiple cycles. Red curve: this variance is synthetically constructed from the frequency-domain simulations (it is used for comparison with the black curve). It is constructed as follows: (1) the variance of the amplitude of each frequency component of the parallelized frequency-domain simulation is measured. (2) Noise is then added to the Fourier reconstructed time-series in the left panel (red), using a random number generator of variance equal to that of point 1. (3) A moving window of 20 points is used on the noisy Fourier reconstruction from point 2 and averaged over multiple cycles, as was performed for the direct time-domain simulation. This variance (red) can then be compared directly to the variance in the time-domain simulation (black). The small mismatch between the red and black curves is due to the truncation of the Fourier series to 6 terms. Simulation parameters: $\mathrm{Kn}=1$, $\theta_{\mathrm{osc}}=0.1$, $\mathrm{Kn}/\Delta y^*=20$, $1/\Delta\tau^*, 1/\Delta t^*=64$, and $N_p=40\times10^9$.

component) rapidly diminishes due to (i) the increasing degree of "overlap" (see above) between the set of $\tilde{g}_j$ as more simulations are added, and (ii) the decreasing magnitude of the Fourier components with increasing j, which reduces the number of simulator particles required.

Results using the above-described algorithm for a triangular wave Couette flow simulation are given in Fig. 17. In the left panel, direct time-domain deviational method simulations of the flow are compared to the result of simulating the first six Fourier terms in Eq. (61) using the parallelized frequency-domain method (summed in Fourier space to give a time-domain response). In the right panel, the stochastic noise in the time-domain solution is compared to the combined noise of each Fourier component from the frequency-domain simulation; see the figure caption for further details. The very similar results between the two methods indicate that the stochastic noise in a transient time-domain simulation is given by the sum of the noise in the Fourier components. The results of Section IV are therefore applicable to the individual Fourier components of a general transient flow, regardless of whether the frequency-domain method or time-domain method are used.

For this example flow, the parallelized frequency-domain method displays superior statistical convergence compared to the time-domain simulation; the predicted relative convergence for the first two frequency components is shown in Fig. 18. A single virtual-time step of the combined frequency-domain simulations, i.e., one step in all six frequency components (used in this example), takes approximately 1.5 times more computational time than a step in the time-domain simulation. The similarity of this figure to the predicted maximum of 1.57, given in Section IV B, is coincidental — the bounds calculated in Section IV B apply to simulations of individual frequency components, but not to the combined parallelized simulation. This longer computational time per step in the combined frequency-domain simulations is greatly outweighed by the superior statistical convergence of the

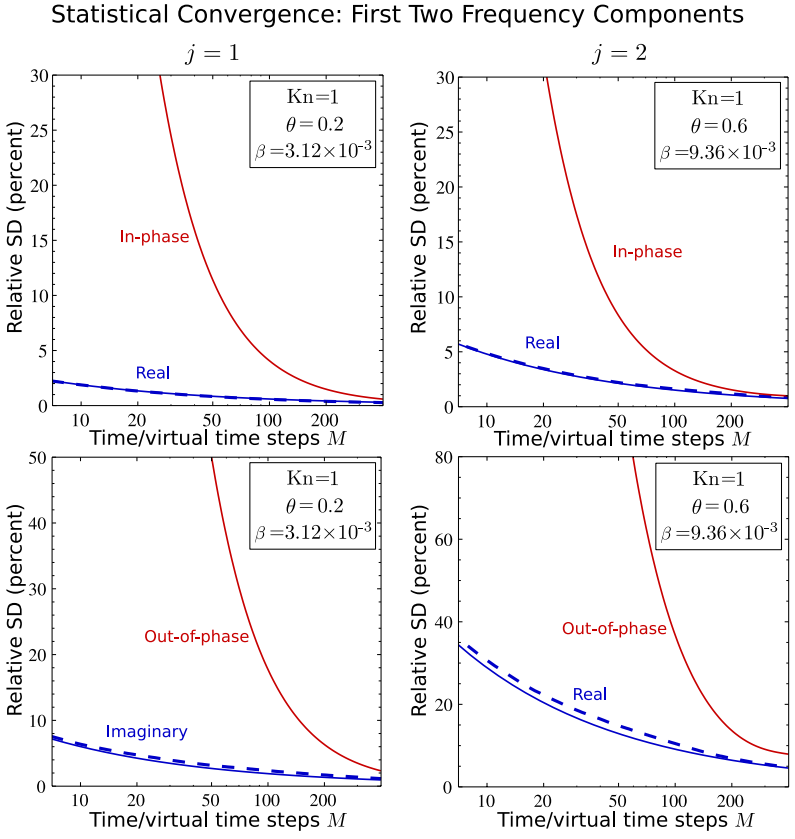

FIG. 18. [BGK model, triangular time-dependent Couette flow]: percentage relative standard deviation in the shear stress at the oscillating wall for the first two frequency components. Left column: first component, $j = 1$. Right column: second component, $j = 2$. Blue dashed lines: frequency-domain simulations. Solid lines: analytical predictions due to equations given in Ref. 34.

frequency-domain method. As such, the parallelized frequency-domain simulation displays much better computational efficiency than the time-domain simulation.

This general parallelized approach thus enables trivial and efficient extension of the frequency-domain method to boundary conditions of periodic and arbitrary time-dependence.

VI. CONCLUSIONS

We have proposed a deviational frequency-domain Monte Carlo method to simulate oscillatory, low Mach number gas flows. This is based on the linearized Boltzmann equation and easily accommodates a wider range of collision operators than previous work using a weight-based method;[34] both BGK relaxation and HS collision operators were considered in this article. This deviational frequency-domain method is formulated by extending current deviational time-domain Monte Carlo methods[17,20] to the frequency-domain, using the underlying complex-variable solution of Ref. 34. The accuracy of this deviational frequency-domain method was verified using two model flows: oscillatory Couette flow and solid sphere oscillating in an unbounded gas. Comparison to existing high accuracy numerical solutions of the Boltzmann-BGK equation for oscillatory Colette flow, and to conventional time-domain simulations for both flows, was performed. Excellent agreement was achieved in all cases.

We also compared predictions of the BGK and HS models for both isothermal (Couette) and non-isothermal (solid sphere) oscillatory flows. Strikingly, the BGK model was found to produce highly accurate results relative to the more realistic HS model in both cases, despite the Prandtl numbers differing between the two kinetic models.

The relative performance of the deviational time and frequency-domain methods was analyzed by considering (i) the number of simulator particles required and (ii) the statistical convergence properties of each method. The frequency-domain method requires a slightly larger number of particles (maximum of approximately 1.57×) than the time-domain method, leading to a greater computational expense on a per step basis. However, due to superior statistical convergence, the frequency-domain method requires fewer steps, leading to a significant improvement in overall computational efficiency.

ACKNOWLEDGMENTS

The authors acknowledge financial support from an Australian Postgraduate Award and the Australian Research Council Grants Scheme.

APPENDIX A: HARD SPHERE COLLISION OPERATOR

The collision kernel, K, and the relaxation frequency, $\hat{\nu}$, are derived in Ref. 36, based on a form of the HS collision operator given by Hilbert.[43] The collision kernel is

$$K(\mathbf{c},\mathbf{c}_1) = K_1(\mathbf{c},\mathbf{c}_1) - K_2(\mathbf{c},\mathbf{c}_1), \tag{A1}$$

where

$$K_1(\mathbf{c},\mathbf{c}_1) = \frac{1}{\lambda\sqrt{\pi^3 RT}}\frac{1}{|\mathbf{c}-\mathbf{c}_1|}\exp\left(-\frac{[\mathbf{c}\cdot(\mathbf{c}-\mathbf{c}_1)]^2}{2RT\,|\mathbf{c}-\mathbf{c}_1|^2}\right), \tag{A2a}$$

$$K_2(\mathbf{c},\mathbf{c}_1) = \frac{1}{4\lambda(\pi RT)^{3/2}}|\mathbf{c}-\mathbf{c}_1|\exp\left(-\frac{\mathbf{c}^2}{2RT}\right). \tag{A2b}$$

Note that K_1 has a singularity at $\mathbf{c} = \mathbf{c}_1$. To avoid difficulties associated with generating particles from an unbounded distribution, the mean value is used in the region $|\mathbf{c}-\mathbf{c}_1| < c_m$, given by

$$\bar{K}_1(\mathbf{c},\mathbf{c}_1) = \frac{3}{2\pi\lambda\sqrt{RT}}\,\mathrm{erf}\left(\frac{|\mathbf{c}_1|}{\sqrt{2RT}}\right)\frac{1}{c_m|\mathbf{c}_1|}, \tag{A3}$$

where the error function is defined as

$$\mathrm{erf}(\chi) = \frac{2}{\sqrt{\pi}}\int_0^{\chi}\exp\left(-\upsilon^2\right)d\upsilon.$$

The collision frequency term is

$$\hat{\nu}(\mathbf{c}) = \sqrt{\frac{RT}{\pi}}\frac{1}{2\lambda}\left\{\sqrt{\pi}\,\mathrm{erf}\left(\frac{|\mathbf{c}|}{\sqrt{2RT}}\right)\left(\frac{\sqrt{2RT}}{|\mathbf{c}|}+\frac{2|\mathbf{c}|}{\sqrt{2RT}}\right)+2\exp\left(-\frac{c^2}{2RT}\right)\right\}. \tag{A4}$$

It follows from conservation of mass, momentum, and energy that

$$\begin{aligned}
\int_{-\infty}^{\infty}\int_{-\infty}^{\infty} K(\mathbf{c},\mathbf{c}_1)\tilde{g}(\mathbf{c}_1)d\mathbf{c}_1 d\mathbf{c} &= \int_{-\infty}^{\infty}\hat{\nu}(\mathbf{c})\tilde{g}(\mathbf{c})d\mathbf{c},\\
\int_{-\infty}^{\infty}\int_{-\infty}^{\infty} \mathbf{c}K(\mathbf{c},\mathbf{c}_1)\tilde{g}(\mathbf{c}_1)d\mathbf{c}_1 d\mathbf{c} &= \int_{-\infty}^{\infty}\mathbf{c}\hat{\nu}(\mathbf{c})\tilde{g}(\mathbf{c})d\mathbf{c},\\
\int_{-\infty}^{\infty}\int_{-\infty}^{\infty} c^2K(\mathbf{c},\mathbf{c}_1)\tilde{g}(\mathbf{c}_1)d\mathbf{c}_1 d\mathbf{c} &= \int_{-\infty}^{\infty}c^2\hat{\nu}(\mathbf{c})\tilde{g}(\mathbf{c})d\mathbf{c}.
\end{aligned} \tag{A5}$$

These moments are employed when calculating the mass, bulk velocity, and temperature perturbations used in Eqs. (46)-(48).

APPENDIX B: ADDITIONAL ANALYSIS OF STATISTICAL CONVERGENCE

In this appendix, we discuss the approach used to generate the variance (standard deviation) data given in Section IV A and provide some additional examples.

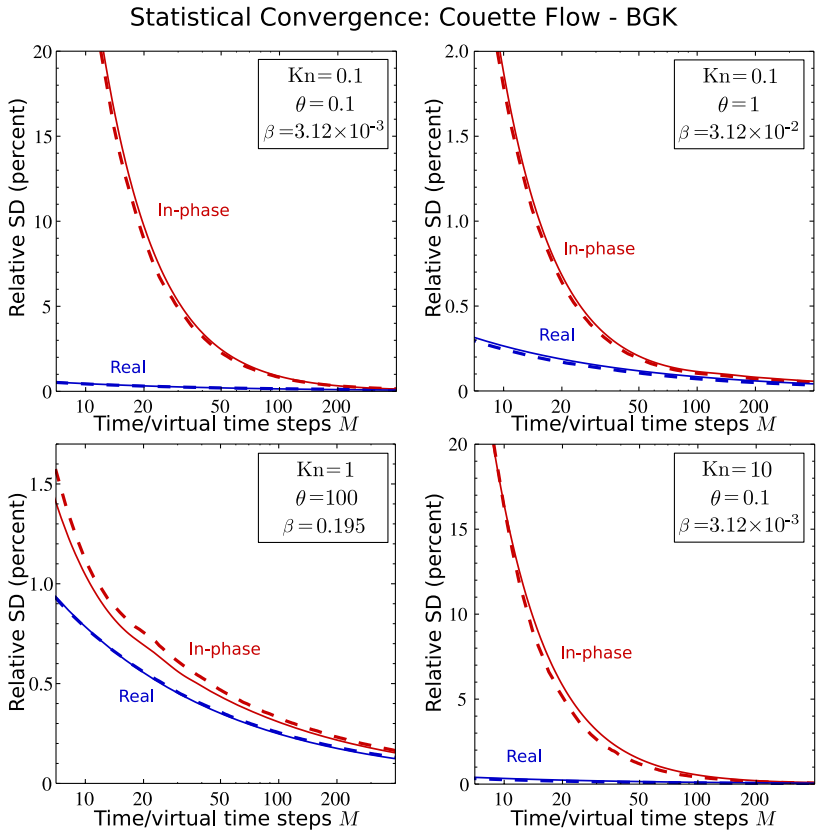

FIG. 19. [BGK model, oscillatory Couette flow]: percentage relative standard deviation in the shear stress at the oscillating wall. Red dashed lines: time-domain simulations, in-phase component. Blue dashed lines: frequency-domain simulations, real component. Solid lines: analytical predictions of equations in Ref. 34. Ratio of the standard deviations of the time and frequency-domain method results is theoretically predicted to be greater than 1.22 for large steps. After 400 steps, data in the top right and bottom left panels ($\beta M \gg 1$) give ratios of 1.49 and 1.28, respectively. The top left and bottom right panels ($\beta M = 1.25$) give 1.64 and 1.76, respectively.

For Couette flow, shear stress at the oscillating wall is determined, whilst for the oscillating sphere, we solve for the drag. Using the same approach as Ref. 34, at least 8000 simulations are used in both the time and frequency-domain simulations. For each simulation, fits of $M = 1$ to 400 steps are made, starting at a random time/virtual-time after the flow has fully developed. For each number of steps, the standard deviation in the fit parameters is then calculated. The time and frequency-domain results are compared to predictions of the equations in Ref. 34.

For the frequency-domain method, choice of the starting virtual-time (after the flow has reached steady-state) has no effect on the rate of statistical convergence. However, this is not the case for the time-domain method, where the location of the starting point within an oscillation period influences the variance of the fit. The time-domain method simulation data shown here and in Section IV A therefore give the rate of statistical convergence *averaged* over all choices of starting time. For further discussion of this point, see Ref. 34. Figure 19 gives the standard deviation in the Couette flow BGK simulations, using the parameter range in Fig. 5, whilst Fig. 20 shows the corresponding HS results. Figures 13 and 21 give results for the oscillating sphere, using the parameters from Figs. 9 and 10, respectively.

It is important to emphasize that the equations in Ref. 34 are derived under the ideal assumption of statistical independence between steps; however, in practice, there is some (unavoidable) small degree of correlation. This accounts for the small discrepancies between the predicted and observed rates of convergence.

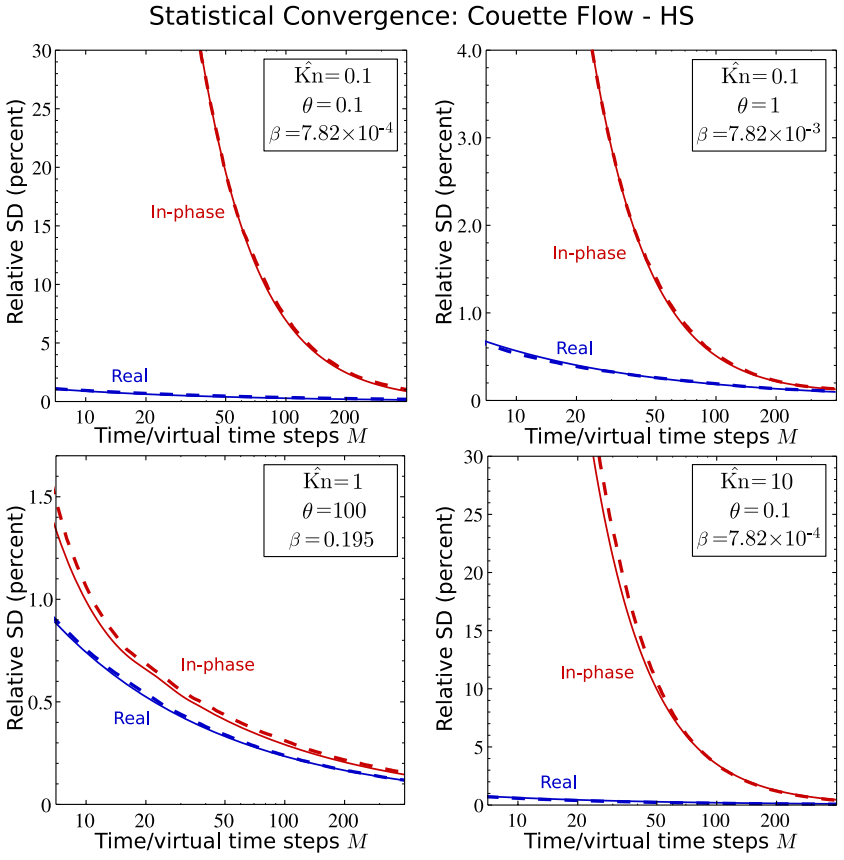

FIG. 20. [HS model, oscillatory Couette flow]: percentage relative standard deviation in the shear stress at the oscillating wall. Red dashed lines: time-domain simulations, in-phase component. Blue dashed lines: frequency-domain simulations, real component. Solid lines: analytical predictions of equations in Ref. 34. Ratio of the standard deviations of the time and frequency-domain method results is theoretically predicted to be greater than 1.22 for large steps. After 400 steps, data in the top right and bottom left panels ($\beta M \gg 1$) give ratios of 1.30 and 1.32, respectively. The top left and bottom right panels ($\beta M \ll 1$) give 5.35 and 5.57, respectively.

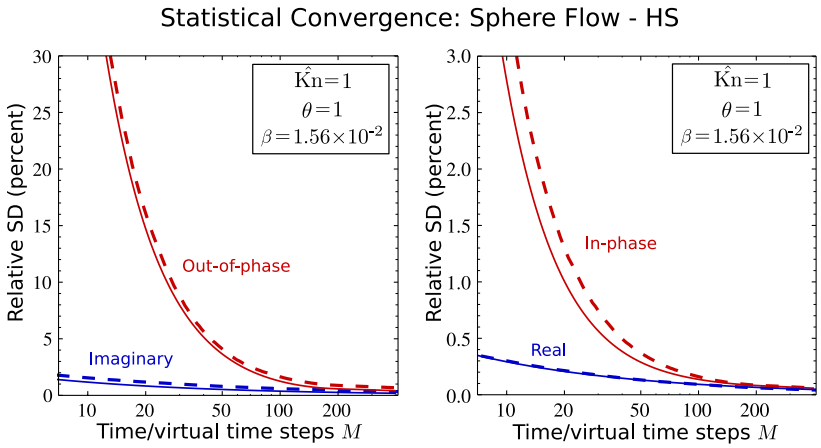

FIG. 21. [HS model, oscillating sphere]: percentage relative standard deviation in the drag. Red dashed lines: time-domain simulations. Blue dashed lines: frequency-domain simulations. Solid lines: analytical predictions due to equations given in Ref. 34. Left: imaginary and out-of-phase components. Right: real and in-phase components. Ratio of the imaginary and out-of-phase standard deviations after 400 steps is 2.73, whilst the ratio of the real and in-phase standard deviations is 1.22. Each of these values agrees with the theoretically predicted minimum of 1.22 when $\beta M \gg 1$.

[1] Y. T. Yang, C. Callegari, X. L. Feng, K. L. Ekinci, and M. L. Roukes, "Zeptogram-scale nanomechanical mass sensing," Nano Lett. **6**(4), 583–586 (2005).
[2] I. Bargatin, I. Kozinski, and M. L. Roukes, "Efficient electrothermal actuation of multiple modes of high-frequency nanoelectromechanical resonators," Appl. Phys. Lett. **90**(9), 093116 (2007).
[3] K. Jenses, K. Kim, and A. Zettl, "An atomic-resolution nanomechanical mass sensor," Nat. Nanotechnol. **3**, 533–537 (2008).
[4] G. Binnig, C. F. Quate, and C. Gerber, "Atomic force microscope," Phys. Rev. Lett. **56**(9), 930–933 (1986).
[5] F. J. Giessibl, "Advances in atomic force microscopy," Rev. Mod. Phys. **75**(3), 949–983 (1986).
[6] J. E. Sader, "Frequency response of cantilever beams immersed in viscous fluids with applications to the atomic force microscope," J. Appl. Phys. **84**(1), 64 (1998).
[7] S. Boskovic, J. W. M. Chon, P. Mulvaney, and J. E. Sader, "Rheological measurements using microcantilevers," J. Rheol. **86**(4), 891 (2002).
[8] M. Pelton, D. Chakraborty, E. Malachosky, P. Guyot-Sionnest, and J. E. Sader, "Viscoelastic flows in simple liquids generated by vibrating nanostructures," Phys. Rev. Lett. **111**, 244502 (2013).
[9] R. Berger, C. Gerber, H. P. Lang, and J. K. Gimzewski, "Micromechanics: A toolbox for femtoscale science: 'Towards a laboratory on a tip'," Microelectron. Eng. **35**, 373–379 (1997).
[10] N. V. Lavrik, M. J. Sepaniak, and P. G. Datskos, "Cantilever transducers as a platform for chemical and biological sensors," Rev. Sci. Instrum. **75**(7), 2229–2253 (2004).
[11] C. Cercignani, *Rarefied Gas Dynamics: From Basic Concepts to Actual Calculations* (Cambridge University Press, 2000), Vol. 21.
[12] Y. Sone, *Kinetic Theory and Fluid Dynamics* (Birkhäuser, 2000).
[13] J. Nassios and J. E. Sader, "Asymptotic analysis of the Boltzmann–BGK equation for oscillatory flows," J. Fluid Mech. **708**, 197–249 (2012).
[14] G. A. Bird, *Molecular Gas Dynamics and the Direct Simulation of Gas Flows*, Oxford Engineering Science Series (Clarendon Press, 1994).
[15] S. Asmussen and P. W. Glynn, *Stochastic Simulation: Algorithms and Analysis* (Springer, 2007), Vol. 57.
[16] G. A. Radtke, N. G. Hadjiconstantinou, and W. Wagner, "Low-noise Monte Carlo simulation of the variable hard sphere gas," Phys. Fluids **23**, 030606 (2011).
[17] 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**, 056711 (2009).
[18] G. A. Radtke, "Efficient simulation of molecular gas transport for micro- and nanoscale applications," Ph.D. thesis, Massachusetts Institute of Technology, 2011.
[19] N. G. Hadjiconstantinou, G. A. Radtke, and L. L. Baker, "On variance reduced simulations of the Boltzmann transport equation for small-scale heat transfer applications," J. Heat Transfer **132**, 112401 (2010).
[20] T. M. M. Homolle and N. G. Hadjiconstantinou, "Low-variance deviational simulation Monte Carlo," Phys. Fluids **19**, 041701 (2007).
[21] T. M. M. Homolle and N. G. Hadjiconstantinou, "A low-variance deviational simulation Monte Carlo for the Boltzmann equation," J. Comput. Phys. **226**, 2341–2358 (2007).
[22] C. D. Landon and N. G. Hadjiconstantinou, "Variance-reduced direct simulation Monte Carlo with the Bhatnagar-Gross-Krook collision operator," AIP Conf. Proc. **1333**, 277–282 (2010).
[23] S. Ramanathan and D. L. Koch, "An efficient direct simulation Monte Carlo method for low Mach number noncontinuum gas flows based on the Bhatnagar-Gross-Krook model," Phys. Fluids **21**, 033103 (2009).
[24] H. A. Al-Mohssen and N. G. Hadjiconstantinou, "Low-variance direct Monte Carlo simulations using importance weights," ESAIM: Math. Modell. Numer. Anal. **44**, 1069–1083 (2010).

[25] H. A. Al-Mohssen, "An excursion with the Boltzmann equation at low speeds: Variance-reduced DSMC," Ph.D. thesis, MIT, 2010.
[26] 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**, 107107 (2005).
[27] Y. Shi and J. E. Sader, "Lattice Boltzmann method for oscillatory Stokes flow with applications to micro- and nanodevices," Phys. Rev. E **81**, 036706 (2010).
[28] G. R. McNamara and G. Zanetti, "Use of the Boltzmann equation to simulate lattice-gas automata," Phys. Rev. Lett. **61**, 2332–2335 (1988).
[29] S. Chen, H. Chen, D. Martnez, and W. Matthaeus, "Lattice Boltzmann model for simulation of magnetohydrodynamics," Phys. Rev. Lett. **67**, 3776 (1991).
[30] X. He and L.-S. Luo, "Theory of the lattice Boltzmann method: From the Boltzmann equation to the lattice Boltzmann equation," Phys. Rev. E **56**, 6811–6817 (1997).
[31] Y. Qian, D. d'Humières, and P. Lallemand, "Lattice BGK models for Navier-Stokes equation," EPL **17**, 479 (1992).
[32] H. Chen, S. Chen, and W. H. Matthaeus, "Recovery of the Navier-Stokes equations using a lattice-gas Boltzmann method," Phys. Rev. A **45**, R5339 (1992).
[33] Y. Shi, P. L. Brookes, Y. W. Yap, and J. E. Sader, "Accuracy of the lattice Boltzmann method for low-speed noncontinuum flows," Phys. Rev. E **83**, 045701 (2011).
[34] D. R. Ladiges and J. E. Sader, "Frequency-domain Monte Carlo method for linear oscillatory gas flows," J. Comput. Phys. **284**, 351–366 (2015).
[35] Y. W. Yap and J. E. Sader, "High accuracy numerical solutions of the Boltzmann Bhatnagar-Gross-Krook equation for steady and oscillatory Couette flows," Phys. Fluids **24**, 032004 (2012).
[36] T. M. M. Homolle, "Efficient particle methods for solving the Boltzmann equation," Master's thesis, Massachusetts Institute of Technology, 2007.
[37] J.-P. M. Péraud, C. D. Landon, and N. G. Hadjiconstantinou, "Monte Carlo methods for solving the Boltzmann transport equation," Annu. Rev. Heat Transfer **17**, 205–265 (2014).
[38] G. A. Radtke, J.-P. M. Péraud, and N. G. Hadjiconstantinou, "On efficient simulations of multiscale kinetic transport," Philos. Trans. R. Soc., A **371**, 20120182 (2013).
[39] J. Wakefield, A. Gelfand, and A. Smith, "Efficient generation of random variates via the ratio-of-uniforms method," Stat. Comput. **1**, 129–133 (1991).
[40] D. Rader, M. Gallis, J. Torczynski, and W. Wagner, "Direct simulation Monte Carlo convergence behavior of the hard-sphere-gas thermal conductivity for Fourier heat flow," Phys. Fluids **18**, 077102 (2006).
[41] N. G. Hadjiconstantinou, "The limits of Navier-Stokes theory and kinetic extensions for describing small-scale gaseous hydrodynamics," Phys. Fluids **18**, 111301 (2006).
[42] In the asymptotic work of Sone[12] and Nassios,[13] this value is represented by k.
[43] C. Cercignani, *The Boltzmann Equation and its Applications* (Springer, 1988).
[44] F. Sharipov and V. Seleznev, "Data on internal rarefied gas flows," J. Phys. Chem. Ref. Data **27**, 657–706 (1998).
[45] J. E. Sader, B. D. Hughes, J. A. Sanelli, and E. J. Bieske, "Effect of multiplicative noise on least-squares parameter estimation with applications to the atomic force microscope," Rev. Sci. Instrum. **83**, 055106 (2012).
[46] Note that there are additional co-dependencies between Var(γ), Var(ζ), Var(ξ), Var(μ), and a. See Ref. 34 for further details.