A HYBRID PARTICLE-CONTINUUM METHOD FOR HYDRODYNAMICS OF COMPLEX FLUIDS*

ALEKSANDAR DONEV[†§], JOHN B. BELL[†], ALEJANDRO L. GARCIA[‡], AND
BERNI J. ALDER[§]

Abstract. A previously developed hybrid particle-continuum method [J. B. Bell, A. Garcia, and S. A. Williams, *Multiscale Model. Simul.*, 6 (2008), pp. 1256–1280] is generalized to dense fluids and two- and three-dimensional flows. The scheme couples an explicit fluctuating compressible Navier–Stokes solver with the isotropic direct simulation Monte Carlo (DSMC) particle method [A. Donev, A. L. Garcia, and B. J. Alder, *J. Stat. Mech. Theory Exp.*, 2009 (2009), article P11008]. To achieve bidirectional dynamic coupling between the particle (microscale) and continuum (macroscale) regions, the continuum solver provides state-based boundary conditions to the particle subdomain, while the particle solver provides flux-based boundary conditions for the continuum subdomain. This type of coupling ensures both state and flux continuity across the particle-continuum interface analogous to coupling approaches for deterministic parabolic partial differential equations; here, when fluctuations are included, a small ($< 1\%$) mismatch is expected and observed in the mean density and temperature across the interface. By calculating the dynamic structure factor for both a "bulk" (periodic) and a finite system, it is verified that the hybrid algorithm accurately captures the propagation of spontaneous thermal fluctuations across the particle-continuum interface. The equilibrium diffusive (Brownian) motion of a large spherical bead suspended in a particle fluid is examined, demonstrating that the hybrid method correctly reproduces the velocity autocorrelation function of the bead but only if thermal fluctuations are included in the continuum solver. Finally, the hybrid is applied to the well-known adiabatic piston problem, and it is found that the hybrid correctly reproduces the slow nonequilibrium relaxation of the piston toward thermodynamic equilibrium but, again, only if the continuum solver includes stochastic (white-noise) flux terms. These examples clearly demonstrate the need to include fluctuations in continuum solvers employed in hybrid multiscale methods.

Key words. fluctuating hydrodynamics, hybrid methods, direct simulation Monte Carlo, adiabatic piston

AMS subject classifications. 65Z, 76N

DOI. 10.1137/090774501

1. Introduction. With the increased interest in nano- and microfluidics, it has become necessary to develop tools for hydrodynamic calculations at the atomistic scale [1, 2, 3]. While the Navier–Stokes–Fourier continuum equations have been surprisingly successful in modeling microscopic flows [4], there are several issues present in microscopic flows that are difficult to account for in models relying on a purely PDE approximation. For example, it is well known that the Navier–Stokes equations fail to describe flows in the kinetic regions (large Knudsen number flows) that appear in small-scale gas flows [5]. It is also not a priori obvious how to account for the bidirectional coupling between the flow and embedded microgeometry or complex boundaries. Furthermore, it is not trivial to include thermal fluctuations in Navier–Stokes solvers [6, 7, 8, 9], and, in fact, most of the time the fluctuations are omitted

†Center for Computational Science and Engineering, Lawrence Berkeley National Laboratory, Berkeley, CA 94720 (aleks.donev@gmail.com, JBBell@lbl.gov).
‡Department of Physics and Astronomy, San Jose State University, San Jose, CA 95192 (algarcia@algarcia.org).
§Lawrence Livermore National Laboratory, P.O. Box 808, Livermore, CA 94551-9900 (alder1@llnl.gov).

even though they can be important at instabilities [10] or in driving polymer dynamics [11, 12]. An alternative is to use particle-based methods, which are explicit and unconditionally stable, robust, and simple to implement. The fluid particles can be directly coupled to the microgeometry; for example, they can directly interact with the beads of a polymer chain. Fluctuations are naturally present and can be tuned to have the correct spatio-temporal correlations.

Several particle methods have been described in the literature, such as molecular dynamics (MD) [13], direct simulation Monte Carlo (DSMC) [14], dissipative particle dynamics (DPD) [15], and multiparticle collision dynamics (MPCD) [16, 17]. Here we use the isotropic DSMC (I-DSMC) stochastic particle method described in [18]. In the I-DSMC method, deterministic interactions between the particles are replaced with stochastic momentum exchange (collisions) between nearby particles. The I-DSMC method preserves the essential hydrodynamic properties of expensive MD: local momentum conservation, linear momentum exchange on length scales comparable to the particle size, and a similar fluctuation spectrum. At the same time, the I-DSMC fluid is ideal and structureless, and as such is much simpler to couple to a continuum solver.

However, even particle methods with coarse-grained dynamics, such as I-DSMC, lack the efficiency necessary to study realistic problems because of the very large numbers of particles needed to fill the required computational domain. Most of the computational effort in such a particle method would, however, be expended on particles far away from the region of interest, where a description based on the Navier–Stokes equations is likely to be adequate.
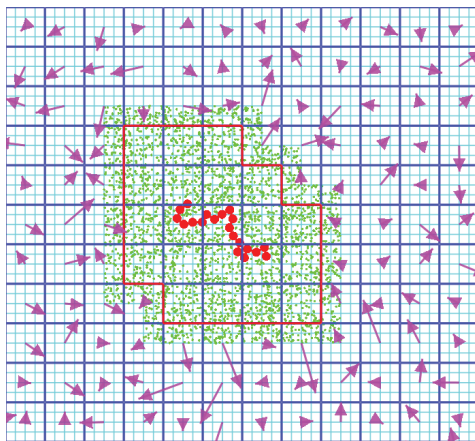

FIG. 1. *A two-dimensional illustration of the use of the MAR hybrid to study a polymer chain (larger red circles) suspended in an I-DSMC particle fluid. The region around the chain is filled with particles (smaller green circles), while the remainder is handled using a fluctuating hydrodynamic solver. The continuum (macro) solver grid is shown (thick blue lines), along with the (micro) grid used by the particle method (thin blue lines). The fluctuating velocities in the continuum region are shown as vectors originating from the center of the corresponding cell (purple). The interface between the particle and continuum regions is highlighted (thicker red line). (Color figures available in electronic version of article.)*

Hybrid methods are a natural candidate to combine the best features of the particle and continuum descriptions [19, 20]. A particle method can be used in regions where the continuum description fails or is difficult to implement, and a more efficient continuum description can be used around the particle domain, as illustrated in Figure 1. For example, a continuum description fails or is inaccurate in regions of

extreme gradients such as shocks [21], in rarefied regions [5], near singularities such as corners [22], and contact lines [23]. Unlike particle methods, continuum solvers require special and often difficult-to-implement techniques in order to handle flow near suspended structures or complex boundaries [12], or to handle sharp interfaces in fluid mixing [24, 25].

This type of hybrid algorithm fits in the multialgorithm refinement (MAR) simulation approach to the modeling and simulation of multiscale problems [26, 27]. In this approach the problem is represented at various levels of refinement, and, in addition to refining the spatio-temporal resolution at successive levels, potentially different algorithms are used at the different levels, as appropriate for the particular level of resolution. The general idea is to perform detailed calculations using an accurate but expensive (e.g., particle) algorithm in a small region, and couple this computation to a simpler, less expensive (e.g., continuum) method applied to the rest. The challenge is to ensure that the numerical coupling between the different levels of the algorithm hierarchy, and especially the coupling of the particle and continuum computations, is self-consistent, stable, and, most importantly, does not adversely impact the underlying physics.

Many hybrid methods coupling a particle subdomain with a continuum domain have been described in the literature; however, it is important to point out that the vast majority are meant for a different class of problems from our own. We will not attempt to review this large field but rather focus on the main differences between our method and those developed by other groups. The most important distinction to keep in mind is whether the system under consideration is a gas [26, 28, 29, 30], liquid [19, 20, 31], or solid [27]. The method we describe aims to bridge the region between gas and liquid flows, which we refer to as dense fluid flows. We are able to use a relatively simple coupling methodology because we neglect the structure of the particle fluid by virtue of using the I-DSMC particle method instead of the more commonly used (and expensive) MD. For a recent review of various methods for coupling MD with continuum methods for dense fluid flows, see [20, 31].

Another important classification is between static or *quasi-static coupling*, which is appropriate for finding steady-state solutions or time-dependent solutions when there is a large separation between the time scales of the macroscopic deformation and the microscopic dynamics, and *dynamic coupling*, which is capable of representing fast macroscopic processes such as the propagation of small-wavelength sound waves across a particle-continuum interface [32, 33]. Our method is one of a few coupling methods for fluid flows that belongs to the second category. A further distinction to keep in mind is based on the type of continuum model that is used (elliptic, parabolic, hyperbolic conservation law, or mixture thereof), and also what range of physical phenomena present in the particle dynamics is kept in the description. For example, the majority of methods for fluid flows use an incompressible approximation for the fluid equations, which is suitable only under an assumption of large separation of scales. Methods that do include density fluctuations often make an isothermal approximation and do not include heat transfer (i.e., the energy equation). The method we present here is dynamic and includes the complete compressible Navier–Stokes system in the continuum model, as was also done in [34].

Furthermore, there are different ways in which the continuum and particle subdomains can exchange information at their common interface. This important issue is discussed further in section 2.3.1, but for a more complete review the reader is referred to [20]. Broadly speaking, an important distinction is whether *state* or *flux* information is used to ensure continuity at the particle-continuum interface (for a historical

review, see [20]). Flux-based coupling between MD and a Navier–Stokes solver has been developed in a series of works by a group of collaborators starting with [35], further developed to include heat transfer [34] and also some state exchange to improve the stability for unsteady flows [34]. Our algorithm is unique in combining flux and state coupling to ensure continuity of both fluxes and state at the interface.

A crucial feature of our hybrid algorithm is that the continuum solver includes thermal fluctuations in the hydrodynamic equations consistent with the particle dynamics, as previously investigated in one dimension in [36] for DSMC and in [2, 37, 38] for MD. Thermal fluctuations are likely to play an important role in describing the state of the fluid at microscopic and mesoscopic scales, especially when investigating systems where the microscopic fluctuations drive a macroscopic phenomenon such as the evolution of instabilities, or where the thermal fluctuations drive the motion of suspended microscopic objects in complex fluids. Some examples in which spontaneous fluctuations may significantly affect the dynamics include the breakup of droplets in jets [39, 40, 41], Brownian molecular motors [42], Rayleigh–Bernard convection [43], Kolmogorov flows [44, 45], Rayleigh–Taylor mixing [10], and combustion and explosive detonation [46]. In our algorithm, the continuum solver is a recently developed three-stage Runge–Kutta integration scheme for the Landau–Lifshitz Navier–Stokes (LLNS) equations of fluctuating hydrodynamics in three dimensions [9], although other finite-volume explicit schemes can trivially be substituted.

As summarized in section 2, the proposed hybrid algorithm is based on a fully dynamic bidirectional state-flux coupling between the particle and continuum regions. In this coupling scheme the continuum method provides state-based boundary conditions to the particle subdomain through reservoir particles inserted at the boundary of the particle region at every particle time step. During each continuum time step a certain number of particle time steps are taken and the total particle flux through the particle-continuum interface is recorded. This flux is then imposed as a flux-based boundary condition on the continuum solver, ensuring strict conservation [26, 36]. Section 3 describes the technical details of the hybrid algorithm, focusing on components that are distinct from those described in [26, 36]. Notably, the use of the I-DSMC particle method instead of the traditional DSMC method requires accounting for the interactions among particles that are in different continuum cells.

In section 4 we thoroughly test the hybrid scheme in both equilibrium and non-equilibrium situations, and in both two and three dimensions. In section 4.1 we study the continuity of density and temperature across the particle-continuum interface and identify a small mismatch of order $1/N_0$, where N_0 is the number of particles per continuum cell, that can be attributed to the use of fluctuating values instead of means. In section 4.2 we compute dynamic structure factors in periodic ("bulk") and finite quasi two- and one-dimensional systems and find that the hybrid method seamlessly propagates thermal fluctuations across the particle-continuum interface.

In section 4.3 we study the diffusive motion of a large spherical buoyant bead suspended in a bead of I-DSMC particles in three dimensions by placing a mobile particle region around the suspended bead. This example is of fundamental importance in complex fluids and microfluidics, where the motion of suspended objects such as colloidal particles or polymer chains has to be simulated. Fluctuations play a critical role since they are responsible for the diffusive motion of the bead. The velocity-autocorrelation function (VACF) of a diffusing bead has a well-known power law tail of hydrodynamic origin, and its integral determines the diffusion coefficient. Therefore, computing the VACF is an excellent test for the ability of the hybrid method to capture the influence of hydrodynamics on the macroscopic properties of complex fluids.

Finally, in section 4.4 we study the slow relaxation toward thermal equilibrium of an adiabatic piston with the particle region localized around the piston. In the formulation that we consider, the system is bounded by adiabatic walls on each end and is divided into two chambers by a mobile and thermally insulating piston that can move without friction. We focus on the case when the initial state of the system is in mechanical equilibrium but not thermodynamic equilibrium: the pressures on the two sides are equal but the temperatures are not. Here we study the slow equilibration of the piston towards the state of thermodynamic equilibrium, which happens because asymmetric fluctuations on the two sides of the piston slowly transfer energy from the hotter to the colder chamber.

We assess the performance of the hybrid by comparing it to purely particle simulations, which are assumed to be "correct." Unlike in particle methods, in continuum methods we can trivially turn off fluctuations by not including stochastic fluxes in the Navier–Stokes equations. By turning off fluctuations in the continuum region we obtain a *deterministic hybrid* scheme, to be contrasted with the *stochastic hybrid* scheme in which fluctuations in the continuum region are consistent with those in the particle region. By comparing results between the deterministic and stochastic hybrid we are able to assess the importance of fluctuations. We find that the deterministic hybrid gives the wrong long-time behavior for both the diffusing spherical bead and the adiabatic piston, while the stochastic hybrid correctly reproduces the purely particle runs at a fraction of the computational cost. These examples demonstrate the need to include thermal fluctuations in the continuum solvers in hybrid particle-continuum methods.

2. Brief overview of the hybrid method. In this section we briefly introduce the basic concepts behind the hybrid method, relegating further technical details to later sections. Our scheme is based on an adaptive mesh and algorithm refinement (AMAR) methodology developed over the last decade in a series of works in which a DSMC particle fluid was first coupled to a deterministic compressible Navier–Stokes solver in three dimensions [26, 47], and then to a stochastic (fluctuating) continuum solver in one dimension [36]. This section presents the additional modifications to the previous algorithms necessary to replace the traditional DSMC particle method with the I-DSMC method [18], and a full three-dimensional dynamic coupling of a complex particle fluid [48] with a robust fluctuating hydrodynamic solver [9]. These novel techniques are discussed further in section 3.

Next, we briefly describe the two components of the hybrid, namely, the particle *microscopic* model and the continuum *macroscopic* solver, and then outline the domain decomposition used to couple the two, including a comparison with other proposed schemes. Both the particle algorithm and the macroscopic solver have already been described in detail in the literature, and, furthermore, both can easily be replaced by other methods. Specifically, any variant of DSMC and MPCD can be used as a particle algorithm, and any explicit finite-volume method can be used as a continuum solver. For this reason, in this paper we focus on the coupling algorithm.

2.1. Particle model. The particle method that we employ is the isotropic DSMC (I-DSMC) method, a *dense fluid*[1] generalization of the direct simulation Monte Carlo (DSMC) algorithm for rarefied gas flows [14]. The I-DSMC method is described in detail in [18], and here we only briefly summarize it. It is important to note that,

---

[1]Note that by a dense fluid we mean a fluid where the mean free path is small compared to the typical fluid interatomic distance.

like the traditional DSMC fluid, the I-DSMC fluid is an *ideal* fluid, that is, it has the equation of state (EOS) and structure of an ideal gas; it can be thought of as a viscous ideal gas. As we will see shortly, the lack of structure in the I-DSMC fluid significantly simplifies coupling to a continuum solver while retaining many of the salient features of a dense fluid.

In the I-DSMC method, the effect of interatomic forces is replaced by stochastic collisions between nearby particles. The interaction range is controlled via the *collision diameter D* or, equivalently, the density (hard-sphere volume fraction) $\phi = \pi N D^3/(6V)$, where N is the total number of particles in the simulation volume V. The strength of the interaction is controlled through the dimensionless *cross-section prefactor* χ, which is on the order of unity [18]. Stochastic collisions are processed at the end of every particle time step of duration Δt_P, and between collisions each particle i is streamed advectively with constant velocity $\boldsymbol{v}_i = \dot{\boldsymbol{r}}_i$. The spatial domain of the simulation, typically a rectangular region, is divided into *micro cells* of length $L_c \gtrsim D$, which are used to efficiently identify all particles that are within distance L_c of a given particle by searching among the particles in neighboring cells (each cell has 3^d neighboring cells, counting itself, where d is the spatial dimension). At each time step, *conservative collisions* occur between randomly chosen pairs of particles that are closer than a distance D apart; specifically, a random amount of momentum and energy is exchanged between the two particles with the probabilities of the collision outcomes obeying detailed balance. The I-DSMC algorithm can be viewed as a stochastic alternative to deterministic hard-sphere molecular dynamics (MD) [49], where hard spheres of diameter D collide when they touch.

In addition to the fluid (I-DSMC) particles there may be a number of additional spherical particles, which we refer to as *beads*, suspended in the I-DSMC fluid. The beads interact with each other and the fluid particles either deterministically, as *hard* spheres impermeable to and colliding with touching particles, or stochastically, as *permeable* spheres colliding stochastically with overlapping particles. Note that the sphere radii used for determining the fluid-fluid, fluid-bead, and bead-bead interaction distances may, in general, be different. Combining deterministic hard-sphere collisions with stochastic collisions requires a mixed time-driven and event-driven treatment, as in the stochastic event-driven MD (SEDMD) algorithm developed in [48]. We will use the SEDMD algorithm in the examples presented in this paper.

We have also developed an alternative purely time-driven fluid-bead coupling in which the fluid is allowed to permeate the beads and all of the particle interactions are stochastic. This leads to a much simpler algorithm that can also easily be parallelized. This distinction between hard and permeable spheres has analogues in other methods for complex fluids. For example, in lattice-Boltzmann simulations [50], beads can be modeled either as hard spheres (using a bounce-back collision rule at the lattice sites on the surface of the bead) or, more efficiently and commonly, as permeable spheres that let the fluid pass through them but experience a frictional force due to the fluid motion (exerting the opposite force back on the lattice sites with which they overlap).

2.2. Continuum model. At length scales and time scales larger than the molecular ones, the dynamics of the particle fluid can be coarse grained [51, 52] to obtain evolution equations for the slow macroscopic variables. Specifically, we consider the continuum conserved fields

(1)
$$\boldsymbol{U}(\boldsymbol{r},t) = \begin{bmatrix} \rho \\ \boldsymbol{j} \\ e \end{bmatrix} \cong \widetilde{\boldsymbol{U}}(\boldsymbol{r},t) = \sum_i \begin{bmatrix} m_i \\ \boldsymbol{p}_i \\ e_i \end{bmatrix} \delta\left[\boldsymbol{r} - \boldsymbol{r}_i(t)\right] = \sum_i \begin{bmatrix} 1 \\ \boldsymbol{v}_i \\ v_i^2/2 \end{bmatrix} m_i \delta\left[\boldsymbol{r} - \boldsymbol{r}_i(t)\right],$$

where the *conserved variables*, namely, the densities of mass ρ, momentum $\boldsymbol{j} = \rho\boldsymbol{v}$, and energy $e = \epsilon(\rho, T) + \frac{1}{2}\rho v^2$, can be expressed in terms of the *primitive variables*, mass density ρ, velocity $\boldsymbol{v}$, and temperature T; here ϵ is the internal energy density. Here the symbol $\cong$ means that we consider a stochastic field $\boldsymbol{U}(\boldsymbol{r}, t)$ that approximates the behavior of the true atomistic configuration $\widetilde{\boldsymbol{U}}(\boldsymbol{r}, t)$ over long length and time scales (compared to atomistic scales) in a certain integral average sense; notably, for sufficiently large cells the integral of $\boldsymbol{U}(\boldsymbol{r}, t)$ over the cell corresponds to the total particle mass, momentum, and kinetic energy contained inside the cell.

The evolution of the field $\boldsymbol{U}(\boldsymbol{r}, t)$ is modeled with the Landau–Lifshitz Navier–Stokes (LLNS) system of stochastic partial differential equations (SPDEs) in d dimensions, given in conservative form by

$$(2) \qquad \partial_t \boldsymbol{U} = -\boldsymbol{\nabla} \cdot [\boldsymbol{F}(\boldsymbol{U}) - \boldsymbol{\mathcal{Z}}(\boldsymbol{U}, \boldsymbol{r}, t)],$$

where the deterministic flux is taken from the traditional compressible Navier–Stokes–Fourier equations,

$$\boldsymbol{F}(\boldsymbol{U}) = \begin{bmatrix} \rho\boldsymbol{v} \\ \rho\boldsymbol{v}\boldsymbol{v}^T + P\boldsymbol{I} - \boldsymbol{\sigma} \\ (e + P)\boldsymbol{v} - (\boldsymbol{\sigma} \cdot \boldsymbol{v} + \boldsymbol{\xi}) \end{bmatrix},$$

where $P = P(\rho, T)$ is the pressure, the viscous stress tensor is $\boldsymbol{\sigma} = 2\eta\left[\frac{1}{2}(\boldsymbol{\nabla}\boldsymbol{v} + \boldsymbol{\nabla}\boldsymbol{v}^T) - \frac{(\boldsymbol{\nabla}\cdot\boldsymbol{v})}{d}\boldsymbol{I}\right]$ (we have assumed zero bulk viscosity), and the heat flux is $\boldsymbol{\xi} = \mu\boldsymbol{\nabla}T$. As postulated by Landau and Lifshitz [52, 53], the *stochastic flux*

$$\boldsymbol{\mathcal{Z}} = \begin{bmatrix} \boldsymbol{0} \\ \boldsymbol{\Sigma} \\ \boldsymbol{\Sigma} \cdot \boldsymbol{v} + \boldsymbol{\Xi} \end{bmatrix}$$

is composed of the stochastic stress tensor $\boldsymbol{\Sigma}$ and stochastic heat flux vector $\boldsymbol{\Xi}$, assumed to be mutually uncorrelated random Gaussian fields with a covariance

$$\langle \boldsymbol{\Sigma}(\boldsymbol{r}, t)\boldsymbol{\Sigma}^{\star}(\boldsymbol{r}', t')\rangle = \boldsymbol{C_{\Sigma}}\delta(t - t')\delta(\boldsymbol{r} - \boldsymbol{r}'),$$

$$\text{where } C_{ij,kl}^{(\boldsymbol{\Sigma})} = 2\bar{\eta}k_B\overline{T}\left(\delta_{ik}\delta_{jl} + \delta_{il}\delta_{jk} - \frac{2}{d}\delta_{ij}\delta_{kl}\right),$$

$$(3) \quad \langle \boldsymbol{\Xi}(\boldsymbol{r}, t)\boldsymbol{\Xi}^{\star}(\boldsymbol{r}', t')\rangle = \boldsymbol{C_{\Xi}}\delta(t - t')\delta(\boldsymbol{r} - \boldsymbol{r}'), \text{ where } C_{i,j}^{(\boldsymbol{\Xi})} = 2\bar{\mu}k_B\overline{T}\delta_{ij},$$

where overbars denote mean values.

As discussed in [9], the LLNS equations do not quite make sense written as a system of nonlinear SPDEs; however, they can be linearized to obtain a well-defined linear system whose equilibrium solutions are Gaussian fields with known covariances. We use a *finite-volume* discretization, in which space is discretized into N_c identical *macro cells* $\mathcal{V}_j$ of volume V_c, and the value $\boldsymbol{U}_j$ stored in cell $1 \leq j \leq N_c$ is the average of the corresponding variable over the cell

$$(4) \qquad \boldsymbol{U}_j(t) = \frac{1}{V_c}\int_{\mathcal{V}_j} \boldsymbol{U}(\boldsymbol{r}, t)d\boldsymbol{r} = \frac{1}{V_c}\int_{\mathcal{V}_j} \widetilde{\boldsymbol{U}}(\boldsymbol{r}, t)d\boldsymbol{r},$$

where $\widetilde{\boldsymbol{U}}$ is defined in (1). Time is discretized with a time step Δt_C, approximating $\boldsymbol{U}(\boldsymbol{r}, t)$ pointwise in time with $\boldsymbol{U}^n = \{\boldsymbol{U}_1^n, \ldots, \boldsymbol{U}_{N_c}^n\}$,

$$\boldsymbol{U}_j^n \approx \boldsymbol{U}_j(n\Delta t_C),$$

where $n \geq 0$ enumerates the macroscopic time steps. While not strictly necessary, we will assume that each macro cell consists of an integer number of micro cells (along each dimension of the grid), and similarly each macro time step consists of an integer number n_{ex} of micro time steps, $\Delta t_C = n_{ex} \Delta t_P$.

In addition to the cell averages $\boldsymbol{U}_j^n$, the continuum solver needs to store the continuum normal flux $\boldsymbol{F}_{j,j'}^n$ through each interface $I = \mathcal{V}_j \cap \mathcal{V}_{j'}$ between touching macro cells j and j' during a given time step,

$$(5) \qquad \boldsymbol{U}_j^{n+1} = \boldsymbol{U}_j^n - \frac{\Delta t}{V_c} \sum_{j'} S_{j,j'} \boldsymbol{F}_{j,j'}^n,$$

where $S_{j,j'}$ is the surface area of the interface, and $\boldsymbol{F}_{j,j'}^n = -\boldsymbol{F}_{j',j}^n$. Here we will absorb the various prefactors into a total transport (surface and time integrated flux) through a given macro cell interface I, $\boldsymbol{\Phi}_I^n = V_c^{-1} \Delta t S_I \boldsymbol{F}_I^n$, which simply measures the total mass, momentum, and energy transported through the surface I during the time interval from time t to time $t + \Delta t$. We arbitrarily assign one of the two possible orientations (direction of the normal vector) for each cell-cell interface. How the (integrated) fluxes $\boldsymbol{\Phi}_I^n$ are calculated from $\boldsymbol{U}_j^n$ does not formally matter; all that the hybrid method uses to advance the solution for one macro time step are $\boldsymbol{U}_j^n$, $\boldsymbol{U}_j^{n+1}$, and $\boldsymbol{\Phi}_I^n$. Therefore, any explicit conservative finite-volume method can be substituted trivially. Given this generality, we do not describe in any detail the numerical method used to integrate the LLNS equations; readers can consult [9] for further information.

2.3. Coupling between particle and continuum subdomains. The hybrid method we use is based on domain decomposition and is inspired by adaptive mesh refinement (AMR) methodology for conservation laws [54, 55]. Our coupling scheme closely follows previously developed methodology for coupling a traditional DSMC gas to a continuum fluid, first proposed in the deterministic setting in [26] and then extended to a fluctuating continuum method in [36]. The key new ingredient is the special handling of the collisional momentum and energy transport across cell interfaces, not found in traditional DSMC. For completeness, we describe the coupling algorithm in detail, including components already described in the literature.

We split the whole computational domain into *particle* and *continuum* subdomains, which communicate with each other through information near the particle-continuum interface I, assumed here to be oriented such that the flux $\boldsymbol{\Phi}_I$ measures the transport of conserved quantities from the particle to the continuum regions. In AMR implementations, subdomains are usually logically rectangular *patches*; in our implementation, we simply label each macro cell as either a particle cell or a continuum cell based on whatever criterion is appropriate, without any further restrictions on the shape or number of the resulting subdomains. For complex fluids applications, macro cells near beads (suspended solute) and sometimes near complex boundaries will be labeled as particle cells. The continuum solver is completely oblivious to what happens inside the particle subdomain, and thus it need not know how to deal with complex moving boundaries and suspended objects. Instead, the continuum solution feels the influence of boundaries and beads through its coupling with the particle subdomains.

The dynamic coupling between particle and continuum subdomains is best viewed as a mutual exchange of boundary conditions between the two regions. Broadly speaking, domain-decomposition coupling schemes can be categorized based on the type of boundary conditions each subdomain specifies for the other [56]. Our scheme

is closest to a *state-flux* coupling scheme based on the classification proposed in [56] for incompressible solvers (the term "velocity-flux" is used there since velocity is the only state variable). A *state-flux* coupling scheme is one in which the continuum solver provides to the particle solver the conserved variables $\boldsymbol{U}$ in the continuum *reservoir macro cells* near the particle-continuum interface I; that is, the continuum state is imposed as a boundary condition on the particle region. The particle solver provides to the continuum solver the flux $\boldsymbol{\Phi}_I$ through the interface I; that is, the particle flux is imposed as a boundary condition on the continuum subdomain. This aims to achieve continuity of both state variables and fluxes across the interface and ensures strict conservation, thus making the coupling rather robust. Note that state/flux information is only exchanged between the continuum and particle subdomains every n_{ex} particle (micro) time steps, at the beginning/end of a macro time step. A more detailed description of the algorithm is given in section 3.

2.3.1. Comparison with other coupling schemes. There are several hybrid methods in the literature coupling a particle method, and in particular, MD, with a continuum fluid solver [19, 20]. There are two main types of applications of such hybrids [57]. The first type is problems where the particle description is localized to a region of space where the continuum description fails, such as, for example, a complex boundary, flow near a corner, a contact line, a drop pinchoff region, etc. The second type is problems where the continuum method needs some transport coefficients, e.g., stress-strain relations, that are not known a priori and are obtained via localized MD computations. In the majority of existing methods a stationary solution is sought [30], and a deterministic incompressible or isothermal formulation of the Navier–Stokes equations is used in the continuum [57, 58]. By contrast, we are interested in a fully dynamic bidirectional coupling capable of capturing the full range of hydrodynamic effects, including sound and energy transport. We also wish to minimize the size of the particle regions and only localize the particle computations near suspended objects, making it important to minimize the artifacts at the interface. As we will demonstrate in this work, including thermal fluctuations in the continuum formulation is necessary to obtain a proper coupling under these demanding constraints.

The only other work we are aware of that develops a coupling between a fluctuating compressible continuum solver and a particle method, specifically MD, is a coupling scheme developed over the last several years by Coveney, Flekkoy, de Fabritiis, Delgado-Buscallioni, and collaborators [2, 37, 38], as reviewed in [20]. There are two important differences between their method and our algorithm. First, their method is (primarily, but not entirely) a flux-flux coupling scheme, unlike our state-flux coupling. Second, we do not use MD but rather I-DSMC, which, as discussed below, significantly simplifies the handling of the continuum-particle interface.

It is not difficult to impose boundary conditions for the continuum subdomain based on the particle data, and any consistent boundary condition for the PDE being solved can in principle be imposed. When the continuum solver is deterministic the fluctuations (often inappropriately referred to as "noise") in the particle data need to be filtered using some sort of spatial and temporal averaging [30], or the continuum solver needs to be robust [26, 47].

The difficult part in coupling schemes is the handling of the boundary conditions for the particle subdomain. It is very difficult to truncate a fluid particle region without introducing artifacts in the structure of the particle fluid, and the transport coefficients of the particle fluid (e.g., the pressure and viscosity) critically depend on the fluid structure. In most molecular simulations periodic boundary conditions are

used to avoid such artifacts; however, this is not possible in our case. In order to minimize artifacts at the boundary of the particle subdomain, most coupling schemes add an *overlap region* in addition to the reservoir region that we described above. In the overlap region, particles are simulated even though the region belongs to the continuum subdomain. The *structure* of the fluid in the overlap region is left to adjust to that in the particle subdomain, thus minimizing the artifacts. At the same time, the *dynamics* of the particles in the overlap region is somehow constrained to match the underlying continuum subdomain dynamics using, for example, various forms of constrained Langevin-type thermostats or (nonholonomic) constrained MD [58, 59]. To prevent particles from flying outside of the particle subdomain, various artificial constraining forces are added in the reservoir region, and typically some particle insertion/deletion scheme is used as well. The details can be very different and tricky, and we are not aware of any detailed comparison or even rudimentary mathematical analysis of the different types of coupling other than the stability analysis of four idealized coupling schemes presented in [56].

We do not use an overlap region in the coupling scheme we present here. First, and most importantly, because the I-DSMC fluid is structureless, an overlap region and the associated algorithmic complications are simply not required. The addition of an overlap region (added fluid mass) almost always introduces some delay and errors in the coupling, and it is usually assumed that this surface effect is negligible when the size of the overlap region is small compared to the particle region and when the hydro time steps are much larger than the particle time step. These are assumptions that we do not wish to make as we try to minimize the size of the particle subdomain.

The simple and direct coupling scheme we have presented, however, does not work if a structured stochastic particle fluid is used, such as, for example, a structured stochastic fluid like the stochastic hard-sphere dynamics (SHSD) fluid [18]. The SHSD fluid is weakly structured, so that at least particle insertion/deletion in the reservoir region is not a problem. However, an overlap region is necessary to smoothly match the fluid structure at the particle-continuum interface. Constraining the dynamics in the overlap region can be done in I-DSMC by introducing additional one-particle collisions (dissipation) that scatter the particle velocities so that their mean equals the continuum field values. However, it is not clear how to do this consistently when fluctuations are included in the continuum solver, and in this paper we restrict our focus on structureless (ideal) stochastic fluids.

3. Details of the coupling algorithm. The basic ideas behind the state-flux coupling were already described in section 2.3. In this section we describe in detail the two components of the particle-continuum coupling method, namely, the imposition of the continuum state as a boundary condition for the particle subdomain and the imposition of the particle flux as a boundary condition for the continuum subdomain. At the same time, we will make clear that our coupling is not purely of the state-flux form. The handling of the continuum subdomain is essentially unchanged from the pure continuum case, with the only difference being the inclusion of a refluxing step. The handling of the particle subdomain is more complex and explained in greater detail, including pseudocodes for several steps involved in taking a micro time step, as well as insertion of reservoir particles and the tracking of the particle fluxes.

3.1. State exchange. We first explain how the state in the *reservoir* macro cells bordering the particle subdomain, denoted by $\boldsymbol{U}_H^{(B)}$, is used by the particle algorithm. The micro cells that are inside the reservoir macro cells and are sufficiently close to the particle subdomain to affect it during a time interval Δt_P are labeled as *reservoir*

micro cells. For I-DSMC fluids, assuming the length of the micro cells along each dimension is $\boldsymbol{L}_c \gtrsim D$, the micro cells immediately bordering the particle subdomain as well as all of their neighboring micro cells need to be included in the reservoir region. An illustration of the particle and reservoir regions is given in Figures 1 and 2.
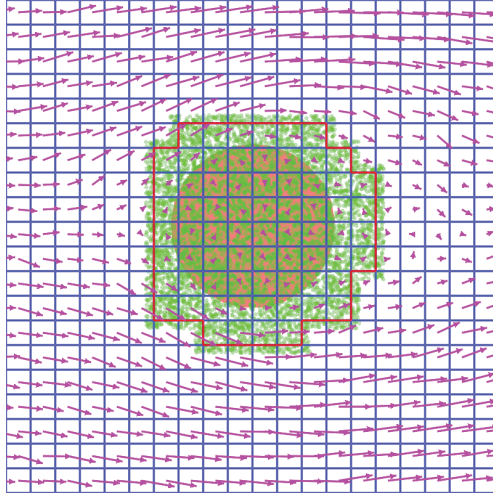

FIG. 2. *Illustration of a hybrid simulation of two-dimensional plug flow around a permeable stationary disk (red). The macroscopic grid is shown (dark blue lines), with each macro cell composed of 6×6 micro cells (not shown). The particle subdomain surrounds the disk, and the particle-continuum interface is shown (red). A snapshot of the I-DSMC particles is also shown (green), including the reservoir particles outside of the particle subdomain. The time-averaged velocity in each continuum cell is shown, revealing the familiar plug flow velocity field that is smooth across the interface. This example clearly demonstrates that the continuum solver feels the stationary disk through the particle subdomain even though the continuum solver is completely oblivious to the existence of the disk.*

At the beginning of each particle time step, *reservoir particles* are inserted randomly into the reservoir micro cells. The number of particles inserted is based on the target density in the corresponding reservoir macro cell. The velocities of the particles are chosen from a Maxwell–Boltzmann or Chapman–Enskog [60] distribution (see discussion in section 3.4.2) with mean velocity and temperature, and also their gradients if the Chapman–Enskog distribution is used, chosen to match the momentum and energy densities in the corresponding macro cell. The positions of the reservoir particles are chosen randomly uniformly (i.e., sampled from a Poisson spatial distribution) inside the reservoir cells, which does not introduce any artifacts in the fluid structure next to the interface because the I-DSMC fluid is ideal and structureless. This is a major advantage of ideal fluids over more realistic structured nonideal fluids. Note that in [48] we used a particle reservoir to implement open boundary conditions in purely particle simulations; the difference here is that the state in the reservoir cells comes from a continuum solver instead of a prespecified stationary flow solution.

The reservoir particles are treated just like the rest of the particles for the duration of a particle time step. First they are advectively propagated (streamed) along with all the other particles, and the total mass, momentum, and energy transported by particles advectively through the particle-continuum interface, $\boldsymbol{\Phi}_P^{(I)}$, is recorded. Particles that, at the end of the time step, are not in either a reservoir micro cell or in the particle subdomain are discarded, and then stochastic collisions are processed

between the remaining particles. In the traditional DSMC algorithm, collisions occur only between particles inside the same micro cell, and thus all of the particles outside the particle subdomain can be discarded [26, 36]. However, in the I-DSMC algorithm particles in neighboring cells may also collide, and thus the reservoir particles must be kept until the end of the time step. Collisions between a particle in the particle subdomain and a particle in a reservoir cell lead to collisional exchange of momentum and energy through the interface as well, and this contribution must also be included in $\boldsymbol{\Phi}_P^{(I)}$. Note that the reservoir particles at the very edge of the reservoir region do not have an isotropic particle environment, and thus there are artifacts in the collisions which they experience; however, this does not matter since it is only essential that the particles in the particle subdomain not feel the presence of the interface.

3.2. Flux exchange. After the particle subdomain is advanced for n_{ex} (micro) time steps, the particle flux $\boldsymbol{\Phi}_P^{(I)}$ is imposed as a boundary condition in the continuum solver so that it can complete its (macro) time step. This flux exchange ensures strict conservation, which is essential for long-time stability. Assume that the continuum solver is a one-step explicit method that uses only stencils of width one, that is, calculating the flux for a given macro cell-cell interface only uses the values in the adjacent cells. Under such a scenario, the continuum solver only needs to calculate fluxes for the cell-cell interfaces between continuum cells, and once the particle flux $\boldsymbol{\Phi}_P^{(I)}$ is known, the continuum time step (5) can be completed. It is obvious that in this simplified scenario the coupling is purely of the state-flux form. However, in practice we use a method that combines pieces of state and flux exchange between the particle and continuum regions.

First, the particle state is partly used to advance the continuum solver to the next time step. Our continuum solver is a multistage method and uses stencils of width two in each stage, thus using an effective stencil that can be significantly larger than one cell wide [9]. While one can imagine modifying the continuum solver to use specialized boundary stencils near the particle-continuum interface (e.g., one-sided differencing or extrapolation), this is not only more complex to implement but it is also less accurate. Instead, the continuum method solves for the fields over the whole computational domain (continuum patch in AMAR terminology) and uses hydrodynamic values for the particle subdomain (particle patch) obtained from the particle solver. These values are then used to calculate *provisional* fluxes $\boldsymbol{\Phi}_H^{(I)}$ and take a *provisional* time step, as if there were no particle subdomain. This makes the implementation of the continuum solver essentially oblivious to the existence of the particle regions; however, it does require the particle solver to provide reasonable conserved values for all of the macro particle cells. This may not be possible for cells where the continuum hydrodynamic description itself breaks down, for example, cells that overlap with or are completely covered by impermeable beads or features of a complex boundary. In such *empty cells* the best that can be done is to provide reasonable hydrodynamic values, for example, values based on the steady state compatible with the specified macroscopic boundary conditions. For partially empty cells, that is, macro cells that are only partially obscured, one can use the uncovered fraction of the hydro cell to estimate hydrodynamic values for the whole cell. In practice, we have found that as long as empty cells are sufficiently far from the particle-continuum interface (in particular, empty cells must not border the continuum subdomain) the exact improvised hydrodynamic values do not matter much. Note that for permeable beads there is no problem with empty cells since the fluid covers the whole domain.

Second, the provisional continuum fluxes are partly used to advance the particle subdomain to the same point in time as the continuum solver. Specifically, a linear interpolation between the current continuum state and the provisional state is used as a boundary condition for the reservoir particles. This temporal interpolation is expected to improve the temporal accuracy of the coupling, although we are not aware of any detailed analysis. Note that if $n_{ex} = 1$, this interpolation makes no difference and the provisional continuum fluxes are never actually used.

Once the particle solver advances n_{ex} time steps, a particle flux $\boldsymbol{\Phi}_P^{(I)}$ is available, and it is imposed in the continuum solver to finalize the provisional time step. Specifically, hydrodynamic values in the particle macro cells are overwritten based on the actual particle state, ignoring the provisional prediction. In order to correct the provisional fluxes, a *refluxing* procedure is used in which the state $\boldsymbol{U}_H^{(B)}$ in each of the continuum cells that border the particle-continuum interface are changed to reflect the particle $\boldsymbol{\Phi}_P^{(I)}$ rather than the provisional flux $\boldsymbol{\Phi}_H^{(I)}$,

$$\boldsymbol{U}_H^{(B)} \leftarrow \boldsymbol{U}_H^{(B)} - \boldsymbol{\Phi}_H^{(I)} + \boldsymbol{\Phi}_P^{(I)}.$$

This refluxing step ensures strict conservation and ensures continuity of the fluxes across the interface, in addition to continuity of the state.

3.3. Taking a macro time step. Algorithm 1 summarizes the hybrid algorithm and the steps involved in advancing both the simulation time by one macro time step $\Delta t_C = n_{ex} \Delta t_P$. Note that at the beginning of the simulation, we initialize the hydrodynamic values for the continuum solver, consistent with the particle data in the particle subdomain and generated randomly from the known (Gaussian) equilibrium distributions in the continuum subdomain.

Algorithm 1. Take a macro time step by updating the continuum state $\boldsymbol{U}_H$ from time t to time $t + \Delta t_C$.

---

1. *Provisionally advance the continuum solver*: Compute a *provisional* macro solution $\boldsymbol{U}_H^{next}$ at time $t + \Delta t_C$ everywhere, including the particle subdomain, with an estimated (integrated) provisional flux $\boldsymbol{\Phi}_H^{(I)}$ through the particle-continuum interface. Reset the particle flux $\boldsymbol{\Phi}_P^{(I)} \leftarrow \boldsymbol{0}$.

2. *Advance the particle solver*: Take n_{ex} micro time steps (see Algorithm 2):
   (a) At the beginning of each particle time step reservoir particles are inserted at the boundary of the particle subdomain with positions and velocities based on a linear interpolation between $\boldsymbol{U}_H$ and $\boldsymbol{U}_H^{next}$ (see Algorithm 4). This is how the continuum state is imposed as a boundary condition on the particle subdomain.
   (b) All particles are propagated advectively by Δt_P, and stochastic collisions are processed, accumulating a particle flux $\boldsymbol{\Phi}_P^{(I)}$ (see Algorithm 3).

3. *Synchronize* the continuum and particle solutions:
   (a) *Advance*: Accept the provisional macro state, $\boldsymbol{U}_H \leftarrow \boldsymbol{U}_H^{next}$.
   (b) *Correct*: The continuum solution in the particle subdomain $\boldsymbol{U}_H^{(P)}$ is replaced with cell averages of the particle state at time $t + \Delta t_C$, thus forming a composite state over the whole domain.
   (c) *Reflux*: The continuum solution $\boldsymbol{U}_H^{(B)}$ in the macro cells bordering the particle subdomain is corrected based on the particle flux, $\boldsymbol{U}_H^{(B)} \leftarrow \boldsymbol{U}_H^{(B)} - \boldsymbol{\Phi}_H^{(I)} + \boldsymbol{\Phi}_P^{(I)}$. This effectively imposes the particle flux as a boundary condition on the continuum and ensures conservation in the hybrid update.

(d) *Update* the partitioning between particle and continuum cells if necessary. Note that this step may convert a continuum cell into an *unfilled* (devoid of particles) particle cell.

3.4. Taking a micro time step. Taking a micro time step is described in Algorithm 2, and the remaining subsections give further details on the two most important procedures used. The initial particle configuration can most easily be generated by marking all macro cells in the particle domain as unfilled.

Algorithm 2. Take an I-DSMC time step. We do not include details about the handling of the non-DSMC (solute) particles here. Note that the micro time step counter n_P should be reinitialized to zero after every n_{ex} time steps.

1. Visit all macro cells that overlap the reservoir region or are unfilled (i.e., recently converted from continuum to particle) one by one, and insert trial particles in each of them based on the continuum state $\boldsymbol{U}_H$, as described in Algorithm 4.

2. Update the clock $t \leftarrow t + \Delta t_P$ and advance the particle subdomain step counter $n_P \leftarrow n_P + 1$. Note that when an event-driven algorithm is used this may involve processing any number of events that occur over the time interval Δt_P [48].

3. Move all I-DSMC particles to the present time, updating the total kinetic mass, momentum, and energy flux through the coupling interface $\boldsymbol{F}_P^{(I)}$ whenever a particle crosses from the particle into the continuum subdomain or vice versa, as detailed in Algorithm 3.

4. Perform stochastic collisions between fluid particles, including the particles in the reservoir region. Keep track of the total *collisional* momentum and energy flux through the coupling interface by accounting for the amount of momentum $\Delta \boldsymbol{p}_{ij} = m\Delta \boldsymbol{v}_{ij}$ and kinetic energy Δe_{ij} transferred from a regular particle i (in the particle subdomain) to a reservoir particle j (in the continuum subdomain), $\boldsymbol{\Phi}_P^{(I)} \leftarrow \boldsymbol{\Phi}_P^{(I)} + \left[0, \Delta \boldsymbol{p}_{ij}, \Delta e_{ij}\right]$.

5. Remove all particles from the continuum subdomain.

6. Linearly interpolate the continuum state to the present time,
$$\boldsymbol{U}_H \leftarrow \frac{(n_{ex} - n_P)\, \boldsymbol{U}_H + \boldsymbol{U}_H^{next}}{n_{ex} - n_P + 1}.$$

3.4.1. Particle flux tracking. As particles are advected during a particle time step, they may cross from a particle to a continuum cell and vice versa, and we need to keep track of the resulting fluxes. Note that a particle may cross up to d cell interfaces near corners, where d is the spatial dimension, and may even recross the same interface twice near hard-wall boundaries. Therefore, ray tracing is the most simple and reliable way to account for all particle fluxes correctly. For the majority of the particles in the interior, far from corners and hard walls, the usual quick DSMC update will, however, be sufficient.

Algorithm 3. Move the fluid particle i from time t to time $t+\Delta t_P$ and determine whether it crosses the coupling interface. We will assume that during a particle time step no particle can move more than one macro cell length along each dimension (in practice particles typically move only a fraction of a micro cell).

1. Store information on the macro cell c_{old} to which the particle belongs, and then tentatively update the position of the particle $\boldsymbol{r}_i \leftarrow \boldsymbol{r}_i + \boldsymbol{v}_i \Delta t_P$, and find the tentative macro cell c_i and micro cell b_i to which the particle moves, taking into account periodic boundary conditions.

2. If c_{old} and c_i are near a boundary, go to step 3. If $c_{old} \equiv c_i$, or if c_{old} and c_i are both continuum or are both particle cells and at least one of them is not at a corner, accept the new particle position and go to step 4.

3. Undo the tentative particle update, $\boldsymbol{r}_i \leftarrow \boldsymbol{r}_i - \boldsymbol{v}_i \Delta t_P$, and then ray trace the path of the particle during this time step from one macro cell-cell interface I_c to the next, accounting for boundary conditions (e.g., wrapping around periodic boundaries and colliding the particle with any hard walls it encounters [48]). Every time the particle crosses from a particle to a continuum cell, update the particle flux at the cell interface I_c, $\boldsymbol{\Phi}_P^{(I_c)} \leftarrow \boldsymbol{\Phi}_P^{(I_c)} + \left[m, m\boldsymbol{v}_i, mv_i^2/2\right]$, and similarly if the particle crosses from a continuum to a particle cell, $\boldsymbol{\Phi}_P^{(I_c)} \leftarrow \boldsymbol{\Phi}_P^{(I_c)} - \left[m, m\boldsymbol{v}_i, mv_i^2/2\right]$.

4. If the new particle micro cell b_i is neither in the particle subdomain nor in the reservoir region, remove the particle from the system.

3.4.2. Inserting reservoir particles. At every particle time step, *reservoir particles* need to be inserted into the reservoir region or unfilled continuum cells. These particles may later enter the particle subdomain or they may be discarded, while the trial particles in unfilled cells will be retained unless they leave the particle subdomain (see Algorithm 3). When inserting particles in an unfilled macro cell, it is important to maintain strict momentum and energy conservation by ensuring that the inserted particles have exactly the same total momentum and energy as the previous continuum values. This avoids global drifts of momentum and energy, which will be important in several of the examples we present in section 4. It is not possible to ensure strict mass conservation because of quantization effects, but by using smart (randomized) rounding one can avoid any spurious drifts in the mass.

The velocity distribution for the trial particles should, at first approximation, be chosen from a Maxwell–Boltzmann distribution. However, it is well known from kinetic theory that the presence of shear and temperature gradients skews the distribution, specifically, to first order in the gradients the Chapman–Enskog distribution is obtained [60, 61]. It is important to note that only the *kinetic* contribution to the viscosity enters into the Chapman–Enskog distribution and not the full viscosity, which also includes a *collisional* viscosity for dense gases [60]. Previous work on deterministic DSMC hybrids has, as expected, found that using the Chapman–Enskog distribution improves the accuracy of the hybrids [26, 30]. However, in the presence of transient fluctuations and the associated transient gradients, it becomes less clear what the appropriate distribution to use is, as we observe numerically in section 4.1.

Certainly at pure equilibrium we know that the Maxwell–Boltzmann distribution is correct despite the presence of fluctuating gradients. At the same time, we know that the CE distribution ought to be used in the presence of constant macroscopic

gradients, as is required in order to obtain the correct kinetic contribution to the viscous stress tensor [62]. The inability to estimate time-dependent mean gradients from just a single fluctuating realization forces the use of instantaneous gradients, which are unreliable due to the statistical uncertainty and can become unphysically large when N_0 is small (say, $N_0 < 50$). In some cases the background macroscopic gradients may be known a priori (for example, Couette or Fourier flows) and they can be used without trying to numerically estimate them; otherwise, they can be assumed to be zero or numerically estimated by performing some spatio-temporal averaging [36]. Note that here we assume that the particles are uniformly distributed (Poisson spatial distribution) as in an ideal gas, and we do not try to take the density gradient into account (but see the procedure described in the appendix of [36]).

Algorithm 4. Insert trial particles into the reservoir or unfilled macro cell c with centroid $\boldsymbol{r}_c$, taking into account the target density ρ_c, velocity $\boldsymbol{v}_c$, and temperature T_c in cell c, calculated from the conserved cell state $\boldsymbol{U}_c$. If using the Chapman–Enskog distribution, also take into account the estimated local shear rate $\boldsymbol{\nabla}_c\boldsymbol{v}$ and local temperature gradient $\boldsymbol{\nabla}_c T$.

1. Build a list $\mathcal{L}_{rc}$ of the N_{rc} micro cells contained in the macro cell c that need to be filled with particles. This typically excludes micro cells that are partially covered by impermeable beads or boundaries so as to avoid generating overlaps.

2. Determine the total number N_p of trial particles to insert into the reservoir portion of c by sampling from the binomial distribution

$$P(N_p) = \left(\begin{array}{c} \bar{N}_p \\ N_p \end{array} \right) p^{N_p} (1-p)^{\bar{N}_p - N_p},$$

where $\bar{N}_p = \lfloor \rho_c V_c / m \rfloor$ is the total expected number of particles in c, V_c is the volume of c, and $p = N_{rc}/N_{sc}$, where N_{sc} is the number of micro subcells per macro cell. For sufficiently large $\bar{N}_p$ this can be well approximated by a Gaussian distribution, which can be sampled more quickly.

3. For each of the N_p trial particles to be inserted, do:
   (a) Choose a micro cell b uniformly from the N_{rc} micro cells in the list $\mathcal{L}_{rc}$.
   (b) Generate a random particle position $\boldsymbol{r}_i \leftarrow \boldsymbol{r}_c + \boldsymbol{r}_{rel}$ uniformly inside micro cell b.
   (c) Generate a random relative velocity for the particle $\boldsymbol{v}_{rel}$ from the Maxwell–Boltzmann or Chapman–Enskog distribution [61], and set the particle velocity by taking into account the desired continuum state in cell c and, if available, its estimated gradient, $\boldsymbol{v}_i \leftarrow \boldsymbol{v}_c + (\boldsymbol{\nabla}_c\boldsymbol{v})\,\boldsymbol{r}_{rel} + \boldsymbol{v}_{rel}$.
   (d) If the cell c is an *unfilled* macro cell, keep track of the total momentum $\boldsymbol{P}$ and energy E of the N_p trial particles, to be adjusted in step 4 for conservation.

4. If cell c was unfilled and $N_p > 1$, then correct the particle velocities to match the desired total momentum $\boldsymbol{P}_c = \boldsymbol{p}_c V_c$ and energy $E_c = V_c e_c$ inside macro cell c, thus ensuring exact conservation:
   (a) Calculate the scaling factor

$$\alpha^2 = \frac{E_c - P_c^2/(2mN_p)}{E - P^2/(2mN_p)}$$

   and velocity shift $\Delta\boldsymbol{v} = (\boldsymbol{P} - \boldsymbol{P}_c)/(\alpha m N_p)$.
   (b) Scale and shift the velocity for every trial particle i, $\boldsymbol{v}_i \leftarrow \alpha\,(\boldsymbol{v}_i - \Delta\boldsymbol{v})$.

---

4. Results. In this section we provide extensive tests of the hybrid scheme, in both equilibrium and nonequilibrium situations, and in both two and three dimensions. Our goal is to access how well the hybrid method can reproduce results obtained with a purely particle method, which we consider as a baseline for comparison with hybrid runs.

We have implemented our hybrid method in a code that can handle both two- and three-dimensional systems. Of course, one can study one- and two-dimensional flows with the three-dimensional code by using periodic boundaries along the remaining dimensions. We refer to this as quasi one- or two-dimensional simulations. At the same time, both (I-)DSMC and the LLNS equations have a truly two-dimensional formulation, which we have also implemented for testing purposes. Transport coefficients for two-dimensional particle models formally diverge in the infinite-time limit (see the discussion in [63]), and it is not obvious that the Navier–Stokes equations are a proper coarse-graining of the microscopic dynamics. However, this divergence is very slow (logarithmic) and will be mollified (bounded) by finite system size, and we will therefore not need to concern ourselves with these issues. In the first three examples, we use the three-dimensional particle and continuum codes, and use the two-dimensional code only for the adiabatic piston example for computational reasons.

We have also implemented continuum solvers for both the full nonlinear and the *linearized* LLNS equations. As discussed in more detail in [9], the nonlinear LLNS equations are mathematically ill-defined and this can lead to breakdown in the numerical solution such as negative densities or temperatures. At the same time, the linearized equations are not able to describe a wide range of physical phenomena such as the effect of fluctuations on the mean flow; they also omit a number of terms of order N_0^{-1}, where N_0 is the average number of particles per continuum cell (e.g., the center-of-mass kinetic energy). If the number of particles per continuum cell is sufficiently large (in our experience, $N_0 > 75$) the fluctuations are small and the difference between the linear and nonlinear hydrodynamic solvers is very small, and we prefer to use the nonlinear solver. We will use $N_0 \sim 100$ in our hybrid simulations.

In our implementation, the continuum solver can be either the more accurate third-order Runge–Kutta (RK3) temporal integrator developed in [9] or a more classical stochastic MacCormack integrator [6]. The analysis in [9] shows that obtaining reasonably accurate equilibrium fluctuations with predictor-corrector methods for the diffusive fluxes, as used in the MacCormack scheme, requires using a continuum time step Δt_C that is a fraction of the CFL stability limit Δt_{CFL}. In the simulations we present here we have typically used a time step $\Delta t \approx 0.2\Delta t_{CFL}$, which is typically still about 5 times larger than the particle time step Δt_P, and we have found little impact of the exact value of Δt_C.

The hybrid method requires estimates of the transport coefficients of the particle fluid, notably the viscosity and the thermal conductivity. For traditional DSMC at low densities there are rather accurate theoretical estimates of the viscosity and thermal conductivity [64, 65]; however, it is nontrivial to obtain reasonably accurate theoretical values at the higher densities we use in I-DSMC because of the importance of multiparticle correlations. We therefore estimate the transport coefficients numerically. For this purpose we simulate a system with periodic boundaries along two of the directions and isothermal stick wall boundaries along the other direction. To estimate the viscosity we apply a shear flow by moving one of the wall boundaries at constant speed, inducing a Couette flow with an approximately constant shear gradient $\bar{\boldsymbol{\nabla}}\boldsymbol{v} = \frac{1}{2}(\boldsymbol{\nabla}\boldsymbol{v} + \boldsymbol{\nabla}\boldsymbol{v}^T)$. We then calculate the steady-state stress tensor $\boldsymbol{\sigma}$

by averaging over all particles i and colliding pairs of particles ij over a long-time interval Δt,

$$\boldsymbol{\sigma} = \boldsymbol{\sigma}_k + \boldsymbol{\sigma}_c = m \langle \boldsymbol{v}_i \otimes \boldsymbol{v}_i \rangle + \frac{\langle \boldsymbol{r}_{ij} \otimes \Delta \boldsymbol{p}_{ij} \rangle_c}{\Delta t},$$

where $\boldsymbol{\sigma}_k = P\boldsymbol{I} + 2\eta_k \bar{\boldsymbol{\nabla}} \boldsymbol{v}$ is the kinetic contribution giving the *kinetic viscosity* η_k, and $\boldsymbol{\sigma}_c = 2\eta_c \bar{\boldsymbol{\nabla}} \boldsymbol{v}$ is the collisional contribution giving the *collisional viscosity* η_c, $\eta = \eta_k + \eta_c$. We exclude particles that are close to the wall boundaries when calculating these averages to minimize finite-size effects. Similarly, for the thermal conductivity we apply a small constant temperature gradient $\boldsymbol{\nabla} T$ by setting the two walls at different temperatures, and we also impose the required density gradient to maintain mechanical equilibrium (constant pressure). We then calculate the steady-state heat flux vector $\boldsymbol{\xi} = \mu \boldsymbol{\nabla} T$,

$$\boldsymbol{\xi} = \boldsymbol{\xi}_k + \boldsymbol{\xi}_c = m \left\langle \frac{v_i^2}{2} \boldsymbol{v}_i \right\rangle + \frac{\langle (\Delta e_{ij}) \, \boldsymbol{r}_{ij} \rangle_c}{\Delta t},$$

from which we obtain the kinetic and collisional contributions to the thermal conductivity. There are alternative methods that one can use to calculate the transport coefficients, using both equilibrium and nonequilibrium settings; however, we have found the above method to be most accurate for a given computational effort if only moderate accuracy is desired. Results from different nonequilibrium methods are found to be within 5–10% of each other.

4.1. Mismatch at the interface. Previous work has studied a hybrid scheme very similar to the one described here for several quasi one-dimensional situations [36]. Reference [36] first studied a pure equilibrium situation in which one part of a periodic domain was covered by a particle subdomain, and found that the stochastic hybrid scheme was able to reproduce the spatio-temporal correlations in equilibrium fluctuations very well, although some mismatch at the particle/continuum interface was found. Here we explore this mismatch more carefully by studying a quasi one-dimensional periodic system where the middle portion of the domain is filled with particles and the rest is continuum. In each macro cell, we compute the average density, temperature, and velocity and their variances with high accuracy.

We first calculate the mean conserved quantities (mean density, momentum density, and energy density) in each macro cell, and then calculate the mean velocity and temperature from those, for example, $\langle v \rangle = \langle j \rangle / \langle \rho \rangle$ instead of averaging the instantaneous velocities, $\langle v \rangle = \langle j/\rho \rangle$. As shown in [66], the approach we use leads to an unbiased estimate of the mean, while the latter has a bias when there are correlations between the fluctuations of the different hydrodynamic variables. The variances are estimated from the instantaneous values, e.g., $\langle \delta v^2 \rangle = \langle (j/\rho)^2 \rangle - \langle j/\rho \rangle^2$. It is possible to construct unbiased estimates for the variances as well [66]; however, this is somewhat more involved and the bias is rather small compared to the artifacts we are focusing on.

In Figure 3 we show the means and variances along the length of the system, normalized by the expected values [6]. For velocity, the mean is zero to within statistical uncertainty in both the particle and continuum subdomains, and we do not show it in the figure. However, a small mismatch is clearly seen in Figure 3(a) between the density and temperature in the particle and continuum subdomains. The mismatch is such that the pressure $\langle p \rangle = \langle \rho \rangle R \langle T \rangle$ is constant across the interface; that is,
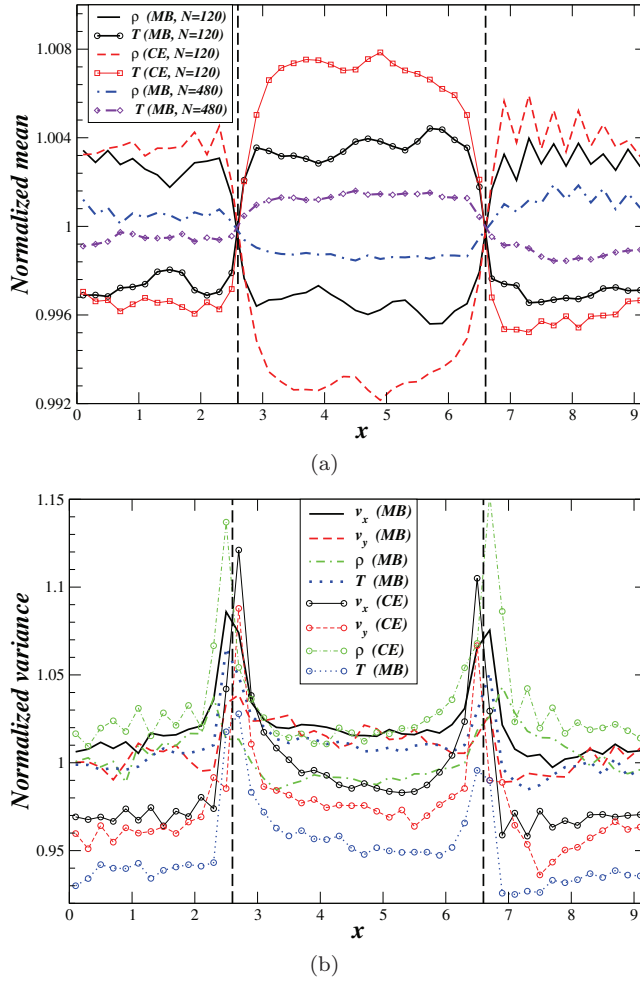
Fig. 3. *Normalized means and variances of the hydrodynamic variables in each of the 46 macro-scopic cells of a quasi one-dimensional periodic system where the middle portion (20 continuum cells, between dashed vertical lines) is the particle subdomain. The velocities of the reservoir particles are either samples from a Maxwell–Boltzmann (MB, lines only) or a Chapman–Enskog (CE, lines and symbols) distribution. (a) Unbiased [66] estimates of the mean density and temperature (we averaged about $2.5 \cdot 10^6$ samples taken every macro step) for two different sizes of the continuum cells, ones containing $N_0 = 120$ particles on average, and ones containing $N_0 = 480$ particles. Only the Maxwell–Boltzmann results are shown for the larger cells for clarity, with similar results observed for Chapman–Enskog. (b) Estimates of the variances of the hydrodynamic variables for $N_0 = 120$ particles per continuum cell (also an average over $2.5 \cdot 10^6$ samples).*

the particle and continuum subdomains are in mechanical equilibrium but not in true thermodynamic equilibrium. In Appendix A we show that this kind of mismatch is expected because the average particle fluxes coming from the reservoir particles inserted into the continuum subdomains have a bias of order N_0^{-1}, where N_0 is the average number of particles per macroscopic cell. Because our coupling matches both the state and the fluxes across the interface, this bias makes it impossible for the particle and continuum to reach true thermodynamic equilibrium. The theory in Appendix A suggests that the size of the mismatch is of order N_0^{-1}, consistent with the results in

Figure 3(a); however, the crude theoretical estimates do not actually give the steady state since they assume equilibrium to begin with. That the cells near the interface are not in equilibrium is reflected in the variances of the hydrodynamic variables, which have a spike near the particle-continuum interface, as shown in Figure 3(b). We have observed that the relative magnitude of this spike does not depend on N_0.

The cause of the mismatch is the fact that we use the instantaneous values of the local density, velocity, and temperature when generating the velocities for the reservoir particles. This is necessary because we cannot in general obtain estimates of the time-dependent mean values from running a single realization and are forced to use the instantaneous values. One can use some sort of spatio-temporal averaging to obtain estimates of the local means; however, this introduces additional time and length scales into the algorithm that do not have an obvious physical interpretation. For steady-state problems it may be possible to use running means to avoid the mismatch we observe, but this is not possible in a general dynamic context. Deeper theoretical understanding of the connection between the microscopic dynamics and the LLNS equations is necessary to design a more consistent approach.

Another complex issue that arises in the fluctuating hybrid is whether to use the Maxwell–Boltzmann or the Chapman–Enskog distribution when generating the velocities of the reservoir particles (see discussion in section 3.4.2). In Figure 3(a) we compare the size of the mismatch at the interface when the Maxwell–Boltzmann and Chapman–Enskog distributions are used. For the Chapman–Enskog distribution we obtained a local estimate of the gradient using simple centered differences of the instantaneous hydrodynamic variables. As seen in the figure, there is a greater discrepancy when the Chapman–Enskog distribution is used, especially for the variances. We will therefore adopt a compromise in which we use the Maxwell–Boltzmann distribution for all calculations reported here. In cases when there is a macroscopic background gradient that is specified a priori (e.g., shear flows) we can use that gradient in the Chapman–Enskog distribution.

Note that we have performed numerous additional quasi one-dimensional tests of the coupling that we do not describe here for the sake of brevity. For example, we have tested the matching of the shear stress tensor in a shear flow parallel to the particle-continuum interface by verifying that a linear velocity profile is obtained without any slope discontinuity at the interface (see Appendix D.2 in [37]). We have also reproduced the results reported in [36], such as the presence of unphysical long-range correlations in the fluctuations in the particle region when the deterministic instead of the stochastic hybrid is used.

4.2. Dynamic structure factors. The hydrodynamics of the spontaneous thermal fluctuations in the I-DSMC fluid is expected to be described by the LLNS equations for the fluctuating field $\boldsymbol{U} = (\rho_0 + \delta\rho, \delta v, T_0 + \delta T)$ linearized around a reference equilibrium state $\boldsymbol{U}_0 = (\rho_0, \boldsymbol{v}_0 = \boldsymbol{0}, T_0)$ [9, 53, 67]. By solving these equations in the Fourier wavevector-frequency $(\boldsymbol{k}, \omega)$ domain for $\widehat{\boldsymbol{U}}(\boldsymbol{k}, \omega) = (\widehat{\delta\rho}, \widehat{\delta v}, \widehat{\delta T})$ and performing an ensemble average over the fluctuating stresses, we can obtain the equilibrium (stationary) spatio-temporal correlations (covariance) of the fluctuating fields. We express these correlations in terms of the 3×3 symmetric positive-definite *hydrodynamic structure factor matrix* $\boldsymbol{S}_H(\boldsymbol{k}, \omega) = \left\langle \widehat{\boldsymbol{U}}\widehat{\boldsymbol{U}}^* \right\rangle$ [9], which is essentially the spatio-temporal spectrum of the fluctuating fields. Integrating $\boldsymbol{S}_H(\boldsymbol{k}, \omega)$ over frequencies gives the *hydrostatic structure factor matrix* $\boldsymbol{S}_H(\boldsymbol{k})$, which turns out to be diagonal since in any given snapshot the hydrodynamic variables are uncorrelated at equilibrium.

We nondimensionalize $\boldsymbol{S}_H(\boldsymbol{k}, \omega)$ so that $\boldsymbol{S}_H(\boldsymbol{k})$ is the identity matrix. For example, the density-density correlations are given by the dimensionless structure factor

$$S_\rho(\boldsymbol{k}, \omega) = \left(\rho_0 c_0^{-2} k_B T_0\right)^{-1} \left\langle \hat{\rho}(\boldsymbol{k}, \omega) \hat{\rho}^\star(\boldsymbol{k}, \omega) \right\rangle,$$

and we express the spatio-temporal cross-correlation between density and velocity through the dimensionless structure factor

$$S_{\rho,v}(\boldsymbol{k}, \omega) = \left(\rho_0 c_0^{-2} k_B T_0\right)^{-\frac{1}{2}} \left(\rho_0^{-1} k_B T_0\right)^{-\frac{1}{2}} \left\langle \hat{\rho}(\boldsymbol{k}, \omega) \hat{v}^\star(\boldsymbol{k}, \omega) \right\rangle,$$

where $c_0^2 = k_B T_0 / m$ is the isothermal speed of sound. Reference [36] demonstrated that a hybrid method very similar to the one we described here correctly reproduces the density-density time-correlation function $S_\rho(\boldsymbol{k}, t)$ for large wavelengths in a quasi one-dimensional periodic system. The density-density dynamic structure factor $S_\rho(\boldsymbol{k}, \omega)$ is often the only one considered because it is accessible experimentally via light scattering measurements and thus most familiar [67]. The full dynamic structure factor matrix $\boldsymbol{S}_H(\boldsymbol{k}, \omega)$ is a more complete measure of the spatio-temporal evolution of the thermal fluctuations and includes both sound (hyperbolic) and dissipative (diffusive) effects. It is therefore important to show that the hybrid scheme correctly reproduces $\boldsymbol{S}_H(\boldsymbol{k}, \omega)$ as compared to a purely particle simulation and demonstrate that the hybrid is capable of capturing the propagation of spontaneous thermal fluctuations across the particle-continuum interface.

4.2.1. Bulk dynamic structure factor. For a bulk fluid, i.e., for periodic boundary conditions, it is well known [67] that the density-density component $S_\rho(\boldsymbol{k}, \omega)$ and the temperature-temperature component $S_T(\boldsymbol{k}, \omega)$ of $\boldsymbol{S}_H(\boldsymbol{k}, \omega)$ exhibit three peaks for a given wavevector $\boldsymbol{k}$. There is one central Rayleigh peak at $\omega = 0$ which comes from entropic fluctuations. There are also two symmetric Brillouin peaks at $\omega \approx c_s k$, where c_s is the adiabatic speed of sound, which come from the isoentropic propagation of sound waves induced by the fluctuations. For the components of the velocity parallel to the wavevector the dynamic structure factors $S_{\boldsymbol{v}_\parallel}(\boldsymbol{k}, \omega)$ exhibit all three peaks, while for the component perpendicular to the wavevector $S_{v_\perp}(\boldsymbol{k}, \omega)$ lacks the central peak.

Figure 4 shows selected dynamic structure factors for a quasi two-dimensional system with periodic boundary conditions. The simulation box is composed of $10 \times 10 \times 1$ macro cells, each cell containing about 120 particles. Finite-volume averages of the hydrodynamic conserved variables and the corresponding spatial discrete Fourier transforms (DFTs) were then calculated for each cell every 10 macro time steps, and a temporal DFT was used to obtain discrete dynamic structure factors [9] for several wavenumbers. In the first set of particle runs, the whole domain was filled with particles and the macro cells were used only to sample hydrodynamic fields. In the second set of hybrid runs, the central portion of the simulation box was split along the x axes and designated as a particle subdomain, and the remaining two-thirds of the domain were continuum.

In Figure 4 we show the results for a wavevector $\boldsymbol{k}$ that is neither parallel nor perpendicular to the particle-continuum interface so as to test the propagation of both perpendicular and tangential fluctuations across the interface. The results show very little discrepancy between the pure particle and the hybrid runs, and they also conform to the theoretical predictions based on the LLNS equations. Perfect agreement is not expected because the theory is for the spectrum of the continuum field while the numerical results are discrete spectra of cell averages of the field, a distinction that
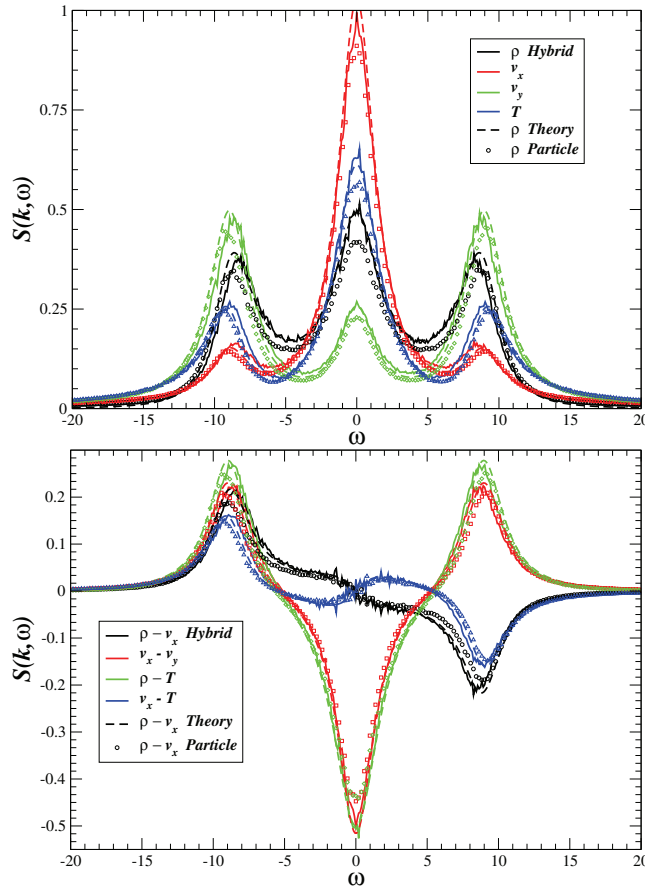
FIG. 4. *Discrete dynamic structure factors for waveindices* $q_x = k_x L_x/(2\pi) = 1$, $q_y = 2$, *and* $q_z = 0$ *for a quasi two-dimensional system with lengths* $L_x = L_y = 2$, $L_z = 0.2$, *split into a grid of* $10 \times 10 \times 1$ *macro cells. Each macro cell contains about* 120 *I-DSMC particles of diameter* $D = 0.04$ *(density* $\phi = 0.5$ *and collision frequency* $\chi = 0.62$*). An average over 5 runs each containing* 10^4 *temporal snapshots is employed. We perform pure particle runs and also hybrid runs in which only a strip of 4 macro cells along the* x *axes was filled with particles and the rest handled with a continuum solver. The different hydrodynamic pairs of variables are shown with different colors, using a solid line for the result from the hybrid runs, symbols for the results of the pure particle runs, and a dashed line for the theoretical predictions based on the linearized LLNS equations (solved using the computer algebra system Maple).* (Top) *The diagonal components* $S_\rho(\boldsymbol{k}, \omega)$, $S_{v_x}(\boldsymbol{k}, \omega)$, $S_{v_y}(\boldsymbol{k}, \omega)$, *and* $S_T(\boldsymbol{k}, \omega)$. (Bottom) *The off-diagonal components (cross-correlations)* $S_{\rho, v_x}(\boldsymbol{k}, \omega)$, $S_{v_x, v_y}(\boldsymbol{k}, \omega)$, $S_{\rho, T}(\boldsymbol{k}, \omega)$, *and* $S_{v_x, T}(\boldsymbol{k}, \omega)$.

becomes important when the wavelength is comparable to the cell size. Additionally, even purely continuum calculations do not reproduce the theory exactly because of spatio-temporal discretization artifacts.

4.2.2. Dynamic structure factors for finite systems. The previous section discussed the "bulk" dynamic structure factors, as obtained by using periodic boundary conditions. For nonperiodic (i.e., finite) systems, equilibrium statistical mechanics requires that the static structure factor be oblivious to the presence of walls. However, the dynamic structure factors exhibit additional peaks due to the reflections of sound waves from the boundaries. At a hard-wall boundary surface $\partial\Omega$ with normal

vector $\boldsymbol{n}$ either Dirichlet or von Neumann boundary conditions need to be imposed on the components of the velocity and the temperature (the boundary condition for density follows from these two). The following are two particularly common types of boundaries:

- *Thermal walls*, for which a stick condition is imposed on the velocity, $\boldsymbol{v}_{\partial\Omega} = 0$, and the temperature is fixed, $T_{\partial\Omega} = T_0$.
- *Adiabatic walls*, for which a slip condition is imposed on the velocity, $\boldsymbol{n}\cdot\boldsymbol{\nabla}\boldsymbol{v}_{\parallel} = 0$ and $v_{\perp} = 0$, and there is no heat conduction through the wall, $\boldsymbol{n}\cdot\boldsymbol{\nabla}T = 0$.

In particle simulations, these boundary conditions are imposed by employing standard rules for particle reflection at the boundaries [48]. We describe the corresponding handling in continuum simulations in Appendix C. In Appendix B we derive the form of the additional peaks in the dynamic structure factor for adiabatic walls by solving the linearized LLNS equations with the appropriate conditions.

In Figure 5 we show dynamic structure factors for a quasi one-dimensional system bounded by adiabatic walls. As for the bulk (periodic) case in the previous section, we perform both purely particle runs and also hybrid runs in which the middle third of the domain is designated as a particle subdomain. Additional peaks due to the reflections of sound waves from the boundaries are clearly visible and correctly predicted by the LLNS equations and also accurately reproduced by both the purely continuum solution (not shown) and the hybrid. Similar agreement (not shown) is obtained among the particle, the continuum, and the hybrid runs for thermal walls. These results show that the hybrid is capable of capturing the dynamics of the fluctuations even in the presence of boundaries. Note that when the deterministic hybrid scheme is used, one obtains essentially the correct shape of the peaks in the structure factor (not shown); however, the magnitude is smaller (by a factor of 2.5 for the example in Figure 5) than the correct value due to the reduced level of fluctuations.

4.3. Bead VACF. As an illustration of the correct hydrodynamic behavior of the hybrid algorithm, we study the velocity autocorrelation function (VACF) $C(t) = \langle v_x(0)v_x(t)\rangle$ for a large *neutrally buoyant* impermeable *bead* of mass M and radius R diffusing through a dense Maxwell I-DSMC stochastic fluid [18] of particles with mass $m \ll M$, collision diameter $D \ll R$, and density (volume fraction) ϕ (mass density $\rho = 6m\phi/(\pi D^3)$). The VACF problem is relevant to the modeling of polymer chains or (nano)colloids in solution (i.e., *complex fluids*), in particular, the integral of the VACF determines the diffusion coefficient, which is an important macroscopic quantity. Furthermore, the very first MD studies of the VACF for fluid molecules led to the discovery of a long power-law tail in $C(t)$ [68] which has since become a standard test for hydrodynamic behavior of methods for complex fluids [69, 70, 71, 72, 73, 74, 75].

The fluctuation-dissipation principle [76] points out that $C(t)$ is exactly the decaying speed of a bead that initially has a unit speed, if only viscous dissipation is present without fluctuations, and the equipartition principle tells us that $C(0) = \langle v_x^2 \rangle = kT/2M$. Using these two observations and assuming that the dissipation is well described by a continuum approximation with stick boundary conditions on a sphere of radius R_H, $C(t)$ has been calculated from the linearized (compressible) Navier–Stokes equations [77, 78]. The results are analytically complex even in the Laplace domain; however, at short times an inviscid compressible approximation applies. At large times the compressibility does not play a role and the incompressible Navier–Stokes equations can be used to predict the long-time tail [78, 79]. At short times, $t < t_c = 2R_H/c_s$, the major effect of compressibility is that sound waves
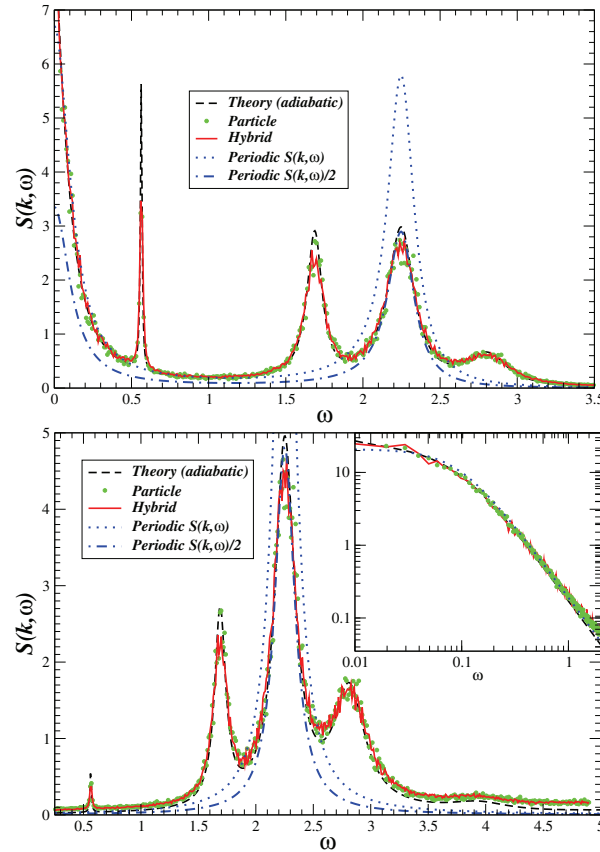
FIG. 5. *Discrete dynamic structure factors for a quasi one-dimensional system with length $L = 7.2$ (corresponding to 36 continuum cells) bounded by adiabatic walls, for waveindex $q = 2$ and wavevector $k = 2\pi q/L \approx 1.75$. The I-DSMC fluid parameters are as in Figure 4. We perform purely particle runs and also hybrid runs in which the middle third of the domain is filled with particles and the rest handled with a continuum solver. The results from purely particle runs are shown with symbols, while the results from the hybrid are shown with a solid line. The theoretical predictions based on the linearized LLNS equations and the equations in Appendix B (solved using the computer algebra system Maple) are shown with a dashed line for adiabatic boundaries and a dotted line for periodic boundaries. Since the magnitude of the Brillouin peaks shrinks to one-half the bulk value in the presence of adiabatic walls, we also show the result with periodic boundaries scaled by $1/2$ (dashed-dotted line). (Top) Dynamic structure factor for density, $S_\rho(k, \omega)$, showing the Rayleigh peak and the multiple Brillouin peaks. (Bottom) Dynamic structure factor for the component of velocity perpendicular to the wall, $S_{v_\perp}(k, \omega)$, which lacks the Rayleigh peak. The corresponding correlations for either of the parallel velocity components, $S_{v_\parallel}(k, \omega)$, have only a Rayleigh peak, shown in the inset on a log-log scale.*

generated by the motion of the suspended particle carry away a fraction of the momentum with the sound speed c_s, so that the VACF quickly decays from its initial value $C(0) = k_B T/M$ to $C(t_c) \approx k_B T/M_{eff}$, where $M_{eff} = M + 2\pi R^3 \rho/3$ [78]. At long times, $t > t_{visc} = 4\rho R_H^2/3\eta$, the VACF decays as with an asymptotic power-law tail $(k_B T/M)(8\sqrt{3\pi})^{-1}(t/t_{visc})^{-3/2}$, in disagreement with predictions based on the Langevin equation (Brownian dynamics), $C(t) = (k_B T/M) \exp(-6\pi R_H \eta t/M)$.

We performed purely particle simulations of a diffusing bead in various I-DSMC fluids in [18, 75]. In purely particle methods the length of the runs necessary to achieve

sufficient accuracy in the region of the hydrodynamic tail is often prohibitively large for beads much larger than the fluid particles themselves. It is necessary to use hybrid methods and limit the particle region to the vicinity of the bead in order to achieve a sufficient separation of the molecular, sonic, viscous, and diffusive time scales and study sufficiently large box sizes over sufficiently long times. In the results we report here we have used an impermeable hard bead for easier comparison with existing theory; similar results are obtained using permeable beads. The interaction between the I-DSMC fluid particles and the bead is treated as if the fluid particles are hard spheres of diameter D_s, chosen to be somewhat smaller than their interaction diameter with other fluid particles (specifically, we use $D_s = D/4$) for computational efficiency reasons, using an event-driven algorithm [48]. Upon collision with the bead, the relative velocity of the fluid particle is reversed in order to provide a no-slip condition at the surface of the suspended sphere [48, 72] (slip boundaries give qualitatively identical results). We have estimated the effective (hydrodynamic) colloid radius R_H from numerical measurements of the Stokes friction force $F = -6\pi R_H \eta v$ and found it to be somewhat larger than the hard-core collision radius $R + D_s/2$, but for the calculations below we use $R_H = R + D_s/2$. Since we used periodic boundary conditions with a box of length L, there are artifacts in $C(t)$ after about the time at which sound waves generated by its periodic images reach the particle, $t_L = L/c_s$. The averaging procedure that we used in order to eliminate some of the noise from the tail does not properly resolve these sound effects,[2] even though they are visible in the results shown below. Furthermore, there are strong finite-size effects that manifest themselves as a rapid decay of the VACF for times longer than the viscous time scale $\rho L^2/\eta$ [73].

For the hybrid calculations, we localize the particle subdomain to the continuum cells that overlap or are close to the moving bead. The location of the particle subdomain is updated periodically as the bead moves. The algorithm that we use tries to fit the particle subdomain as closely around the bead but ensuring that there is a certain number of micro cells between the surface of the bead and the particle-continuum interface. The exact shape of the particle subdomain thus changes as the bead moves and the number of particles employed by the hybrid fluctuates, especially when the bead is small compared to the continuum cells. Obtaining reasonably accurate results for the VACF at long times requires very long runs. We found that it is crucial to strictly conserve momentum in the hybrid when unfilled continuum cells are transformed into particle cells. Otherwise very slow drifts in the momentum of the system appear due to the use of periodic boundary conditions, and this drift changes the tail of $C(t)$, especially for massive beads where the typical bead speed is already small compared to the typical fluid particle speed. We used the MacCormack solver and the linearized formulation of the LLNS equations for these simulations; however, similar results are obtained with the nonlinear Runge–Kutta solver as long as the macro cells are sufficiently large. As discussed earlier, the use of the linearized formulation makes it possible to reduce the size of the continuum cells without introducing numerical problems due to the nonlinearity of the equations.

In the calculations reported in Figure 6, the I-DSMC fluid has a density (volume

---

[2]To calculate the VACF, we first calculated the average mean-square displacement $\Delta r(t)$ by averaging over a long run and numerically differentiated this to obtain a time-dependent diffusion coefficient $D(t)$. We then smoothed $D(t)$ by fitting a quadratic polynomial over short-time intervals spaced on a logarithmic scale (so that more points were averaged over in the tail), and obtained the VACF by differentiating the smoothed $D(t)$. This procedure produces well-resolved tails; however, it obscures features over short time scales compared to the smoothing interval. A more direct velocity autocorrelation calculation was used at very short times.
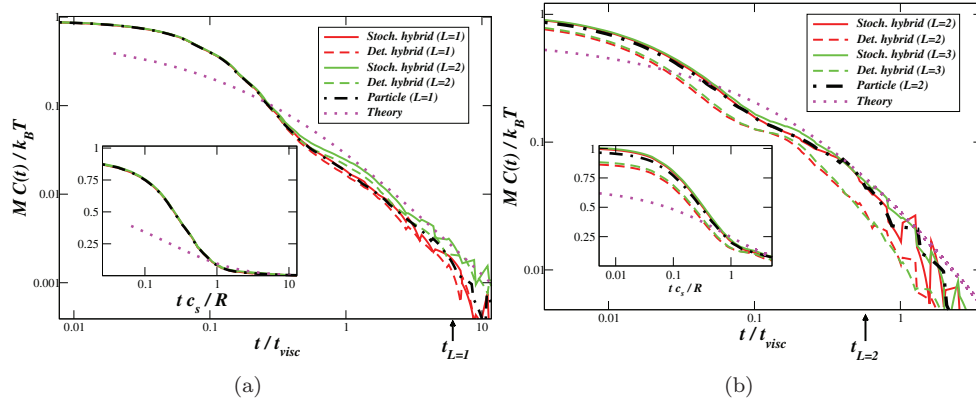
FIG. 6. *Normalized VACF $C(t)/(k_B T/M)$ for a neutrally buoyant hard spherical bead of mass M suspended in a fluid of I-DSMC particles of diameter D, for two different bead sizes, a small bead of radius $R = 1.25D$ (a) and a large bead of radius $R = 6.25D$ (b). A log-log scale is used to emphasize the long-time power law tail, and the time is normalized by the viscous time t_{visc}, so that the results should be approximately independent of the actual bead radius. The inset shows the initial decay of the VACF on a semi-log scale, where the time is normalized by the sonic time scale t_c. Periodic boundary conditions with a cubic cell of length L are employed, and the sound crossing time t_L is indicated. Results from purely particle simulations are shown with a dashed-dotted line, and the incompressible hydrodynamic theory is shown in with a dotted line. Results from hybrid runs are also shown with a solid line for the stochastic hybrid and dashed line for the deterministic hybrid, for both the same box size as the particle runs (red) and a larger simulation box (green).*

fraction) $\phi = 0.5$ and collision frequency prefactor $\chi = 0.62$. The adiabatic sound speed is $c_s = \sqrt{5k_B T/(3m)}$ and viscosity is $\eta = \tilde{\eta} D^{-2} \sqrt{m k_B T}$, where we measured $\tilde{\eta} \approx 0.75$. Note that in atomistic time units $t_0 = D\sqrt{m/k_B T}$ the viscous time scale is $t_{visc}/t_0 \approx 6\phi(R_H/D)^2/(3\pi\tilde{\eta}) \approx 0.4(R_H/D)^2$.

As a first test case, in Figure 6(a) we compare against the particle data from [75], for which the size of the bead is $R = 1.25D$, $M = 7.81m$, and the simulation box is $L = 1 = 25D$, which corresponds to 24^3 micro cells and about $N \approx 1.5 \cdot 10^4$ particles. The hybrid runs used macro cells each composed of 4^3 micro cells, which corresponds to about $N_0 = 80$ particles per cell, and the size of the particle subdomain fluctuated between about $3 \cdot 10^3$ and $6 \cdot 10^3$ particles due to the change of the location of the bead relative to the continuum grid. The particle result in Figure 6(a) is the average over 10 runs, each of length $T/t_{visc} \approx 2 \cdot 10^5$, while the hybrid results are from a single run of length $T/t_{visc} \approx 7.5 \cdot 10^5$. It is seen in the figure that both the deterministic and the fluctuating hybrids reproduce the particle results closely, with a small but visible difference at long times where the deterministic hybrid underpredicts the magnitude of the tail in the VACF. We also show results from a hybrid run with a twice larger simulation box, $L = 2$, which marginally increases the computational effort in the hybrid runs, but would increase the length of purely particle runs by an order of magnitude. The hydrodynamic tail becomes pronounced and closer to the theoretical prediction for an infinite system, as expected. We have observed (not shown) that using small continuum cells composed of only 3^3 micro cells, which corresponds to about $N_0 = 35$ particles per cell, leads to an overprediction of the magnitude of the VACF at short times, that is, to an excess kinetic energy for the bead by as much as 20%, depending on the exact parameters used.

As a second, more difficult test, in Figure 6(b) we report results from particle

simulations for a much larger bead, $R = 6.25D$, $M = 976m$, and the simulation box is $L = 2 = 50D$, which corresponds to $N \approx 1.2 \cdot 10^5$ particles. We have performed a variety of hybrid runs, and in Figure 6(b) we report results from runs with macro cells each composed of 3^3 micro cells, as well as results for a larger simulation box, $L = 3$, and macro cells each composed of 4^3 micro cells. The particle results are the average over 5 runs each of length $T/t_{visc} \approx 2.5 \cdot 10^3$, while the hybrid results are from a single run of length $T/t_{visc} \approx 7.5 \cdot 10^3$. The hybrid runs had a particle subdomain containing about $2 \cdot 10^4$ particles. We observed little impact of the size of the continuum cell size or the size of the particle subdomain. The results in Figure 6(b) show that the stochastic hybrid correctly reproduces the tail in the VACF, while it slightly overestimates the VACF at short times. The deterministic hybrid, on the other hand, strongly underestimates the magnitude of $C(t)$ at both short and long times. It is particularly striking that the deterministic hybrid fails to reproduce the magnitude of the long-time tail (and thus the diffusion coefficient), vividly demonstrating the importance of including fluctuations in the continuum domain.

Computing the VACF for a diffusing bead has become a standard test for micro- and nanoscale fluid-structure coupling methods and has been performed for a suspended bead in a wide range of particle and (semi)continuum compressible and incompressible fluids [68, 69, 70, 71, 72, 73, 74, 75]. However, these tests often do not correctly test for all of the components necessary to match the VACF over all relevant time scales: equipartition, a fluctuation-dissipation relation, and hydrodynamics. Purely continuum fluid methods allow for using a much larger time step than particle (and thus hybrid) methods, especially if an incompressible formulation is directly coupled to the equations of motion of the suspended bead [71, 73, 74]. When fluctuations are not included in continuum methods, the VACF is often obtained by considering the deterministic decay of the velocity of a bead. This, however, assumes a priori that proper thermodynamic equilibrium exists with the correct fluctuation-dissipation relation, without actually testing this explicitly. Alternatively, coupling methods between a continuum fluid and a suspended particle often have some arbitrary coupling parameters that are tuned to reproduce the desired diffusion coefficient without producing a physically consistent VACF, especially at short times [80]. In particular, incompressible formulations cannot reproduce the initial value or the decay of the VACF and should instead aim to produce an average kinetic energy of the bead of $k_B T$ rather than the statistical mechanics result of $3k_B T/2$ [79]. It is therefore important to more carefully access the ability of other methods in the literature to correctly reproduce the full VACF for a true equilibrium simulation with a large bead.

4.4. Adiabatic piston. The problem of how thermodynamic equilibrium is established has a long history in statistical mechanics [81]. The *adiabatic piston problem* is one of the examples used to study the fluctuation-driven relaxation toward equilibrium [82, 83, 84, 85] that is simple enough to be amenable to theoretical analysis but also sufficiently realistic to be relevant to important problems in nanoscience such as Brownian motors [42, 86]. We study the following formulation of the adiabatic piston problem. A long quasi one-dimensional box with adiabatic walls is divided into two halves with a thermally insulating piston of mass $M \gg m$ that can move along the length of the box without friction. Each of the two halves of the box is filled with a fluid that is, initially, at a different temperature T and density ρ, here assumed to follow the ideal EOS $P = \rho k_B T/m$. If the macroscopic pressures in the two halves are different, $\rho_L T_L \neq \rho_R T_R$, then the pressure difference will push the piston to perform macroscopic oscillations with a period that can be estimated by assuming that

each half undergoes an adiabatic transformation ($PV^{\gamma} = $ const.). These oscillations are damped by viscous friction and lead to the piston essentially coming to rest in a state of *mechanical equilibrium*, $\rho_L T_L = \rho_R T_R$. This stage of the relaxation from nonequilibrium to mechanical equilibrium has been shown to be well described by deterministic hydrodynamics [84].

The state of mechanical equilibrium is not, however, a state of true thermodynamic equilibrium, which also requires equality of the temperatures on the two sides of the piston. Reaching full equilibrium requires heat transfer through the piston, but the piston is adiabatic and does not conduct heat. In classical deterministic hydrodynamics the piston would just stand still and never reach full equilibrium. It has been realized long ago that heat is slowly transferred through the mechanical asymmetric fluctuations of the piston due to its thermal motion, until the temperatures on both sides of the piston equilibrate and the fluctuations become symmetric. This equilibration takes place through a single degree of freedom (the piston position) coupling the two large reservoirs, and it would be astronomically slow in a laboratory setting. While various Langevin or kinetic theories have been developed for the effective heat conduction of the adiabatic piston (see [82, 83, 84, 85] and the references therein), there is no complete theoretical understanding of the effective heat conductivity, especially in dense fluids. Molecular dynamic simulations have been performed in the past [82, 83, 84] using hard-disk fluids, but the very long runs required to reach thermodynamic equilibrium for massive pistons have limited the size of the systems that could be studied. Furthermore, there have been no studies applying fluctuating hydrodynamics to this problem.

Here we apply our hybrid method to the adiabatic piston problem in two dimensions, using a nonlinear two-dimensional implementation of the Runge–Kutta integrator described in [9] as the continuum solver. The choice of two dimensions is for purely computational reasons. First, the number of particles required to fill a box of sufficient size is much smaller, thus allowing for long particle simulations. Second, in order to implement the piston in our particle scheme we reused the same mixed event-driven/time-driven handling [48] as we used for the VACF computations in section 4.3. Namely, we made a piston out of N_b small impermeable beads, connected together to form a barrier between the two box halves, as illustrated in Figure 7. In two dimensions, by ensuring that two piston beads never separate by more than a given distance we can ensure that two I-DSMC particles on opposite sides of the piston cannot possibly collide, and thus the piston will be insulating. We have studied two different types of pistons, a *flexible* piston where the the beads are tethered together to form a chain [48] that is stretched but where each individual bead can still move independently of the others, and a *rigid* piston that is obtained with a slight modification of the event loop to move all of the piston beads in unison. While at the macroscopic level the exact shape of the piston should not make a big difference, we have found that increasing the number of degrees of freedom of the piston from one to N_b makes a significant difference in the thermal conductivity of the piston, and therefore we will focus here on rigid pistons as in the traditional formulation. We use specular collisions of the fluid particles with the piston beads, although qualitatively identical (but not quantitatively identical) results are obtained using bounce-back collisions as well.

The hybrid method setup for the adiabatic piston is illustrated in Figure 7. We use a two-dimensional Maxwell I-DSMC particle fluid ($\phi = 1$, $\chi = 1$) with collision diameter $D = 0.1$ (hard-sphere diameter $D_s = D/2$) and a piston composed of $N_b = 40$ beads of diameter $D_b = 0.0955$. The particle subdomain is limited to a few
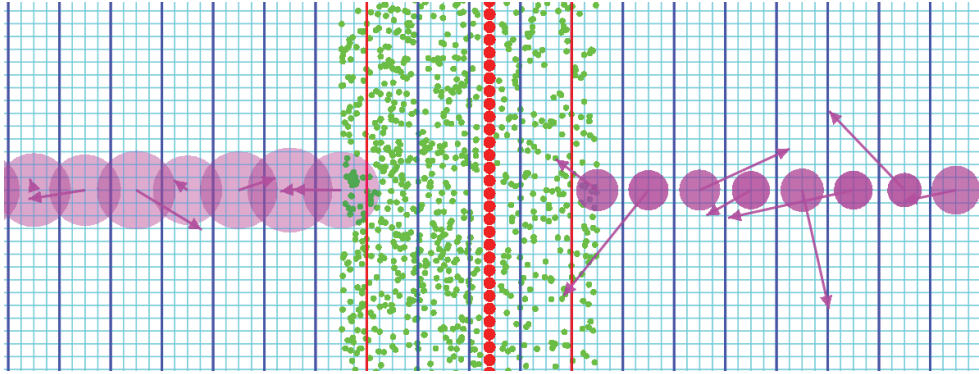
FIG. 7. *An illustration of the computational setup used for the adiabatic piston computations. Only the central portion of the box of aspect ratio 6×1 is shown. Left of the piston the gas is cold and dense; to the right it is hot and dilute. The piston beads (centered column of red disks) separate the box into two halves, and are surrounded on each side by a fluid of I-DSMC particles (smaller green disks), which is twice denser but also twice cooler in the left half than in the right half. The microscopic grid is shown with thinner light lines and the hydrodynamic grid is shown with thicker dark blue lines. The interface between the particle and continuum regions is highlighted with a thick red line. A snapshot of the values of the hydrodynamic variables in each continuum cell is shown using a larger purple disk whose size is proportional to the density and whose opaqueness is proportional to the temperature, with an arrow for the fluctuating velocities.*

continuum cells around the piston, which we keep at about two or more continuum cells on each side of the piston, so that the unreasonable hydrodynamic values in the cells that overlap the piston do not affect the continuum solver appreciably. Periodic boundary conditions are applied along the y dimension (parallel to the piston) with the width of the domain $L_y = 4$ being 40 microscopic cells, while adiabatic walls were placed at the ends of the box whose total length $L_x = 24$ was 240 microscopic cells. We have studied various sizes for the macroscopic cells and report results for a quasi one-dimensional continuum grid in which each macro cell contains 4×40 micro cells, corresponding to about 200 particles per continuum cell. We also present results for a two-dimensional continuum grid where each macro cell contains 8×10 micro cells.

We have performed hybrid runs with both the deterministic and stochastic hybrids. In Figure 8 we show the position of a massive piston of mass $M = 4000m$ that started at a position $x = 8$ that is not in mechanical equilibrium and thus performs rapid damped oscillations until it reaches a state of mechanical equilibrium at $x \approx 7$, from which it slowly relaxes toward true equilibrium. The results in the figure show that the stochastic hybrid reproduces the correct relaxation toward equilibrium, while the deterministic hybrid severely underpredicted the rate of equilibration (effective heat conductivity), even though the initial mechanical stage of the relaxation is correctly captured by both hybrids, as expected. We have observed that the deterministic hybrid fails to give the correct answer whenever a rigid massive piston is used, $M > 250m$. For flexible pistons, we find that even for a large bead mass M_b (overall piston mass $M = N_b M_b$) both the deterministic and stochastic hybrids reproduce the purely particle results for the slow relaxation toward equilibrium.

A more detailed comparison of the particle and hybrid results for a piston of mass $M = 1000m$ that is initially in mechanical equilibrium at position $x = 8 = L_x/3$ is shown in Figure 9. The initial conditions were $k_B T_L = 2/3$, $\rho_L = 2/3$ and $k_B T_R = 4/3$, $\rho_L = 1/3$, so that there is an equal mass on each side of the piston. At the true
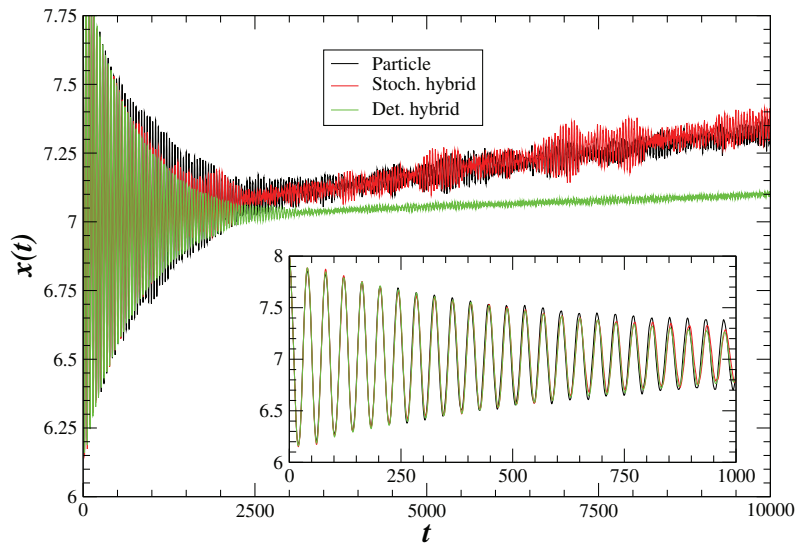
FIG. 8. *Relaxation of a massive rigid piston ($M/m = 4000$) from a position $x = 8$ that is not in mechanical equilibrium. Through rapid damped oscillations, mechanical equilibrium is established at position $x \approx 7$, after which a slow relaxation to true equilibrium is seen. The stochastic hybrid is able to match the particle data very well, to within the expected statistical difference from realization to realization. The deterministic hybrid, on the other hand, clearly underpredicts the rate of relaxation. The inset highlights the initial oscillations and shows that in the regime where fluctuations do not matter the deterministic and stochastic hybrids do not differ appreciably.*
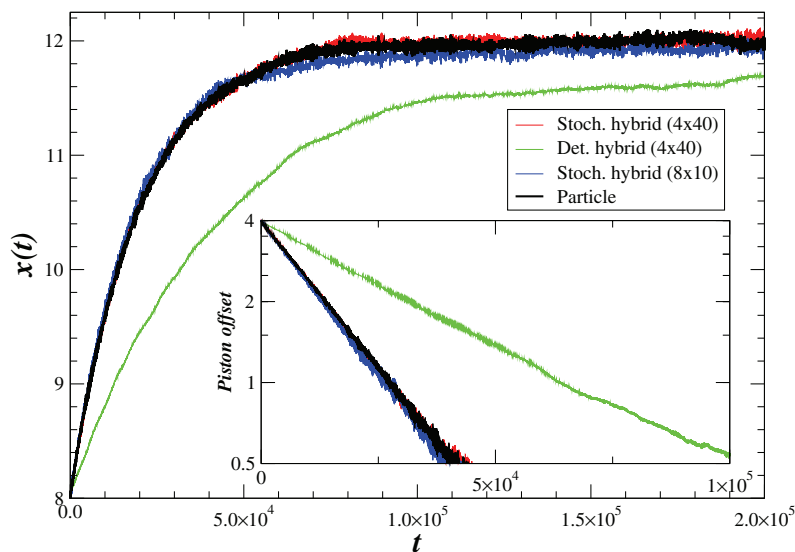

FIG. 9. *Relaxation of a rigid piston of mass $M/m = 1000$ from an initial state of mechanical equilibrium ($x = 8$) to a state of thermodynamic equilibrium ($x = 12$). The inset emphasizes the initial exponential decay on a semi-log scale. The hybrid runs used a particle subdomain of width $w_P = 2$ on each side of the piston and continuum cells that were composed of either 4×40 or 8×10 microscopic cells. For the deterministic hybrid the macro cell size makes little difference so we only show the 4×40 case.*
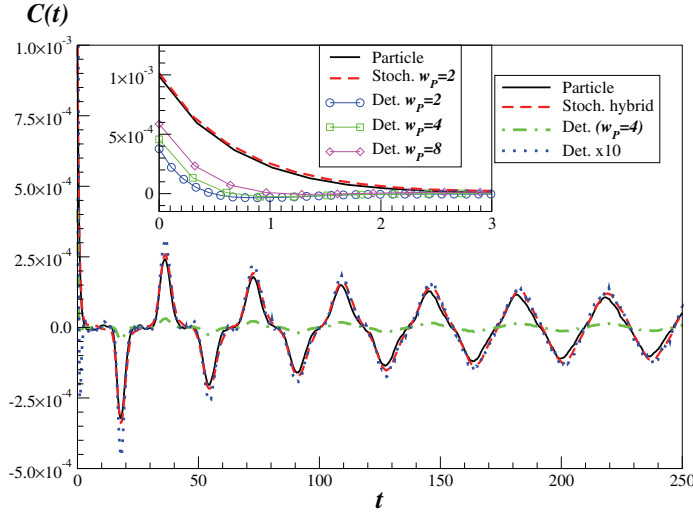
FIG. 10. *The VACF for a rigid piston of mass $M/m = 1000$ at thermal equilibrium at $k_B T = 1$ (thus, the expected initial value is $C(0) = 10^{-3}$) and macro cells of size 4×40 micro cells. The deterministic hybrid gives much smaller effective temperature T_{eff} of the piston and a negative dip in $C(t)$ at short times, but when magnified by an order of magnitude it reveals the correct shape at longer times. The inset focuses on the initial decay of the VACF and shows that even increasing the width of the particle region to $w_P = 8$ (the whole box has a length of $w_L = 60$ continuum cells) does not help the deterministic hybrid much, whereas the stochastic hybrid gives the correct initial decay.*

equilibrium state the piston remains close to the middle of the box, $x_{eq} = L/2 = 12$, with equal density on each side. The results shown are averages over 10 samples, but it should be emphasized that each run exhibits thermal oscillations of the piston[3] position that are diminished by direct averaging because they have different (random) phases. Figure 9 shows that the stochastic hybrid is able to correctly reproduce the rate of exponential relaxation of the piston toward equilibrium with many fewer particles than the purely particle runs, while the deterministic hybrid fails. We have observed a slight dependence on the exact details of the hybrid calculations such as cell size or the width of the particle subdomain; however, in general, the stochastic hybrid has shown to be remarkably robust and successful. At the same time, the importance of including thermal fluctuations in the continuum subdomain is revealed as for the VACF computations in section 4.3.

The relation to the equilibrium VACF computations in section 4.3 is emphasized by computing the VACF $C(t) = \langle v_x(t) v_x(0) \rangle$ for the piston in its state of true equilibrium,[4] $k_B T = 1$. The VACF is rather complex due to the interplay of short-time kinetic effects, dissipation, and sound reflections from the walls; however, our focus here is to simply compare it against the purely particle simulations and not to understand all the features in the VACF. The results, shown in Figure 10, reveal that the

---

[3]At thermodynamic equilibrium with a common temperature T, the frequency of the thermally driven oscillations can be estimated using a quasi-adiabatic harmonic approximation to be $\omega^2 \approx N k_B T / [M x (L_x - x)]$ (see also an alternative derivation in [83]), and the amplitude of the oscillations can be estimated to be on the order of $\Delta x^2 \approx x (L_x - x)/N$, where N is the number of fluid particles per chamber.

[4]In [83] the autocorrelation function for the piston bead position is used to extract the rate of exponential rate toward true equilibrium; however, a direct nonequilibrium calculation such as we perform in Figure 9 is more efficient and illustrative for our purpose.

piston does not equilibrate at the correct effective temperature $T_{\text{eff}} = MC(0)/k_B$ in the deterministic hybrid calculations. Notably, just as we found for the massive bead in section 4.3, the piston has a kinetic energy that is markedly lower (half) than the correct value $C(0) = k_B T/M = 10^{-3}$ when fluctuations are not consistently included in the continuum region. Since the quasi-equilibrium temperature of the piston[5] plays a crucial role in all of the kinetic theories [82, 83, 85] for the effective heat transfer, it is not surprising that the deterministic hybrid gives the wrong answer. What is even more striking is that increasing the width of the particle subdomain w_P on each side of the piston, as measured in units of continuum cells (Figure 7 shows the case $w_P = 2$), barely improves the accuracy of the VACF at short times, showing that the whole spectrum of fluctuations affects the initial rate of decay of the piston. At the same time, the deterministic hybrid does give the correct shape of the VACF at longer times, as revealed by magnifying the VACF for $w_P = 4$ by an arbitrary factor of 10 to bring it in close agreement with the particle result at longer times.[6] This is also not unexpected since the VACF at longer times is dominated by mechanical vibrations and dissipation, and is analogous to what we find for the dynamic structure factor when it is calculated by a deterministic hybrid scheme.

5. Conclusions. We have described a hybrid particle-continuum algorithm for simulating complex flows and applied it to several nontrivial problems. The algorithm couples an ideal stochastic particle fluid algorithm with a fluctuating hydrodynamic continuum solver using a direct dynamic coupling where the continuum solver supplies Dirichlet-like (state) boundary conditions for the particle region, while the particle region supplies von Neumann-like (flux) boundary conditions to the continuum solver. The continuum solver computes a provisional solution over the whole domain that then gets replaced with the particle data in the particle region and also gets corrected (refluxed) in the cells bordering the particle subdomain. We described the components necessary to extend previous variants of this rather general coupling methodology [26, 36, 47] to use a recently developed variant of the DSMC particle method suitable for dense fluids [18, 75], as well as an explicit conservative compressible fluid solver that accurately accounts for thermal fluctuations in the Navier–Stokes equations [9]. By turning the fluctuating fluxes off in the continuum solver we can trivially transform our stochastic hybrid method to a deterministic hybrid method closer to commonly used hybrid schemes.

In section 4 we applied our stochastic hybrid method to several challenging problems and demonstrated that it could obtain the correct answer with significantly fewer particles and thus significantly less computational effort than a purely particle simulation. We used purely particle runs as a baseline standard against which we judge the accuracy of the hybrid method, consistent with using the hybrid method for situations in which a purely continuum description is unable to capture the full physics and a particle method is necessary in some portion of the physical domain. For example, particle methods such as DSMC are essential to correctly resolve the flow at a high Mach number shock [36], or in kinetic flow regions such as the wake behind a fast-moving body [26] or small channels in a microelectromechanical device [5].

In section 4.1 we showed that there is a small mismatch between the density and

---

[5]It is predicted that the piston equilibrates at a temperature that is approximately the geometric mean of the left and right temperatures [87].

[6]The magnification factor required to bring the long-time VACF for the deterministic hybrid in agreement with the correct result decreases with increasing w_P; for example, it is ~ 20 for $w_P = 2$ and ~ 4 for $w_P = 8$.

temperature in the particle and continuum regions, caused by using instantaneous fluctuating values of the hydrodynamic variables when generating velocities for the reservoir particles, and that the effect is of order N_0^{-1}, where N_0 is the average number of particles in one cell. In Appendix A we demonstrated that this is an artifact not of the method but rather of an inherent difficulty in stochastic methods where instantaneous values are used to estimate means. In section 4.2 we calculated the dynamic structure factor using the hybrid method for a quasi two-dimensional situation and a large wavevector $\boldsymbol{k}$ that is obliquely incident to the particle-continuum interface, and we confirmed that spontaneous sound and entropic fluctuations are transmitted correctly through the interface. We also calculated the dynamic structure factor for a finite system bounded by two adiabatic walls and observed excellent agreement with theoretical calculations given in Appendix B.

In section 4.3 we studied the diffusive motion of a large spherical neutrally buoyant bead suspended in a fluid of I-DSMC particles in three dimensions by placing a mobile particle subdomain around the suspended bead. We computed the velocity autocorrelation function (VACF) for two bead sizes using the stochastic and deterministic hybrids as well as purely particle simulations. We found that the stochastic hybrid correctly reproduces the VACF over all time scales, while the deterministic hybrid underestimates both the kinetic energy of the bead and the magnitude of the tail of the VACF for sufficiently large beads. Finally, in section 4.4 we applied the hybrid scheme to a nonequilibrium quasi one-dimensional version of the *adiabatic piston* problem [81, 82], a classic example of the importance of fluctuations in establishing global thermal equilibrium. The two-dimensional particle region was placed around the piston and an event-driven algorithm was used to handle the interaction of the fluid with the piston. We again found that the stochastic hybrid was able to reproduce the purely particle results correctly, while the deterministic calculations underestimated the relaxation substantially for sufficiently massive rigid pistons.

Our results for both the VACF and the adiabatic piston clearly demonstrated that a large massive suspended body cannot equilibrate at the correct Boltzmann distribution unless thermal fluctuations are consistently included in the full domain, including the continuum region in hybrid methods, even if a large particle subdomain is used. This points to an increased importance of the long-wavelength—and thus slowly decaying—hydrodynamic fluctuations. Massive and large suspended bodies have longer relaxation times, and thus it is not surprising that slower-decaying fluctuations play a more prominent role for them than for smaller suspended particles. However, a better theoretical understanding of these observations is necessary in order to establish the importance of hydrodynamic fluctuations in general. At the same time, our results make it clear that fluctuations should be included in the continuum region in hybrid methods, consistent with the particle dynamics, rather than treating the fluctuations as "noise" from which the continuum solver ought to be shielded.

Fluctuating hydrodynamics has successfully accounted for thermal fluctuations in a variety of problems. At the same time, however, the nonlinear stochastic partial differential Landau–Lifshitz Navier–Stokes (LLNS) equations are mathematically ill-defined and require a cutoff length and/or time scale to be interpreted in a reasonable sense. The linearized equations can be given a well-defined interpretation; however, they are physically unsatisfactory in that they require a base state to linearize around, which may itself be time-dependent and unknown or even be affected by fluctuations, as in the adiabatic piston or diffusing shock problems [6, 88]. We have numerically observed that using small continuum cells leads to worse results even if the linearized

LLNS equations are used, thus formally avoiding the difficulties with the increased relative magnitude of the fluctuations. This is consistent with the expectation that a continuum description is applicable only on length scales and time scales sufficiently larger than the molecular size and molecular collision time. In our experience using more than $O(10^2)$ particles per cell leads to a good match between the hybrid and particle runs; however, a better theoretical understanding of the proper inclusion of fluctuations in hydrodynamics is a necessary future development.

An important and challenging task for future research is to understand the trade-off between accuracy and computational cost in hybrid algorithms. This is a nontrivial task because there are multiple sources of inaccuracy that are difficult to separate. First, assuming that the particle description is sufficiently accurate that it can be taken as the "correct" solution, the continuum description itself is an approximation whose region of applicability is only intuitively known. Even assuming that the continuum description is sufficiently accurate, the numerical scheme used to discretize the continuum model has its own spatio-temporal discretization errors. Furthermore, assuming that the discrete continuum model sufficiently matches the particle model in the region near the particle-continuum interface, there is some error associated with the coupling between the two descriptions. Finally, for any given problem it is nontrivial to define what is meant by "accuracy," and the very notion depends on what quantity is of physical interest and how it is being measured, especially when fluctuations are taken into account. Even if one properly defines an error estimate, actually measuring the error is often computationally very expensive due to the long runs required to eliminate statistical noise (for example, to obtain the long-time tail in the VACF), especially in runs that use larger particle subdomains. In this work, we have not attempted to properly quantify the accuracy of the hybrid. Rather, we have used the smallest particle subdomains possible in order to make our implementation of the hybrid algorithm function. Even with such small particle domains we see very good agreement (close to or within statistical accuracy) between the hybrid and purely particle runs for the quantities of physical interest, at least for the VACF and adiabatic piston problems that we studied.

Typically, even for the rather simplified particle model employed in the I-DSMC algorithm, the cost in a hybrid run is dominated by the particle calculations. For the piston problem and runs with different sizes of the particle subdomain shown in the inset of Figure 10, a rough estimate of the total computational cost *per particle time step* (which is assumed to be kept constant) is

$$\text{cost} \approx c_P N_0 (2w_P + 1) + c_H \frac{w_L}{n_{ex}} + \text{coupling and other costs,}$$

where c_P is the cost of updating one particle for one particle time step while c_H is the cost of updating a continuum cell for one time step, N_0 is the average number of particles per continuum cell, w_P is the width (in continuum cells) of the particle subdomain on a given side of the piston (we have added one as an estimate of the width of the reservoir region), n_{ex} is the number of particle time steps per continuum step, and w_L is the length (in continuum cells) of the full simulation domain. For our implementation and the parameters used to obtain the results in Figure 10, we empirically estimated $c_H/c_P \approx 5$. This means that $c_H/n_{ex} \approx c_P$, that is, the cost of updating a hydro cell is close to the cost of updating a single particle, and therefore the speedup obtained by switching to a continuum description is on the order of the number of particles per cell.

Our implementation is at present serial and our runs are therefore limited by the

CPU-intensive collision procedure in the I-DSMC particle algorithm. Some of the examples we presented utilized the mixed event- and time-driven particle algorithm developed in [48]. The event-driven component of this algorithm is notoriously difficult to parallelize. However, a purely time-driven particle algorithm can easily be parallelized, as can the purely continuum solver. We plan to implement a parallel hybrid scheme in the future in order to enable the study of realistic system sizes. At the same time, however, reaching long time scales will necessarily require time steps beyond the small ones required by particle methods and explicit continuum schemes.

Appendix A. Finite-size mismatch between particle and continuum descriptions. Consider a simple particle/continuum hybrid consisting of a one-dimensional system with an ideal particle fluid on one side and a fluctuating continuum solver on the other. At equilibrium the mean density and temperature are ρ_0 and T_0, respectively; the mean fluid velocity is taken as zero. The problem we want to consider is whether the one-sided fluxes of mass, momentum, and energy due to the reservoir particles that enter the continuum subdomain are correct, assuming that the means and variances of the hydrodynamic variables in the reservoir cells are correct.

In the continuum calculation, instantaneous density and temperature, ρ and T, are computed and used to generate particles for injection. The one-sided fluxes of mass, momentum, and energy are known from kinetic theory to be

$$F_\rho(\rho, T) = \sqrt{\frac{k}{2\pi m}}\, \rho T^{1/2},$$

$$F_j(\rho, T) = \frac{k}{2m}\, \rho T,$$

$$F_E(\rho, T) = \sqrt{\frac{4k^3}{m^3\pi}}\, \rho T^{3/2}.$$

The correct mean equilibrium mass flux is $F_\rho(\rho_0, T_0)$; however, the mean value of the particle flux does not equal this correct value since

$$\langle \rho T^{1/2} \rangle = \langle \rho \rangle \langle T^{1/2} \rangle \neq \langle \rho \rangle \langle T \rangle^{1/2};$$

note that density and temperature are instantaneously uncorrelated. Similarly, the mean energy flux is also incorrect; interestingly the momentum flux is correct since it has a linear dependence on temperature.

To find the errors in the fluxes, we write $T = T_0 + \delta T$, and for a given power exponent a we have

$$\langle T^a \rangle = \langle T_0^a \rangle \langle (1 + \delta T/T_0)^a \rangle \approx \langle T_0^a \rangle \left[1 + \frac{1}{2} a(a-1) \frac{\langle \delta T^2 \rangle}{T_0^2} \right] = \langle T_0^a \rangle \left[1 + \frac{a(a-1)(\gamma-1)}{2N_0} \right],$$

where $N_0 = \rho_0 V_c/m$ is the number of particles in a continuum cell (of volume V_c). From this we have that, for a monatomic gas ($\gamma = 5/3$),

$$\langle F_\rho(\rho, T) \rangle = F_\rho(\rho_0, T_0) \left(1 - \frac{1}{12 N_0} \right),$$

$$\langle F_J(\rho, T) \rangle = F_J(\rho_0, T_0),$$

$$\langle F_E(\rho, T) \rangle = F_E(\rho_0, T_0) \left(1 + \frac{1}{4 N_0} \right).$$

This result shows that there is no way for all three mean fluxes to be correct when the particles are generated from the Maxwell–Boltzmann distribution using instantaneous, fluctuating values of density and temperature. The mismatch is of the order N_0^{-1}, where N_0 is the number of particles per macro cell, and therefore the mismatch gets worse as the macroscopic cells become smaller and the (relative) fluctuations become larger. As our numerical results show, because of this mismatch between the particle and hydrodynamic descriptions it will be impossible for the particle and continuum regions to reach a common equilibrium state.

Appendix B. Dynamic structure factor with adiabatic walls. Dynamic structure factors are easily calculated in the bulk by using the spatio-temporal Fourier transform of the linearized LLNS equations [67]. For finite domains, such as slab geometries, there are results in the literature, but they are often restricted to some simplified models or complex nonequilibrium situations [89, 90]. We therefore derive here the equilibrium dynamic structure factors for a fluid between two adiabatic hard walls by solving the LLNS equations with the appropriate boundary conditions.

Boundary conditions change the Hilbert space in which a solution is to be sought and the corresponding basis functions (eigenfunctions of the generator with the specified boundary conditions). For the LLNS equations with adiabatic boundaries (i.e., slip insulating walls) along the x axes, the appropriate basis functions are $\cos(kx)$ for density and temperature, where $k = p\pi/L$ is the wavevector enumerated by the waveindex $p \in \mathbb{Z}^+$, and $\sin(kx)$ for the velocities [89, 90], as compared to the ones for "bulk" conditions (periodic boundaries), $\exp(2q\pi x/L)$, with $k = 2q\pi/L$ and waveindex $q \in \mathbb{Z}$ [9, 67]. For thermal walls (constant-temperature stick boundaries) there does not appear to be a simple basis.

White noise has a trivial expansion in either the sine or cosine basis sets, namely, all of the coefficients are independently and identically distributed Gaussian random variables with mean zero and variance 2. The generator of the Navier–Stokes equations separates wavevectors/frequencies into the same (k, ω) equations as for "bulk" (periodic BCs), and therefore the dynamic structure factor, if expressed in the given basis set, has the same familiar form [89, 90]. In particular,

$$\rho(x,t) = \rho_0 + \sum_{p=1}^{\infty} \rho_p(t) \cos(p\pi x/L) = \rho_0 + \sum_{p=1}^{\infty} \int_{-\infty}^{\infty} d\omega e^{-i\omega t} \rho_{p,\omega} \cos(p\pi x/L),$$

where the different p's and ω's are uncorrelated,

$$\left\langle \rho_{p,\omega} \rho_{p',\omega'}^{\star} \right\rangle = 2\tilde{S}_{p,\omega} \delta_{p,p'} \delta\left(\omega - \omega'\right),$$

where $\tilde{S}_{p,\omega} = \tilde{S}_{k=p\pi/L,\omega}$ denotes the usual bulk dynamic structure factor.

Now, we need to convert this result to the more usual Fourier basis, $\exp(2q\pi x/L)$, since this is how dynamic structure factors are defined. From the Fourier inversion formula, and the orthogonality of the cosine basis functions, we have

$$\rho_q = \frac{1}{L} \int_0^L dx \rho(x,t) e^{-2q\pi x/L} = \frac{1}{2}\rho_{p=2q} - \frac{i}{L} \sum_{p=1}^{\infty} \rho_p \int_0^L dx \cos(p\pi x/L) \sin(2q\pi x/L).$$

By performing the integration explicitly we get

$$\rho_q = \frac{1}{2}\rho_{p=2q} + \frac{4iq}{\pi} \sum_{p \text{ odd}} \frac{\rho_p}{p^2 - 4q^2},$$

giving the dynamic structure factor

$$\text{(B1)} \qquad S_{q,\omega}^{(\rho)} \equiv S_{k=2\pi q/L,\omega}^{(\rho)} = \langle \rho_{q,\omega} \rho_{q,\omega}^{\star} \rangle = \frac{1}{2} \tilde{S}_{k,\omega}^{(\rho)} + \frac{32 q^2}{\pi^2} \sum_{p \text{ odd}} \frac{\tilde{S}_{p,\omega}^{(\rho)}}{\left(p^2 - 4q^2 \right)^2}.$$

Each of the terms under the sum gives an additional peak at $\omega = ck = pc\pi/L$, where p is odd, which arises physically because of the standing waves that appear due to reflections of the sound waves from the walls. Only the first few terms need to be kept to get most of the power (the total integral over ω is one), and it can be shown (by simple numerical comparison or explicit summation) that the above is equivalent to the more opaque equation (12) in [90] (set $\gamma = 0$ for no shear).

Equations identical to (B1) hold for temperature and the velocity components parallel to the wall. For the perpendicular velocity (v_x) a sine basis is more appropriate, and a similar calculation gives the dynamic structure factor for the finite system in terms of the bulk one,

$$S_{k=2\pi q/L,\omega}^{(v_\perp)} = \frac{1}{2} \tilde{S}_{k,\omega}^{(v_\perp)} + \frac{8}{\pi^2} \sum_{p \text{ odd}} \frac{p^2 \tilde{S}_{p,\omega}^{(v_\perp)}}{\left(p^2 - 4q^2 \right)^2}.$$

Appendix C. Handling of adiabatic and thermal walls in the continuum solver. Solving the LLNS equations with nonperiodic boundaries requires some special handling of the stochastic fluxes at the boundaries, which are assumed to coincide with faces of the continuum grid. As discussed in [9, 91], the numerical discretization of the Laplacian operator $\boldsymbol{L}$, the divergence operator $\boldsymbol{D}$, and the gradient operator $\boldsymbol{G}$ should satisfy a *discrete* fluctuation-dissipation balance condition $\boldsymbol{L} = \boldsymbol{DCG} = -\boldsymbol{DCD}^*$, where $\boldsymbol{C}$ is a dimensionless covariance matrix for the stochastic fluxes that are generated using a random number generator on each face of the grid. For one dimension with periodic boundaries it is well known that the standard face-to-cell divergence, cell-to-face gradient, and three-point Laplacian second-order stencils satisfy $\boldsymbol{L} = \boldsymbol{DG}$, and thus $\boldsymbol{C} = \boldsymbol{I}$ (the identity matrix) works and it is in fact trivial to implement algorithmically [9, 91]. When boundaries are present, the stencils near the boundaries are modified to take into account the boundary conditions.

Algorithmically, *ghost cells* extending beyond the boundaries are used to implement modified finite-difference stencils near the boundaries. The numerical scheme continues to use standard divergence $\boldsymbol{D}$ (face-to-cell) and gradient (cell-to-face) $\boldsymbol{G} = -\boldsymbol{D}^*$ stencils but implements a modified Laplacian operator due to special handling of the ghost cells. If a Dirichlet condition is imposed on a given variable (e.g., a fixed wall temperature), then the ghost cell value is obtained by a linear extrapolation of the value in the neighboring interior cell (inverse reflection). If a von Neumann condition is imposed, on the other hand, then the ghost cell value is set equal to the value in the neighboring interior cell (reflection). This gives discrete operators that can be represented by the following banded matrices near the left corner (first cell) of a one-dimensional domain:

$$\boldsymbol{D} = \begin{bmatrix} -1 & 1 & 0 & \cdots \\ 0 & -1 & 1 & \cdots \\ 0 & 0 & -1 & \cdots \\ \vdots & \vdots & \vdots & \ddots \end{bmatrix}, \qquad \boldsymbol{L} = \begin{bmatrix} -(2-\alpha) & 1 & 0 & \cdots \\ 1 & -2 & 1 & \cdots \\ 0 & 1 & -2 & \cdots \\ \vdots & \vdots & \vdots & \ddots \end{bmatrix},$$

where $\alpha = -1$ for a Dirichlet condition and $\alpha = 1$ for a von Neumann condition. It is easy to verify that $\boldsymbol{L} = -\boldsymbol{DCD}^*$ is satisfied with the following diagonal scaling matrix:

$$\boldsymbol{C} = \begin{bmatrix} \beta & 0 & 0 & \cdots \\ 0 & 1 & 0 & \cdots \\ 0 & 0 & 1 & \cdots \\ \vdots & \vdots & \vdots & \ddots \end{bmatrix},$$

where $\beta = 1 - \alpha$.

This direct computation shows that in order to satisfy the discrete fluctuation-dissipation balance condition the diagonal element of $\boldsymbol{C}$ corresponding to the cell face that touches the boundary ought to be set to 2 for a Dirichlet and to 0 for a von Neumann condition. This means that the corresponding component of the stochastic flux needs to be generated using a random normal variate of variance 2 for Dirichlet, and set to zero for a von Neumann condition.

Finally, for density, the ghost cell values are generated so that the pressure in the ghost cells is equal to the pressure in the neighboring interior cell, which ensures that there is no unphysical pressure gradient in the momentum equation across the interface. There is also no stochastic mass flux through faces on the boundary independent of the type of boundary condition at the wall.

REFERENCES

[1] H. Noguchi, N. Kikuchi, and G. Gompper, *Particle-based mesoscale hydrodynamic techniques*, Europhys. Lett., 78 (2007), article 10005.

[2] R. Delgado-Buscalioni, K. Kremer, and M. Praprotnik, *Concurrent triple-scale simulation of molecular liquids*, J. Chem. Phys., 128 (2008), article 114110.

[3] G. Hu and D. Li, *Multiscale phenomena in microfluidics and nanofluidics*, Chem. Engrg. Sci., 62 (2007), pp. 3443–3454.

[4] T. M. Squires and S. R. Quake, *Microfluidics: Fluid physics at the nanoliter scale*, Rev. Modern Phys., 77 (2005), pp. 977–1026.

[5] N. G. Hadjiconstantinou, *The limits of Navier-Stokes theory and kinetic extensions for describing small-scale gaseous hydrodynamics*, Phys. Fluids, 18 (2006), article 111301.

[6] J. B. Bell, A. Garcia, and S. A. Williams, *Numerical methods for the stochastic Landau-Lifshitz Navier-Stokes equations*, Phys. Rev. E, 76 (2007), article 016708.

[7] G. De Fabritiis, M. Serrano, R. Delgado-Buscalioni, and P. V. Coveney, *Fluctuating hydrodynamic modeling of fluids at the nanoscale*, Phys. Rev. E, 75 (2007), article 026307.

[8] N. K. Voulgarakis and J.-W. Chu, *Bridging fluctuating hydrodynamics and molecular dynamics simulations of fluids*, J. Chem. Phys., 130 (2009), article 134111.

[9] A. Donev, E. Vanden-Eijnden, A. L. Garcia, and J. B. Bell, *On the Accuracy of Finite-Volume Schemes for Fluctuating Hydrodynamics*, preprint, 2009; available online from http://arxiv.org/abs/0906.2425v2.

[10] K. Kadau, C. Rosenblatt, J. L. Barber, T. C. Germann, Z. Huang, P. Carles, and B. J. Alder, *The importance of fluctuations in fluid mixing*, Proc. Natl. Acad. Sci., 104 (2007), pp. 7741–7745.

[11] O. B. Usta, A. J. C. Ladd, and J. E. Butler, *Lattice-Boltzmann simulations of the dynamics of polymer solutions in periodic and confined geometries*, J. Chem. Phys., 122 (2005), article 094902.

[12] P. J. Atzberger, P. R. Kramer, and C. S. Peskin, *A stochastic immersed boundary method for fluid-structure dynamics at microscopic length scales*, J. Comput. Phys., 224 (2007), pp. 1255–1292.

[13] C. Aust, M. Kroger, and S. Hess, *Structure and dynamics of dilute polymer solutions under shear flow via nonequilibrium molecular dynamics*, Macromolecules, 32 (1999), pp. 5660–5672.

[14] F. J. Alexander and A. L. Garcia, *The direct simulation Monte Carlo method*, Comput. Phys., 11 (1997), pp. 588–593.

[15] X. Fan, N. Phan-Thien, S. Chen, X. Wu, and T. Y. Ng, *Simulating flow of DNA suspension using dissipative particle dynamics*, Phys. Fluids, 18 (2006), article 063102.

[16] M. Ripoll, K. Mussawisade, R. G. Winkler, and G. Gompper, *Low-Reynolds-number hydrodynamics of complex fluids by multi-particle-collision dynamics*, Europhys. Lett., 68 (2004), pp. 106–112.

[17] S. H. Lee and R. Kapral, *Mesoscopic description of solvent effects on polymer dynamics*, J. Chem. Phys., 124 (2006), article 214901.

[18] A. Donev, A. L. Garcia, and B. J. Alder, *A thermodynamically-consistent non-ideal stochastic hard-sphere fluid*, J. Stat. Mech. Theory Exp., 2009 (2009), article P11008.

[19] H. S. Wijesinghe and N. G. Hadjiconstantinou, *Discussion of hybrid atomistic-continuum methods for multiscale hydrodynamics*, Internat. J. Multiscale Comput. Engrg., 2 (2004), pp. 189–202.

[20] K. M. Mohamed and A. A. Mohamad, *A review of the development of hybrid atomistic–continuum methods for dense fluids*, Microfluidics and Nanofluidics, (2009), pp. 1–20.

[21] D. C. Wadsworth and D. A. Erwin, *One-Dimensional Hybrid Continuum/Particle Simulation Approach for Rarefied Hypersonic Flows*, AIAA Paper 90-1690, 1990.

[22] X. B. Nie, S. Y. Chen, and M. O. Robbins, *Hybrid continuum-atomistic simulation of singular corner flow*, Phys. Fluids, 16 (2004), pp. 3579–3591.

[23] N. G. Hadjiconstantinou, *Combining atomistic and continuum simulations of contact-line motion*, Phys. Rev. E, 59 (1999), pp. 2475–2478.

[24] X. Liu, Y. Li, J. Glimm, and X. L Li, *A front tracking algorithm for limited mass diffusion*, J. Comput. Phys., 222 (2007), pp. 644–653.

[25] H. Jin, X. F. Liu, T. Lu, B. Cheng, J. Glimm, and D. H. Sharp, *Rayleigh–Taylor mixing rates for compressible flow*, Phys. Fluids, 17 (2005), article 024104.

[26] A. L. Garcia, J. B. Bell, W. Y. Crutchfield, and B. J. Alder, *Adaptive mesh and algorithm refinement using direct simulation Monte Carlo*, J. Comput. Phys., 154 (1999), pp. 134–155.

[27] D. K. Datta, R. C. Picu, and M. S. Shephard, *Composite grid atomistic continuum method: An adaptive approach to bridge continuum with atomistic analysis*, Internat. J. Multiscale Comput. Engrg., 2 (2004), pp. 401–419.

[28] D. C. Wadsworth and D. A. Erwin, *Two-Dimensional Hybrid Continuum/Particle Simulation Approach for Rarefied Hypersonic Flows*, AIAA Paper 92-2975, 1992.

[29] D. B. Hash and H. A. Hassan, *Assessment of schemes for coupling Monte Carlo and Navier-Stokes solution methods*, J. Thermophys. Heat Transfer, 10 (2001), pp. 242–249.

[30] T. E. Schwartzentruber and I. D. Boyd, *A hybrid particle-continuum method applied to shock waves*, J. Comput. Phys., 215 (2006), pp. 402–416.

[31] M. Kalweit and D. Drikakis, *Multiscale methods for micro/nano flows and materials*, J. Comput. Theoret. Nanosci., 5 (2008), pp. 1923–1938.

[32] R. Delgado-Buscalioni and A. Dejoan, *Nonreflecting boundaries for ultrasound in fluctuating hydrodynamics of open systems*, Phys. Rev. E, 78 (2008), article 046708.

[33] R. Delgado-Buscalioni, P. V. Coveney, and G. D. Fabritiis, *Towards multi-scale modelling of complex liquids using hybrid particle–continuum schemes*, Proc. Institut. Mech. Engineers C, 222 (2008), pp. 769–776.

[34] R. Delgado-Buscalioni, E. G. Flekkøy, and P. V. Coveney, *Fluctuations and continuity in particle-continuum hybrid simulations of unsteady flows based on flux-exchange*, Eur. Phys. Lett., 69 (2005), pp. 959–965.

[35] E. G. Flekkoy, G. Wagner, and J. Feder, *Hybrid model for combined particle and continuum dynamics*, Europhys. Lett., 52 (2000), pp. 271–276.

[36] S. A. Williams, J. B. Bell, and A. L. Garcia, *Algorithm refinement for fluctuating hydrodynamics*, Multiscale Model. Simul., 6 (2008), pp. 1256–1280.

[37] R. Delgado-Buscalioni and G. De Fabritiis, *Embedding molecular dynamics within fluctuating hydrodynamics in multiscale simulations of liquids*, Phys. Rev. E, 76 (2007), article 036709.

[38] G. Giupponi, G. De Fabritiis, and P. V. Coveney, *Hybrid method coupling fluctuating hydrodynamics and molecular dynamics for the simulation of macromolecules*, J. Chem. Phys., 126 (2007), article 154903.

[39] M. P. Brenner, X. D. Shi, and S. R. Nagel, *Iterated instabilities during droplet fission*, Phys. Rev. Lett., 73 (1994), pp. 3391–3394.

[40] M. Moseler and U. Landman, *Formation, stability, and breakup of nanojets*, Science, 289 (2000), pp. 1165–1169.

[41] J. Eggers, *Dynamics of liquid nanojets*, Phys. Rev. Lett., 89 (2002), article 084502.

[42] C. Van den Broeck, P. Meurs, and R. Kawai, *From Maxwell demon to Brownian motor*,

New J. Phys., 7 (2005), article 10.

[43] M. Wu, G. Ahlers, and D. S. Cannell, *Thermally induced fluctuations below the onset of Rayleigh-Bénard convection*, Phys. Rev. Lett., 75 (1995), pp. 1743–1746.

[44] I. Bena, M. Malek Mansour, and F. Baras, *Hydrodynamic fluctuations in the Kolmogorov flow: Linear regime*, Phys. Rev. E, 59 (1999), pp. 5503–5510.

[45] I. Bena, F. Baras, and M. Malek Mansour, *Hydrodynamic fluctuations in the Kolmogorov flow: Nonlinear regime*, Phys. Rev. E, 62 (2000), pp. 6560–6570.

[46] A. Lemarchand and B. Nowakowski, *Fluctuation-induced and nonequilibrium-induced bifurcations in a thermochemical system*, Molecular Simulation, 30 (2004), pp. 773–780.

[47] S. Wijesinghe, R. Hornung, A. L. Garcia, and N. Hadjiconstantinou, *Three-dimensional hybrid continuum-atomistic simulations for multiscale hydrodynamics*, J. Fluids Engrg., 126 (2004), pp. 768–777.

[48] A. Donev, A. L. Garcia, and B. J. Alder, *Stochastic event-driven molecular dynamics*, J. Comput. Phys., 227 (2008), pp. 2644–2665.

[49] B. J. Alder and T. E. Wainwright, *Studies in molecular dynamics.* I. *General method*, J. Chem. Phys., 31 (1959), pp. 459–466.

[50] B. Dünweg and A. J. C. Ladd, *Lattice Boltzmann simulations of soft matter systems*, in Advanced Computer Simulation Approaches for Soft Matter Sciences III, 2009, pp. 89–166.

[51] K. T. Mashiyama and H. Mori, *Origin of the Landau-Lifshitz hydrodynamic fluctuations in nonequilibrium systems and a new method for reducing the Boltzmann equation*, J. Stat. Phys., 18 (1978), pp. 385–407.

[52] P. Español, *Stochastic differential equations for non-linear hydrodynamics*, Phys. A, 248 (1998), pp. 77–96.

[53] L. D. Landau and E. M. Lifshitz, *Fluid Mechanics*, Course of Theoretical Physics, Vol. 6, Pergamon Press, Oxford, 1959.

[54] M. J. Berger and P. Colella, *Local adaptive mesh refinement for shock hydrodynamics*, J. Comput. Phys., 82 (1989), pp. 64–84.

[55] J. Bell, M. Berger, J. Saltzman, and M. Welcome, *Three-dimensional adaptive mesh refinement for hyperbolic conservation laws*, SIAM J. Sci. Comput., 15 (1994), pp. 127–138.

[56] W. Ren, *Analytical and numerical study of coupled atomistic-continuum methods for fluids*, J. Comput. Phys., 227 (2007), pp. 1353–1371.

[57] W. Ren and W. E, *Heterogeneous multiscale method for the modeling of complex fluids and micro-fluidics*, J. Comput. Phys., 204 (2005), pp. 1–26.

[58] X. B. Nie, S. Y. Chen, W. N. E, and M. O. Robbins, *A continuum and molecular dynamics hybrid method for micro-and nano-fluid flow*, J. Fluid Mech., 500 (2004), pp. 55–64.

[59] J. Liu, S. Y. Chen, X. B. Nie, and M. O. Robbins, *A continuum–atomistic simulation of heat transfer in micro-and nano-flows*, J. Comput. Phys., 227 (2007), pp. 279–291.

[60] S. Chapman and T. G. Cowling, *The Mathematical Theory of Non-Uniform Gases*, Cambridge University Press, Cambridge, UK, 1970.

[61] A. L. Garcia and B. J. Alder, *Generation of the Chapman-Enskog distribution*, J. Comput. Phys., 140 (1998), pp. 66–70.

[62] S. Chou, *Kinetic flux vector splitting for the Navier Stokes equations*, J. Comput. Phys., 130 (1997), pp. 217–230.

[63] C. P. Lowe and D. Frenkel, *The super long-time decay of velocity fluctuations in a two-dimensional fluid*, Phys. A, 220 (1995), pp. 251–260.

[64] F. Alexander, A. L. Garcia, and B. J. Alder, *Cell size dependence of transport coefficients in stochastic particle algorithms*, Phys. Fluids, 10 (1998), pp. 1540–1542. Erratum: Phys. Fluids, 12 (2000), pp. 731–731.

[65] N. G. Hadjiconstantinou, *Analysis of discretization in the direct simulation Monte Carlo*, Phys. Fluids, 12 (2000), pp. 2634–2638.

[66] A. L. Garcia, *Estimating hydrodynamic quantities in the presence of microscopic fluctuations*, Comm. Appl. Math. Comput. Sci., 1 (2006), pp. 53–78.

[67] J. M. O. De Zarate and J. V. Sengers, *Hydrodynamic Fluctuations in Fluids and Fluid Mixtures*, Elsevier Science, Oxford, 2006.

[68] B. J. Alder and T. E. Wainwright, *Decay of the velocity autocorrelation function*, Phys. Rev. A, 1 (1970), pp. 18–21.

[69] M. W. Heemels, M. H. J. Hagen, and C. P. Lowe, *Simulating solid colloidal particles using the Lattice-Boltzmann method*, J. Comput. Phys., 164 (2000), pp. 48–61.

[70] T. Ihle and D. M. Kroll, *Stochastic rotation dynamics.* II. *Transport coefficients, numerics, and long-time tails*, Phys. Rev. E, 67 (2003), article 066706.

[71] N. Sharma and N. A. Patankar, *Direct numerical simulation of the Brownian motion of particles by using fluctuating hydrodynamic equations*, J. Comput. Phys., 201 (2004), pp. 466–486.

[72] J. T. Padding and A. A. Louis, *Hydrodynamic interactions and Brownian forces in colloidal suspensions: Coarse-graining over time and length scales*, Phys. Rev. E, 74 (2006), article 031402.

[73] P. J. Atzberger, *Velocity correlations of a thermally fluctuating Brownian particle: A novel model of the hydrodynamic coupling*, Phys. Lett. A, 351 (2006), pp. 225–230.

[74] T. Iwashita, Y. Nakayama, and R. Yamamoto, *A numerical model for Brownian particles fluctuating in incompressible fluids*, J. Phys. Soc. Japan, 77 (2008), article 074007.

[75] A. Donev, A. L. Garcia, and B. J. Alder, *Stochastic hard-sphere dynamics for hydrodynamics of non-ideal fluids*, Phys. Rev. Lett., 101 (2008), article 075902.

[76] R. Kubo, *The fluctuation-dissipation theorem*, Rep. Progr. Phys., 29 (1966), pp. 255–284.

[77] E. H. Hauge and A. Martin-Lof, *Fluctuating hydrodynamics and Brownian motion*, J. Stat. Phys., 7 (1973), pp. 259–281.

[78] R. Zwanzig and M. Bixon, *Compressibility effects in the hydrodynamic theory of Brownian motion*, J. Fluid Mech., 69 (1975), pp. 21–25.

[79] T. V. Lokotosh and N. P. Malomuzh, *Lagrange theory of thermal hydrodynamic fluctuations and collective diffusion in liquids*, Phys. A, 286 (2000), pp. 474–488.

[80] T. Iwashita, Y. Nakayama, and R. Yamamoto, *Velocity autocorrelation function of fluctuating particles in incompressible fluids: Toward direct numerical simulation of particle dispersions*, Progr. Theoret. Phys. Suppl., 178 (2009), pp. 86–91.

[81] E. Lieb, *Some problems in statistical mechanics that I would like to see solved*, Phys. A, 263 (1999), pp. 491–499.

[82] E. Kestemont, C. Van den Broeck, and M. Malek Mansour, *The 'adiabatic' piston: And yet it moves*, Europhys. Lett., 49 (2000), pp. 143–149.

[83] J. A. White, F. L. Roman, A. Gonzalez, and S. Velasco, *The 'adiabatic' piston at equilibrium: Spectral analysis and time-correlation function*, Europhys. Lett., 59 (2002), pp. 479–485.

[84] M. Malek Mansour, A. L. Garcia, and F. Baras, *Hydrodynamic description of the adiabatic piston*, Phys. Rev. E, 73 (2006), article 016121.

[85] M. Cencini, L. Palatella, S. Pigolotti, and A. Vulpiani, *Macroscopic equations for the adiabatic piston*, Phys. Rev. E, 76 (2007), article 051103.

[86] C. Bustamante, J. Liphardt, and F. Ritort, *The nonequilibrium thermodynamics of small systems*, Phys. Today, 58 (2005), pp. 43–48.

[87] C. Van den Broeck, E. Kestemont, and M. M. Mansour, *Heat conductivity by a shared piston*, Europhys. Lett., 56 (2001), pp. 771–777.

[88] J. B. Bell, J. Foo, and A. L. Garcia, *Algorithm refinement for the stochastic Burgers equation*, J. Comput. Phys., 223 (2007), pp. 451–468.

[89] A. L. Garcia, M. Malek Mansour, G. C. Lie, M. Mareschal, and E. Clementi, *Hydrodynamic fluctuations in a dilute gas under shear*, Phys. Rev. A, 36 (1987), pp. 4348–4355.

[90] M. M. Mansour, A. L. Garcia, J. W. Turner, and M. Mareschal, *On the scattering function of simple fluids in finite systems*, J. Stat. Phys., 52 (1988), pp. 295–309.

[91] P. J. Atzberger, *Spatially adaptive stochastic numerical methods for intrinsic fluctuations in reaction-diffusion systems*, J. Comput. Phys., 229 (2010, pp. 3474–3501.