Inertial Coupling Method for particles in an incompressible fluctuating fluid

Florencio Balboa Usabiaga,[1] Rafael Delgado-Buscalioni,[1] Boyce E. Griffith,[2,3] and Aleksandar Donev[3,*]

[1] *Departamento de Física Teórica de la Materia Condensada, Condensed Matter Physics Center (IFIMAC), Univeridad Autónoma de Madrid, Madrid 28049, Spain*

[2] *Leon H. Charney Division of Cardiology, Department of Medicine, New York University School of Medicine, New York, NY 10016*

[3] *Courant Institute of Mathematical Sciences, New York University, New York, NY 10012*

We develop an inertial coupling method for modeling the dynamics of point-like "blob" particles immersed in an incompressible fluid, generalizing previous work for compressible fluids [*F. Balboa Usabiaga, I. Pagonabarraga, and R. Delgado-Buscalioni, J. Comp. Phys., 235:701-722, 2013*]. The coupling consistently includes excess (positive or negative) inertia of the particles relative to the displaced fluid, and accounts for thermal fluctuations in the fluid momentum equation. The coupling between the fluid and the blob is based on a no-slip constraint equating the particle velocity with the local average of the fluid velocity, and conserves momentum and energy. We demonstrate that the formulation obeys a fluctuation-dissipation balance, owing to the non-dissipative nature of the no-slip coupling. We develop a spatio-temporal discretization that preserves, as best as possible, these properties of the continuum formulation. In the spatial discretization, the local averaging and spreading operations are accomplished using compact kernels commonly used in immersed boundary methods. We find that the special properties of these kernels allow the blob to provide an effective model of a particle; specifically, the volume, mass, and hydrodynamic properties of the blob are remarkably grid-independent. We develop a second-order semi-implicit temporal integrator that maintains discrete fluctuation-dissipation balance, and is not limited in stability by viscosity. Furthermore, the temporal scheme requires only constant-coefficient Poisson and Helmholtz linear solvers, enabling a very efficient and simple FFT-based implementation on GPUs. We numerically investigate the performance of the method on several standard test problems. In the deterministic setting, we find the blob to be a remarkably robust approximation to a rigid sphere, at both low and high Reynolds numbers. In the stochastic setting, we study in detail the short and long-time behavior of the velocity autocorrelation function and observe agreement with all of the known behavior for rigid sphere immersed in a fluctuating fluid. The proposed inertial coupling method provides a low-cost coarse-grained (minimal resolution) model of particulate flows over a wide range of time-scales ranging from Brownian to convection-driven motion.

Keywords: fluctuating hydrodynamics; Brownian motion; inertial coupling; immersed-boundary method; minimally-resolved particulate flows

I. Introduction

The dynamics of small particles immersed in a fluid is key to many applications involving disparate length and time scales [1]: from the dynamics of millimeter particles (dust) in turbulent flow, to multiphase flow with micron and nanoscopic colloidal molecules in quiescent, laminar [2–4], or turbulent regimes [5, 6]. In many engineering applications colloidal particles are exposed to disparate dynamic regimes coexisting in different subdomains of the same chamber [7]. Such processes demand fast computational methods able to efficiently resolve the motion of many ($O(10^5)$) colloidal particles in quite different dynamics ranging from diffusive to inertial dynamics. Such scenarios are paradigmatic of what one might call *multi-regime* systems.

A group of methods such as smooth particle hydrodynamics (SPH) [8], smoothed dissipative particle dynamics [9], and stochastic rotation dynamics (SRD) [2] resolve both the particle and fluid phase using similar discrete Lagrangian descriptions, and as such, seem to be natural candidates to become multi-regime solvers [8, 10]. Particle-particle methods allow for an easy treatment of complex boundary conditions, and offer a natural way to couple moving boundaries or immersed particles to the fluid. However, particle-particle methods have important drawbacks when compared with standard solvers for discretized Computational Fluid Dynamics (CFD). In particular, they

*Electronic address: donev@courant.nyu.edu

offer limited control over the fluid properties and require relatively small time steps compared to, for example, semi-implicit CFD schemes. Moreover, they cannot be adapted to efficiently treat the natural time scales governing the different dynamical regimes (e.g., the Brownian or overdamped limit). Similar advantages and drawbacks also apply to the lattice Boltzmann (LB) method [3], although the LB approach has proven to be a rather flexible framework **[11]**.

Many other approaches use CFD for the solvent flow and couple its dynamics with that of the immersed particles. In the realm of CFD one can still distinguish two large subgroups of methods. The first group of methods involves a Lagrangian description of the computational mesh which self-adapts to follow the particle [12]. The second group uses a fixed (Eulerian) grid and requires converting the particle boundary conditions into body forces or some interaction equations [11, 13, 14]. The present work focuses on this second group, sometimes called mixed Eulerian-Lagrangian methods. These schemes are particularly suited to attack the "multi-regime" problem, because they are faster, more flexible and can work with minimal resolution models (pointwise particle descriptions).

In their seminal work, Maxey and Riley [15] decomposed the fluid velocity as $\mathbf{v}(\mathbf{r},t) = \mathbf{v}_0(\mathbf{r},t) + \mathbf{v}_1(\mathbf{r},t)$, where $\mathbf{v}_0$ is the undisturbed flow (which would result if the boundary conditions at the particle surface were not applied), and $\mathbf{v}_1$ is the perturbative component created by the fluid-particle interaction. In the bulk flow, convection (advection) becomes relevant for $\mathrm{Re}_F = v_0 L\rho/\eta > 1$; where the *fluid* Reynolds number Re_F is defined in terms of the typical flow speed v_0, the fluid density ρ, the dynamic viscosity $\eta = \rho\nu$, and a characteristic length L for velocity variation in the flow. The fluid force on the particle arises from the local fluid inertia (proportional to the local material derivative of $\mathbf{v}_0$) and also from the local stress created by the particle disturbance. The *relaxational* part of the particle inertia is a consequence of its mass resistance to instantaneously follow the velocity of the surrounding fluid; the fluid drag damps the particle velocity towards the local fluid velocity within an inertial time $\tau_P \sim (\rho_P - \rho)R^2/\eta$ which increases with the density contrast $\rho_P - \rho$ and with the particle radius R.

By contrast, *convective* inertia arises from non-linear interactions between the particle dynamics and its perturbative flow [16]. The *particle* Reynolds number $\mathrm{Re}_P = 2wR/\nu$, defined with the particle-fluid relative speed w [15], determines the relative strength of advection by the perturbative flow relative to viscous dissipation. The importance of convective inertia is indicated by the ratio $\mathrm{Re}_F(R/L)^2$ between the characteristic times associated with Stokes drag and convection [15, 16]. At finite values of the non-dimensional groups Re_P and $\mathrm{Re}_F(R/L)^2$ inertia effects due to particle mass and particle size are not similar anymore, especially in the turbulent regime, where non-linear interactions between the mean flow and the particle perturbative field become relevant [5, 17]. Non-linear interaction between particle advection and thermal fluctuations are also possible at small Reynolds number. Some examples are the change in the mobility of colloidal particles ($R \sim 10^{-[5-8]}$ m) over the Stokes limit at low values of the Schmidt number [18, 19], and inertial effects in directional locking, a process to separate nanoparticles at very small Re_P [20].

Computational approaches are usually tailored to tackle some specific dynamical regime and they can be naturally classified according to the range of Re_F, Re_P and R/L they can be safely applied to. In the creeping flow limit, $\mathrm{Re}_F \to 0$ and $\mathrm{Re}_P \to 0$, the perturbative flow $\mathbf{v}_1$ has a negligible effect on the unperturbed field, which is a priori fixed. The perturbative field created by a collection of particles is the linear superposition of the Stokes fields and it determines the multi-body hydrodynamic forces on the particle ensemble. Analytical expressions for these forces are embedded in the mobility matrix of Brownian hydrodynamics (BD) [21, 22] and Stokesian dynamics (SD) [23]. In addition to the stokeslet (monopole) terms, in SD one can include higher terms of the multipole expansion of the perturbative stress [1]. The zero-Reynolds regime resolves the long-time diffusive (Smoluchowski) limit of colloidal motion, in which fluctuations make an important ($O(1)$) contribution. Direct implementation of the fluctuation dissipation (FD) relation between the friction and noise matrices requires $\mathrm{O}(N^3)$ operations, where N is the number of particles. Sophisticated and technically-complex techniques such as the accelerated Stokesian dynamics [24], and the general geometry Ewald-like method for confined geometries [25], reduce the large raw cost to $O(N \ln N)$ operations, albeit with large multiplicative prefactors.

As an alternative to BD and SD methods, two-way coupling algorithms using a Stokes frictional force were developed for mixed Eulerian-Lagrangian dynamics [3, 26, 27]. The idea is to deploy a relative simple and efficient fluid solver to explicitly resolve the perturbative flow responsible for the hydrodynamic coupling between particles. The particle cost is dominated by neighbour searching and

scales (almost) linearly with N, while the (added) fluid solver cost scales like the system volume. The Eulerian-Lagrangian mixed approach permits to work at finite Re_F. However, the Stokes (i.e. frictional) coupling assumption limits the scheme to $\mathrm{Re}_P < 1$ and only resolves far-field hydrodynamics ($R/L < 1$). The Stokes coupling consistently neglects convective inertia and only includes relaxational particle inertia in an approximate way, with a finite particle response time τ_P introduced by a *phenomenological* friction coefficient. Frictional coupling is obviously dissipative and requires introducing an additional noise term in the particle equation, different from that of the fluctuating fluid stress tensor [3, 27].

Other methods for finite Re_F have been restricted to $\mathrm{Re}_P = 0$, where the particle inertia is absent. Two relevant examples are the stochastic Immersed Boundary(IB) method [27, 28] commonly used for fluid-structure interaction at $R/L = O(1)$, and the Force Coupling method (FCM) [16, 29, 30], where each particle is represented by a low-order expansion of force multipoles ($R/L < 1$). Very recently, subsequent to the competition and submission of this paper, an extension of the FCM that includes fluctuations in the overdamped or inertia-less limit has been developed [31]. For $\mathrm{Re}_P = 0$ the relative fluid-particle acceleration is zero and the particle velocity just follows the local fluid velocity. The hydrodynamic force due to the particle-fluid interaction is then equal to the total force exerted on the particle by sources other than the fluid. This permits a fluid-only formulation whereby the net non-hydrodynamic particle force is *spread* from the particle to the surrounding fluid using some compact kernel. This important spreading operation differs substantially from method to method. In FCM two different Gaussian kernels are used to spread the force monopole and force dipole moments (stresslet); their widths are fitted in the *continuum* model to recover the Stokes drag and linear Faxen terms [29]. By contrast, the IBM kernels are specifically designed to minimize the effects of the *discrete* Eulerian mesh in the spreading of Lagrangian point forces (monopole terms) [32].

For $\mathrm{Re}_P \ll 1$ and $R/L \ll 1$ advection of the perturbative flow can be neglected leading to the (analytically solvable) unsteady Stokes equation for the perturbative field [15]. The fluid-particle force can be expressed as some function of the relative velocity field $\mathbf{u}-\mathbf{v}_0$ *interpolated* at the particle site. This forms the basis of *one-way-coupling* schemes for point-particle dynamics frequently used in turbulence research, $\mathrm{Re}_F \gg 1$ [16]. Generalizations to $\mathrm{Re}_P \sim 1$ have been also derived (see e.g. [33]) but, even in the simpler $\mathrm{Re}_P \ll 1$ limit, the evaluation of the fluid-particle force involves cumbersome expressions which require interpolations of the displaced fluid acceleration (the added mass effect), second order spatial derivatives of $\boldsymbol{v}_0$ (Faxen terms), and time-convolved integrals which recast the history of vorticity diffusion around the particle (Basset memory). For a sphere moving with velocity $\mathbf{u}$ at $\mathrm{Re}_P \ll 1$, the leading term is the steady Stokes force $F_{\mathrm{Stokes}} = 6\pi\eta R(\mathbf{u} - \mathbf{v}_0)$, which, due to its simple form, has been overused in *two-way* point-particle approximations of turbulent ($\mathrm{Re}_F \gg 1$) and pre-turbulent regimes [5, 17]. Although the point-particle approach can probably describe the relaxational inertia of very small ($R/L \ll 1$) heavy particles in a light fluid (e.g. aerosol), it has the serious limitation of neglecting the convective inertia arising from the particle finite size [34]. Even at low Re_P, convection of perturbative flow is known to alter the Basset memory and the long time particle dynamics [16]. And vice versa, recent works show that micron-size particles can alter the turbulent spectra at moderate Re_P and $R \sim L$ (with L the Kolmogorov length) [5, 6] due to energy dissipation and vorticity production in the particles wake [34].

Several Eulerian-Lagrangian methods have appeared in recent years to allow for a fully consistent treatment of the coupled particle and fluid inertia. A key issue is the spatial resolution of the particle. In the "direct forcing" method [35], and related extensions to fluctuating hydrodynamics [11, 36–38], the fluid force on the particle is obtained by imposing the no-slip constraint on a well-resolved particle surface, and in some cases, to ensure rigid body motion, also in the particle interior [39]. High spatial resolution requires substantial computational effort; the largest simulations so far reached $O(10^3)$ particles [35, 40, 41]. The smoothed particle method (SPM) [13, 14] works with a mixed (particle-fluid) velocity field constructed with a smooth characteristic function which discriminates particle and fluid cells. This permits an intermediate resolution with a typical particle radius $R \simeq 5h$ (here h is the mesh size) requiring $O(10^3)$ fluid cells per particle. These fully or partially resolved methods are quite far from a point-particle approach, which can require as few as 13 cells to perform a fourth order orthogonal Lagrangian interpolation [42]. "Blob" particle descriptions, more appropriately termed *minimally-resolved models*, offer a way to explore finite particle effects at moderate computational cost. In this work each particle (blob) is described by an unique kernel of small support (27 cells) which is shown to provide a consistent set of particle physical properties (volume, mass, hydrodynamic radius).

It has to be noted that due to the long-ranged nature of hydrodynamic interactions, sufficiently large boxes are required to reduce (to an acceptable extent) the finite size effects; hence a reduction in the linear size of the particle description means a large (cubic) reduction in the overall fluid solver cost.

In a preceding paper [43], some of us proposed an inertial coupling method that directly couples a compressible finite-volume fluctuating hydrodynamic solver [44] to blob particles. A distinguishing feature of the coupling methodology is that it includes the effect of the particle and fluid inertia in the dynamics, while still consistently including thermal fluctuations even in non-trivial geometries. It was numerically demonstrated that the inertial coupling method can reproduce ultrasound forces on colloidal particles, taking place at much faster rates than viscous friction [45]. In this previous work [43], a compressible solver was used because one of the focus applications was the interaction between ultrasound and colloidal particles [45]. In many applications sonic effects can be ignored and the essential hydrodynamic interactions can be captured by using the *isothermal incompressible* Navier-Stokes equations instead of the compressible equations. This eliminates the fast sound waves and allows for a much larger time step size in the fluid solver. Here we develop an *inertial coupling method* (ICM) that directly couples an incompressible finite-volume solver for the fluctuating Navier-Stokes equations [44] with suspended particles, which do not necessarily have the same density as the fluid. We demonstrate that the coupling obeys a continuum and a discrete fluctuation-dissipation balance and study the performance of our algorithm. The ICM is a coarse-grained model for particle hydrodynamics which aims to capture hydrodynamic effects (unsteady forcing, viscous friction and advection) over a broad range of time scales and Re_P: from Brownian motion to convection-driven regimes. In the ICM, the coupling between the particle and the fluid is not assumed to have any functional form (e.g. Stokes drag) but *naturally* arises from the no-slip constraint averaged over the particle (or "blob") domain. The present results (see also Ref. [43]) indicate that this type of (non-linear) coupling permits to take into account both fluid *and* particle inertia beyond the Stokes limit, where advective interactions take place.

It is useful to point out the main similarities and differences between our work and closely related work by others. First, at the level of the formulation, the idea of using a minimally-resolved description of particles immersed in a fluid has a long history in both deterministic and fluctuating hydrodynamics. In the deterministic setting, Maxey and collaborators have explored in extensive detail the coupling between a minimally-resolved spherical particle and fluid flow within the context of the FCM [16, 29, 30]. In the context of fluctuating hydrodynamics, a blob description of particles has been used frequently to couple a Lattice-Boltzmann (LB) fluid solver to immersed particles [3, 26]. There are two distinguishing features of our method from that described in extensive detail in an excellent survey by Dünweg and Ladd [3]. The first difference, inherited from our previous work [43], is the fact that we employ a *direct* frictionless (conservative) coupling that instantaneously enforces a no-slip constraint. This eliminates an artificial frictional time scale from the problem and allows us to obtain physically-accurate short-time dynamics, as we demonstrate in detail in Section IV by examining the velocity autocorrelation function of the blobs for very short times. In principle, a more direct coupling could be done in the formulation of Dünweg and Ladd by taking a very large frictional constant, as explained in more detail by Atzberger [27]. However, the resulting dynamics are stiff and doing this numerically requires using small time steps sizes. The second important distinguishing feature of the work presented here is the use of a semi-implicit incompressible fluid solver instead of the compressible explicit LB solver, as also done in the stochastic IB method [27, 28]. This allows us to maintain stability at much larger time step size than used in typical LB simulations, and, more importantly, allows us to approach the Brownian or Stokesian (overdamped) dynamics limit without any *uncontrolled* approximations. A final distinguishing feature of our work from all work we are aware of is the consistent inclusion of inertial effects in both the formulation and the numerical algorithm. Specifically, inertial effects (excess particle mass) are not included in the FCM or the stochastic IB method. Recently, the very strong coupling limit of the frictional formulation described in Refs. [3, 27] has been considered theoretically by Tabak and Atzberger [46]. The equations obtained by them are identical to the ones we derive here based on physical guiding principles. Finally, we point out that our numerical algorithm, which includes both the spatial and a *second-order semi-implicit* temporal discretizations of the fluid-particle equations, is distinct from any other work we are aware of. The algorithm is a generalization of the first order coupling introduced by the direct forcing method [39, 43, 47]. Our scheme achieves second-order accuracy, avoids pressure-velocity splitting [48], as required for low Reynolds number

flows, and can be extended to other immersed-boundary fluid-structure coupling methods as well.

It is important to point out that in this work we do not consider thermal (heat) transfer between the particles and the fluid and instead use an isothermal description, as most relevant to microfluidic applications. Heat transfer is important in many large-scale particulate flows, and has been included in Direct Numerical Simulation (DNS) algorithms [41]. It is possible to include additional transport processes in the minimally-resolved approach we employ here, however, the specifics of how to express the surface boundary conditions to a volumetric blob condition are very problem specific and need to be carefully constructed on a case by case basis. In this respect, recently, reaction-diffusion processes have been included in the type of method studied here [49], and compressible blobs have been considered and found to adequately describe coupling between ultrasound waves and small particles [50].

In the remainder of this Introduction we introduce some notation and fundamental concepts. In Section II we discuss the continuum equations of the incompressible inertial coupling method. We present both a constrained and a constraint-free formulation, and demonstrate momentum and energy conservation, as well as fluctuation-dissipation balance. In Section III we present a second-order semi-implicit spatio-temporal discretization of the continuum equations, and demonstrate second-order temporal accuracy in the deterministic setting. In Section IV we test and apply the algorithm to a collection of standard test problems. We demonstrate that the algorithm correctly reproduces static equilibrium properties such as the radial distribution function in a suspension of soft spheres, and also reproduces known features of single and pair hydrodynamic interactions at small Reynolds numbers. We also study in detail the short and long-time behavior of the velocity autocorrelation function of a freely diffusing particle. Finally, we study the behavior of the blob particle in high Reynolds number flow. In Section V we offer some conclusions and thoughts on possible extensions of the method and algorithm. Several more technical calculations and algorithmic details are presented in a collection of Appendices.

A. Notation and Basic Concepts

In the beginning, we focus on the continuum formulation of the fluid-particle coupling. However, it is important to point out that most of the notation and conclusions can directly be adopted in the discrete formulation by simply replacing spatial integrals with sums over grid points. We will return to the spatially-discrete formulation in Section III.

Let us consider a particle of physical mass m and size (e.g., radius) a immersed in a fluid with density ρ. In real problems there will be many particles $i = 1, \ldots, N_p$ that interact with each other, for example, in microfluidic applications involving polymers each particle could represent a bead in a bead-spring or bead-link polymer model [22]. Unless otherwise indicated it is straightforward to extend the proposed formulation to a collection of interacting particles by simply adding a summation over the different particles. Therefore, for simplicity of notation, we will typically focus on a single particle and omit the particle index.

The position of the particle is denoted with $\boldsymbol{q}(t)$ and its velocity with $\boldsymbol{u} = \dot{\boldsymbol{q}}$. The shape of the particle and its effective interaction with the fluid is captured through a smooth kernel function $\delta_a(\boldsymbol{r})$ that integrates to unity and whose support is localized in a region of size a. For example, one may choose any one-dimensional "bell-shaped" curve $\delta_a(r)$ with half-width of order a, and define a spherically-symmetric $\delta_a(\boldsymbol{r}) = \delta_a(r)$; alternatively, in d dimensions one may define a tensor-product

$$\delta_a(\boldsymbol{r}) = \prod_{\alpha=1}^{d} \delta_a(r_\alpha). \tag{1}$$

In immersed-boundary methods [32], the kernel function δ_a is considered to be an approximation of the Dirac delta function of purely numerical origin and has the tensor-product form (1). By contrast, in the force-coupling method [29, 30], the shape of the kernel function is chosen to be a spherically-symmetric Gaussian whose width is related to the physical size and properties of the actual particle. We adopt an approach that is intermediate between these two extremes and choose the shape of the function based on numerical considerations, but relate its shape to the physical properties of the particle.

The fluid velocity field is denoted with $\boldsymbol{v}(\boldsymbol{r}, t)$ and is assumed to extend over the whole domain, including the particle interior. The interaction between the fluid and particle is mediated via the kernel function through two crucial local operations. The *local averaging* linear operator $\boldsymbol{J}(\boldsymbol{q})$ averages

the fluid velocity inside the particle to estimate a local fluid velocity

$$\boldsymbol{v}_{\boldsymbol{q}}(t) = \boldsymbol{J}\boldsymbol{v}(\boldsymbol{r},t) = \int \delta_a(\boldsymbol{q}-\boldsymbol{r})\,\boldsymbol{v}(\boldsymbol{r},t)\,d\boldsymbol{r}.$$

The reverse of local averaging is accomplished using the *local spreading* linear operator $\boldsymbol{S}(\boldsymbol{q})$ which takes a force $\boldsymbol{F}$ applied to the particle and spreads it over the extent of the kernel function to return a smooth force density field,

$$\boldsymbol{f}(\boldsymbol{r},t) = \boldsymbol{S}\boldsymbol{F}(t) = \boldsymbol{F}(t)\,\delta_a(\boldsymbol{q}-\boldsymbol{r}).$$

Note that the local spreading operator $\boldsymbol{S}$ has dimensions of inverse volume. For notational simplicity we will slightly abuse notation and assume that the local spreading and interpolation operators can be applied to a scalar, a vector, or a tensor field, with the interpretation that the same local averaging or spreading operation is applied to each component independently. This sort of block-diagonal form of the spreading and interpolation operators is not strictly required for the mathematical formulation [27], but applies to the specific Peskin forms of the operators we use in practice [32].

The physical volume of the particle ΔV is related to the shape and width of the kernel function via $\boldsymbol{J}\boldsymbol{S} = \Delta V^{-1}\boldsymbol{I}$, that is,

$$\Delta V = (\boldsymbol{J}\boldsymbol{S}\,1)^{-1} = \left[\int \delta_a^2(\boldsymbol{r})\,d\boldsymbol{r}\right]^{-1}. \tag{2}$$

Therefore, even though the particle is represented only by the position of its centroid, it is not appropriate to consider it a "point" particle. Rather, it can be thought of as a diffuse sphere that has some physical extent and interacts with the fluid in its interior. For lack of better terminology, we will refer to such a diffuse particle as a "blob". In fluctuating hydrodynamics the fluid velocity is a distribution and cannot be evaluated pointwise, therefore, to obtain well-defined fluctuating equations spatial averaging must be used and a physical volume associated to each blob.

Because fluid permeates the interior of the particle, the effective inertia of the particle is enlarged by $\rho\Delta V$ giving the physical mass

$$m = m_e + \rho\Delta V = m_e + m_f,$$

where m_e is the *excess mass* of the particle over the mass of the entrained fluid $m_f = \rho\Delta V$. In particular, $m_e = 0$ corresponds to a neutrally-buoyant particle, meaning that the inertia of the fluid is unchanged by the presence of the particle. It is a crucial property that ΔV is a constant that only depends on the shape of the kernel function and *not* on the position of the particle. This ensures that the mass of the particle m is constant and can be given a well-defined physical interpretation. Preserving this translational invariance of the physical properties of the blob in the spatially-discrete setting requires using special discrete averaging and spreading operations.

One could alternatively use the dimensionless operator $\widetilde{\boldsymbol{S}} = \boldsymbol{S}\Delta V$, as done in Ref. [43], with the property that $\boldsymbol{J}\widetilde{\boldsymbol{S}} = \boldsymbol{I}$. We prefer to use the dimensional version because the averaging and spreading operators are *adjoint*, $\boldsymbol{S} = \boldsymbol{J}^\star$, i.e., the natural dot products in the particle (Lagrangian) and fluid (Eulerian) domains are related via [27]

$$(\boldsymbol{J}\boldsymbol{v})\cdot\boldsymbol{u} = \int \boldsymbol{v}\cdot(\boldsymbol{S}\boldsymbol{u})\,d\boldsymbol{r} = \int \delta_a(\boldsymbol{q}-\boldsymbol{r})\,(\boldsymbol{v}\cdot\boldsymbol{u})\,d\boldsymbol{r} \tag{3}$$

for any $\boldsymbol{u}$ and $\boldsymbol{v}$. This adjoint property follows from the fact that the same kernel function is used is used for both averaging and spreading, and is crucial in maintaining energy conservation and fluctuation-dissipation balance. This adjoint condition will also be preserved by the discrete local averaging and spreading operators.

B. Fluctuating Incompressible Navier-Stokes Equation

In this work we assume that the fluid can be described via the fluctuating Navier-Stokes equation [51]. Specifically, we model the dynamics of the fluid velocity field $\boldsymbol{v}(\boldsymbol{r},t)$ assuming an isothermal incompressible Newtonian fluid, $\boldsymbol{\nabla}\cdot\boldsymbol{v} = 0$,

$$\rho\left(\partial_t\boldsymbol{v} + \boldsymbol{v}\cdot\boldsymbol{\nabla}\boldsymbol{v}\right) = -\boldsymbol{\nabla}\pi + \boldsymbol{\nabla}\cdot\boldsymbol{\sigma} + \boldsymbol{f} = -\boldsymbol{\nabla}\pi + \eta\boldsymbol{\nabla}^2\boldsymbol{v} + \boldsymbol{\nabla}\cdot\left[(k_BT\eta)^{\frac{1}{2}}\left(\boldsymbol{\mathcal{W}} + \boldsymbol{\mathcal{W}}^T\right)\right] + \boldsymbol{f}, \tag{4}$$

where the stress tensor $\boldsymbol{\sigma}$ includes the viscous $\eta\left(\boldsymbol{\nabla}\boldsymbol{v} + \boldsymbol{\nabla}^T\boldsymbol{v}\right)$ and fluctuating contributions, π is the non-thermodynamic pressure, ρ is the (constant) fluid density, $\eta = \rho\nu$ is the (constant) fluid shear

viscosity, ν is the kinematic viscosity, and $\boldsymbol{f}(\boldsymbol{r},t)$ is an additional force density such as gravity or the force exerted by the particles on the fluid. Note that we prefer to use the standard physics notation instead of the differential notation more common in the mathematics literature since there is no difference between the Ito and Stratonovich interpretations of stochastic integrals for additive noise.

In the momentum conservation law (4), the stochastic momentum flux is modeled using a white-noise random Gaussian tensor field $\boldsymbol{\mathcal{W}}(\boldsymbol{r},t)$, that is, a tensor field whose components are independent (space-time) white noise processes,

$$\langle \mathcal{W}_{ij}(\boldsymbol{r},t)\mathcal{W}_{kl}(\boldsymbol{r}',t')\rangle = (\delta_{ik}\delta_{jl})\,\delta(t-t')\delta(\boldsymbol{r}-\boldsymbol{r}').$$

The form of the stochastic forcing term ensures fluctuation-dissipation balance between the random forcing and the viscous dissipation and gives the correct spectrum for the thermally-induced velocity fluctuations. The symmetrized form of the fluctuating stress $(k_BT\eta)^{1/2}\left(\boldsymbol{\mathcal{W}}+\boldsymbol{\mathcal{W}}^T\right)$ mimics the symmetry of the viscous stress tensor, ensuring fluctuation-dissipation balance even for variable viscosity flows [51, 52]. The discretization and numerical solution of (4) is discussed in more detail in Refs. [44, 52].

It is important to emphasize here that the non-linear fluctuating Navier-Stokes equation forced with white-noise fluxes is ill-defined because the solution should be a distribution rather than a function and the nonlinear term $\boldsymbol{v}\cdot\boldsymbol{\nabla v}$ cannot be interpreted in the sense of distributions. This term needs to be regularized in order to give a precise meaning to (4). Such a regularization has the physical interpretation that the fluctuating fields are only defined from the underlying microscopic dynamics via spatial coarse-graining with a characteristic mesoscopic length σ, as discussed at length in Ref. [53]. In the continuum setting, one can replace the term $\boldsymbol{v}\cdot\boldsymbol{\nabla v}$ with $\tilde{\boldsymbol{v}}\cdot\boldsymbol{\nabla v}+\boldsymbol{\nabla}\tilde{\boldsymbol{v}}^T\cdot\boldsymbol{v}$, where $\tilde{\boldsymbol{v}}$ is a smoothed velocity in which features in $\boldsymbol{v}$ at scales smaller than σ are filtered [1], following the α-Navier-Stokes model [54] in turbulence. An alternative is to filter the stochastic forcing $\boldsymbol{\mathcal{W}}$ with a smoothing kernel of width σ [56] (see also Appendix in Ref. [57]). We are not aware of any careful studies of what regularization is the most appropriate (i.e., produces the best match with molecular dynamics), and we do not attempt to address this complex issue in this work.

In finite-volume or finite-element spatial discretizations, both the nonlinear terms and the stochastic forcing are naturally regularized by the discretization or coarse-graining length scale (grid spacing). In our spatial discretization, the advective term $\boldsymbol{v}\cdot\boldsymbol{\nabla v}$ is discretized using a skew-adjoint (conservative in both momentum and energy) discrete advection operator, as explained in detail in Ref. [52]. This ensures that the inclusion of that term does not alter the equilibrium Gibbs-Boltzmann distribution for the fluctuating velocity field. In practice, we have not observed any measurable effect of the nonlinearity on the results presented here, as tested by simply omitting advective fluxes in the velocity equation. This is consistent with the notion that as long as there are sufficiently many molecules per hydrodynamic cell the fluctuations will be small and the behavior of the nonlinear equations will closely follow that of the *linearized* equations of fluctuating hydrodynamics, which can be given a precise meaning [58]. Specifically, advective terms such as $\boldsymbol{v}\cdot\boldsymbol{\nabla v}$ or $\boldsymbol{u}\cdot\boldsymbol{\nabla v}$ (see, for example, Eq. (8)) scale like the *square* of the magnitude of the fluctuations, and in practice we observe they give unmeasurably small corrections when the blobs are much larger than the fluid molecules.

C. No-Slip Condition

Coupling of a continuum (fluctuating) fluid with point-like (blob) particles has been considered by other researchers. In particular, in Lattice-Boltzmann methods [3, 59, 60] a Stokes *frictional* force between the particle and the fluid is postulated. Specifically, the motion of the particle is described by a Langevin equation in which a phenomenological Stokes frictional force between the particle and the fluid is postulated, proportional to the difference $\boldsymbol{u}-\boldsymbol{J}\boldsymbol{v}$ between the particle and the locally-averaged fluid velocity. A corresponding force is added to the fluid equations to ensure momentum and energy conservation and fluctuation-dissipation balance in the fluid-particle system [3, 27].

An important downside of the inertial Stokes coupling is the imposition of an artificial friction parameter and an associated delay with the response of the particle to changes in the flow. Such a delay is often not physically acceptable unless a very large friction constant is imposed, leading

[1] In the α-Navier-Stokes equations [54] the smoothing is chosen to be an inverse Helmholtz operator, $\boldsymbol{v}=\boldsymbol{u}-\sigma^2\boldsymbol{\nabla}^2\boldsymbol{u}$, with boundary conditions chosen such that $\boldsymbol{u}$ is divergence free in the whole domain of interest [55].

to numerical stiffness. Instead, following Ref. [43], we impose an *instantaneous* coupling between the fluid and the particle in the form of a *no-slip constraint*,

$$\boldsymbol{u} = \dot{\boldsymbol{q}} = \boldsymbol{J}\boldsymbol{v} = \int \delta_a\left(\boldsymbol{q}-\boldsymbol{r}\right)\boldsymbol{v}\left(\boldsymbol{r},t\right)\,d\boldsymbol{r}, \tag{5}$$

The no-slip condition simply states that the velocity of the particle is equal to a local average of the fluid velocity. This is a constraint that formally eliminates the particle velocity from the formulation and leaves only the fluid degrees of freedom. We now demonstrate that the imposition of (5) leads to a physically-consistent coarse-grained model of the coupled fluid-particle system. Notably, the fluid-particle coupling conserves momentum, energy, and obeys a fluctuation-dissipation principle.

It is important to point out that due to the finite extent of the kernel δ_a, the particle velocity (5) differs from that of a point tracer even for a smooth fluid velocity field. As noted by Maxey and Patel [30], to second-order in the gradients of $\boldsymbol{v}$, the particle velocity obeys a Faxen relation [61],

$$\begin{aligned}\boldsymbol{u} \;&\approx\; \int \delta_a\left(\boldsymbol{q}-\boldsymbol{r}\right)\left\{\boldsymbol{v}\left(\boldsymbol{q},t\right)+\boldsymbol{\nabla}\boldsymbol{v}\left(\boldsymbol{q},t\right)\cdot\left(\boldsymbol{r}-\boldsymbol{q}\right)+\frac{1}{2}\boldsymbol{\nabla}\boldsymbol{\nabla}\boldsymbol{v}\left(\boldsymbol{q},t\right):\left[\left(\boldsymbol{r}-\boldsymbol{q}\right)\left(\boldsymbol{r}-\boldsymbol{q}\right)^{T}\right]\right\}d\boldsymbol{r} = \\ &=\; \boldsymbol{v}\left(\boldsymbol{q},t\right)+\frac{1}{2}\boldsymbol{\nabla}\boldsymbol{\nabla}\boldsymbol{v}\left(\boldsymbol{q},t\right):\int \delta_a\left(\boldsymbol{r}\right)\boldsymbol{r}\boldsymbol{r}^{T}\,d\boldsymbol{r}.\end{aligned}$$

If the kernel $\delta_a\left(\boldsymbol{r}\right)$ is spherically-symmetric,

$$\boldsymbol{u} = \boldsymbol{v}\left(\boldsymbol{q},t\right)+\left[\int \frac{r_x^2}{2}\delta_a\left(r\right)d\boldsymbol{r}\right]\boldsymbol{\nabla}^2\boldsymbol{v}\left(\boldsymbol{q},t\right) = \boldsymbol{v}\left(\boldsymbol{q},t\right)+\frac{a_F^2}{6}\boldsymbol{\nabla}^2\boldsymbol{v}\left(\boldsymbol{q},t\right), \tag{6}$$

where a_F can be termed the "Faxen" radius of the blob [30], in general different from the hydrodynamic radius (unlike for a fully-resolved rigid sphere). The same formula applies for the case of a tensor-product kernel (1).

The particle acceleration is

$$\dot{\boldsymbol{u}} = \frac{d}{dt}\left[\boldsymbol{J}\left(\boldsymbol{q}\right)\boldsymbol{v}\right] = \boldsymbol{J}\left(\partial_t\boldsymbol{v}\right)+\left(\boldsymbol{u}\cdot\frac{\partial}{\partial\boldsymbol{q}}\boldsymbol{J}\right)\boldsymbol{v}, \tag{7}$$

where for our choice of interpolation operator we have the explicit form:

$$\left(\boldsymbol{u}\cdot\frac{\partial}{\partial\boldsymbol{q}}\boldsymbol{J}\right)\boldsymbol{v} = \int\left[\boldsymbol{u}\cdot\frac{\partial}{\partial\boldsymbol{q}}\delta_a\left(\boldsymbol{q}-\boldsymbol{r}\right)\right]\boldsymbol{v}\left(\boldsymbol{r},t\right)d\boldsymbol{r}.$$

Observe that in the limit of a "point particle", $a\to 0$, the kernel function approaches a Dirac delta function and one can identify (7) with the advective derivative,

$$\frac{d}{dt}\left(\boldsymbol{J}\boldsymbol{v}\right)\approx\frac{d}{dt}\boldsymbol{v}\left(\boldsymbol{q}(t),t\right) = D_t\boldsymbol{v} = \partial_t\boldsymbol{v}+\left(\boldsymbol{v}\cdot\boldsymbol{\nabla}\right)\boldsymbol{v},$$

which is expected since in this limit the particle becomes a Lagrangian marker. In Ref. [62], the term $\frac{d}{dt}\left(\boldsymbol{J}\boldsymbol{v}\right)$ is replaced with the interpolated Navier-Stokes advective derivative $\boldsymbol{J}\left(D_t\boldsymbol{v}\right)$, thus avoiding the need to differentiate the kernel function. For a blob particle with finite size, however, in general, the relative fluid-particle acceleration is non-zero,

$$\boldsymbol{a}_{\boldsymbol{J}} = \frac{d}{dt}\left(\boldsymbol{J}\boldsymbol{v}\right)-\boldsymbol{J}\left(D_t\boldsymbol{v}\right) = \left(\boldsymbol{u}\cdot\frac{\partial}{\partial\boldsymbol{q}}\boldsymbol{J}\right)\boldsymbol{v}-\boldsymbol{J}\boldsymbol{v}\cdot\boldsymbol{\nabla}\boldsymbol{v}\neq\boldsymbol{0}. \tag{8}$$

II. Incompressible Inertial Coupling Method

Following the discussion in the Introduction and the derivation in Section 2 of Ref. [43] we take the equations of motion for a single particle coupled to a fluctuating fluid to be

$$\rho\left(\partial_t\boldsymbol{v}+\boldsymbol{v}\cdot\boldsymbol{\nabla}\boldsymbol{v}\right) = \rho D_t\boldsymbol{v} \;=\; -\boldsymbol{\nabla}\pi+\boldsymbol{\nabla}\cdot\boldsymbol{\sigma}-\boldsymbol{S}\left(\boldsymbol{q}\right)\boldsymbol{\lambda} \tag{9}$$

$$m_e\dot{\boldsymbol{u}} \;=\; \boldsymbol{F}\left(\boldsymbol{q}\right)+\boldsymbol{\lambda} \tag{10}$$

$$\text{s.t.}\quad \boldsymbol{u} = \boldsymbol{J}\left(\boldsymbol{q}\right)\boldsymbol{v}, \tag{11}$$

where the fluid-particle force $\boldsymbol{\lambda}$ is a Lagrange multiplier that enforces the constraint (11) and $\boldsymbol{F}\left(\boldsymbol{q}\right)$ is the external force applied to the particle. Observe that the total particle-fluid momentum

$$\boldsymbol{P} = m_e\boldsymbol{u}+\int\rho\boldsymbol{v}\left(\boldsymbol{r},t\right)d\boldsymbol{r}$$

is conserved because Newton's third law is enforced; the opposite total force is exerted on the fluid by the particle as is exerted on the particle by the fluid. When there is more than one particle one simply adds the forces from all the particles in the fluid equation.

Note that similar equations apply for both compressible and incompressible fluids. In the compressible case [43], a density equation is added to the system (9,10,11) and the pressure $\pi(\rho)$ obtained from the equation of state. In the incompressible case the divergence-free condition $\boldsymbol{\nabla}\cdot\boldsymbol{v}=0$ is used instead to determine the (non-thermodynamic) pressure as a Lagrange multiplier.

For now, we will silently ignore the fact that the fluctuating equations include a non-smooth white noise component that must be handled with care, and return to a discussion of the stochastic equations later on. For a neutrally-buoyant particle, $m_e=0$, $\boldsymbol{\lambda}=-\boldsymbol{F}$, and the fluid equation is the standard Navier-Stokes with the force on the particle spread back to the fluid as a force density $\boldsymbol{SF}$ [29, 32]. In this case our formulation is equivalent to the Stochastic Immersed Boundary Method [27, 28], and in the deterministic context, it is equivalent to the initial ("monopole") version of the Force Coupling Method [30].

In the determinstic setting, several extensions of the deterministic IB method to include inertial effects have already been developed by Peskin and collaborators [62–65], as well as by Uhlmann in the context of the direct-forcing method [47], to which our method is closely related. In the penalty method of Kim and Peskin [64], a particle (Lagrangian marker in the context of IBM) is represented as a pair of particles, a (neutrally-buoyant) passive tracer that follows the flow, $\dot{\boldsymbol{q}}=\boldsymbol{J}(\boldsymbol{q})\,\boldsymbol{v}$, and an inertial particle of mass m_e connected to the tracer via an elastic spring. In the limit of an infinitely stiff spring (infinite penalty parameter) the spring force $\boldsymbol{\lambda}$ becomes a Lagrange multiplier enforcing the no-slip constraint. An advantage of our constrained formulation is that it does not include the fast dynamics associated with the stiff elastic springs and thereby avoids the time step size restrictions associated with an explicit penalty method [64].

A. Primitive-variable Formulation

In this section we study the properties of (9,10,11) in order to better understand the physics of the fluid-particle coupling. Using (10) to eliminate $\boldsymbol{\lambda}=m_e\dot{\boldsymbol{u}}-\boldsymbol{F}$ and (8) to eliminate $\dot{\boldsymbol{u}}$, the fluid equation (9) becomes,

$$\rho D_t\boldsymbol{v}=\rho\left(\partial_t\boldsymbol{v}+\boldsymbol{v}\cdot\boldsymbol{\nabla}\boldsymbol{v}\right) \;=\; -m_e\boldsymbol{SJ}\left(D_t\boldsymbol{v}\right)-\boldsymbol{\nabla}\pi+\boldsymbol{\nabla}\cdot\boldsymbol{\sigma}-m_e\boldsymbol{S}\boldsymbol{a}_{\boldsymbol{J}}+\boldsymbol{SF}. \tag{12}$$

This gives the effective fluid equation

$$\left(\rho+m_e\boldsymbol{SJ}\right)\partial_t\boldsymbol{v}=-\left[\rho\left(\boldsymbol{v}\cdot\boldsymbol{\nabla}\right)+m_e\boldsymbol{S}\left(\boldsymbol{u}\cdot\frac{\partial}{\partial\boldsymbol{q}}\boldsymbol{J}\right)\right]\boldsymbol{v}-\boldsymbol{\nabla}\pi+\boldsymbol{\nabla}\cdot\boldsymbol{\sigma}+\boldsymbol{SF}, \tag{13}$$

in which the effective fluid inertia is given by the operator $\rho+m_e\boldsymbol{SJ}$, and the kinetic stress term $\rho\boldsymbol{v}\cdot\boldsymbol{\nabla}\boldsymbol{v}$ includes an additional term due to the excess inertia of the particle. When there are many interacting particles one simply adds a summation over all particles in front of all terms involving particle quantities in (13). Note that for a neutrally-buoyant particle $m_e=0$ and one obtains the constant-density Navier-Stokes equation with external forcing $\boldsymbol{SF}$.

Similarly, by eliminating $\boldsymbol{\lambda}$ from (10) we obtain the effective particle equation (see also Section 2 of Ref. [43]),

$$m\dot{\boldsymbol{u}}=\Delta V\,\boldsymbol{J}\left(-\boldsymbol{\nabla}\pi+\boldsymbol{\nabla}\cdot\boldsymbol{\sigma}\right)+\boldsymbol{F}+m_f\boldsymbol{a}_{\boldsymbol{J}}, \tag{14}$$

where $m_f=\rho\Delta V$ is the mass of the fluid dragged with the particle. This equation makes it clear why $m=m_e+m_f$ has the physical interpretation of particle mass (inertia). If the particle were a rigid sphere, the force exerted by the fluid on the particle would be the surface average of the stress tensor. It is sensible that for a blob particle this is replaced by the locally averaged divergence of the stress tensor (first term on right hand side). The last term in the particle equation $m_f\boldsymbol{a}_{\boldsymbol{J}}$ has a less-clear physical interpretation and comes because the fluid is allowed to have a local acceleration different from the particle. It is expected that at small Reynolds numbers the velocity field will be smooth at the scale of the particle size and thus $\boldsymbol{a}_{\boldsymbol{J}}\approx\boldsymbol{0}$ [66]. Nevertheless, we will retain the terms involving $\boldsymbol{a}_{\boldsymbol{J}}$ to ensure a consistent formulation, see Appendix B.

B. Momentum Formulation

Let us define a *momentum field* as the sum of the fluid momentum and the spreading of the particle momentum,

$$\boldsymbol{p}\left(\boldsymbol{r},t\right)=\rho\boldsymbol{v}+m_e\boldsymbol{S}\boldsymbol{u}=\left(\rho+m_e\boldsymbol{SJ}\right)\boldsymbol{v}. \tag{15}$$

The total momentum is $\boldsymbol{P}(t) = \int \boldsymbol{p}(\boldsymbol{r},t)\,d\boldsymbol{r}$ and therefore a local conservation law for $\boldsymbol{p}(\boldsymbol{r},t)$ implies conservation of the total momentum.

By adding the fluid and particle equations (9,10) together we can obtain the dynamics of the momentum field,

$$\begin{aligned}\partial_t \boldsymbol{p} &= \rho\left(\partial_t \boldsymbol{v}\right) + m_e \boldsymbol{S}\dot{\boldsymbol{u}} + m_e\left(\boldsymbol{u}\cdot\frac{\partial}{\partial \boldsymbol{q}}\boldsymbol{S}\right)\boldsymbol{u} \\ &= -\boldsymbol{\nabla}\pi + \boldsymbol{\nabla}\cdot\boldsymbol{\sigma} - \boldsymbol{\nabla}\cdot\left[\rho\boldsymbol{v}\boldsymbol{v}^T + m_e\boldsymbol{S}\left(\boldsymbol{u}\boldsymbol{u}^T\right)\right] + \boldsymbol{S}\boldsymbol{F}, \end{aligned} \tag{16}$$

where we used the fact that $\boldsymbol{S}$ depends on the difference $(\boldsymbol{q}-\boldsymbol{r})$ only, and not on $\boldsymbol{q}$ and $\boldsymbol{r}$ separately. In the absence of applied external forces we can write the right hand side as a divergence of a total stress tensor $-\pi\boldsymbol{I}+\boldsymbol{\sigma}+\boldsymbol{\sigma}_{\rm kin}$, where the kinetic stress tensor includes a contribution from the inertia of the particle,

$$\boldsymbol{\sigma}_{\rm kin} = -\rho\boldsymbol{v}\boldsymbol{v}^T - m_e\boldsymbol{S}\left(\boldsymbol{u}\boldsymbol{u}^T\right). \tag{17}$$

This means that the momentum field obeys a local conservation law, as expected for short-ranged interactions between the particle and the fluid molecules.

The formulation (16) is not only informative from a physical perspective, but was also found very useful in performing adiabatic elimination in the case of frictional coupling in Refs. [27, 46].

C. Pressure-Free Formulation

The equations we wrote so far contain the $\boldsymbol{\nabla}\pi$ term and can easily be generalized to the case of a compressible fluid [43]. For analysis purposes, in the incompressible case it is useful to eliminate the pressure from the equations using a projection operator formalism. This well-known procedure [67] can be understood as follows. The fluid equation (9) is of the form

$$\partial_t\boldsymbol{v} + \rho^{-1}\boldsymbol{\nabla}\pi = \boldsymbol{g}, \quad \boldsymbol{\nabla}\cdot\boldsymbol{v} = 0.$$

By taking the divergence of the evolution equation, we get

$$\partial_t\left(\boldsymbol{\nabla}\cdot\boldsymbol{v}\right) + \rho^{-1}\boldsymbol{\nabla}^2\pi = \rho^{-1}\boldsymbol{\nabla}^2\pi = \boldsymbol{\nabla}\cdot\boldsymbol{g},$$

which is a Poisson equation for the pressure whose solution can be formally written as $\pi = \rho\boldsymbol{\nabla}^{-2}\left(\boldsymbol{\nabla}\cdot\boldsymbol{g}\right)$. This means that

$$\partial_t\boldsymbol{v} = \boldsymbol{g} - \rho^{-1}\boldsymbol{\nabla}\pi = \boldsymbol{g} - \boldsymbol{\nabla}\left[\boldsymbol{\nabla}^{-2}\left(\boldsymbol{\nabla}\cdot\boldsymbol{g}\right)\right] = \boldsymbol{\mathcal{P}}\boldsymbol{g},$$

where $\boldsymbol{\mathcal{P}}$ is a *projection operator* that projects the right-hand side $\boldsymbol{g}$ or a given velocity field onto the space of divergence-free vector fields. Note that the boundary conditions are implicit in the definitions of the gradient, divergence, and Laplacian operators. For periodic boundary conditions, the projection can most easily be implemented using a spatial Fourier transform. Specifically, in Fourier space the projection operator is simply a multiplication by the $d\times d$ matrix $\widehat{\boldsymbol{\mathcal{P}}} = \boldsymbol{I} - k^{-2}(\boldsymbol{k}\boldsymbol{k}^T)$, where d is the dimensionality, $d=2$ or $d=3$, and $\boldsymbol{k}$ is the wavenumber.

By using the projection operator, we can eliminate the pressure from the equations of motion (9,10), to obtain

$$\rho\partial_t\boldsymbol{v} = \boldsymbol{\mathcal{P}}\left[-\rho\boldsymbol{v}\cdot\boldsymbol{\nabla}\boldsymbol{v} + \boldsymbol{\nabla}\cdot\boldsymbol{\sigma} - m_e\boldsymbol{S}\dot{\boldsymbol{u}} + \boldsymbol{S}\boldsymbol{F}\right].$$

If we now use (7) to eliminate $\dot{\boldsymbol{u}}$ we obtain the fluid equation

$$\rho\partial_t\boldsymbol{v} + m_e\boldsymbol{\mathcal{P}}\boldsymbol{S}\boldsymbol{J}\boldsymbol{\mathcal{P}}\left(\partial_t\boldsymbol{v}\right) = \boldsymbol{\mathcal{P}}\left[-\rho\boldsymbol{v}\cdot\boldsymbol{\nabla}\boldsymbol{v} - m_e\boldsymbol{S}\left(\boldsymbol{u}\cdot\frac{\partial}{\partial\boldsymbol{q}}\boldsymbol{J}\right)\boldsymbol{v} + \boldsymbol{\nabla}\cdot\boldsymbol{\sigma} + \boldsymbol{S}\boldsymbol{F}\right],$$

where we used the fact that $\boldsymbol{\mathcal{P}}\boldsymbol{v} = \boldsymbol{v}$ since $\boldsymbol{\nabla}\cdot\boldsymbol{v} = 0$, and we added a $\boldsymbol{\mathcal{P}}$ in front of the second term for symmetry purposes. This shows that the pressure-free form of the fluid-only equation (13) is

$$\boldsymbol{\rho}_{\rm eff}\partial_t\boldsymbol{v} = \boldsymbol{\mathcal{P}}\left\{-\left[\rho\left(\boldsymbol{v}\cdot\boldsymbol{\nabla}\right) + m_e\boldsymbol{S}\boldsymbol{J}\left(\boldsymbol{v}\cdot\frac{\partial}{\partial\boldsymbol{q}}\boldsymbol{J}\right)\right]\boldsymbol{v} + \boldsymbol{\nabla}\cdot\boldsymbol{\sigma}\right\} + \boldsymbol{\mathcal{P}}\boldsymbol{S}\boldsymbol{F} = \boldsymbol{\mathcal{P}}\boldsymbol{f} + \boldsymbol{\mathcal{P}}\boldsymbol{S}\boldsymbol{F}, \tag{18}$$

where the force density $\boldsymbol{f}$ contains the advective, viscous and stochastic contributions to the fluid dynamics. This form of the equation of motion can be shown to be identical to the limiting equation for velocity obtained by Tabak and Atzberger, with the exception of the advective term $\boldsymbol{v}\cdot\boldsymbol{\nabla}\boldsymbol{v}$ which is omitted in Ref. [46]. An important feature of this formulation is that the density ρ in the usual Navier-Stokes equation is now replaced by the effective density *operator*

$$\boldsymbol{\rho}_{\rm eff} = \rho\boldsymbol{I} + m_e\boldsymbol{\mathcal{P}}\boldsymbol{S}\boldsymbol{J}\boldsymbol{\mathcal{P}}, \tag{19}$$

where $\boldsymbol{I}$ is the identity operator or matrix. Notice that the effective density operator for incompressible flow is not $\rho\boldsymbol{I} + m_e\boldsymbol{SJ}$ as one might naively expect based on (13). The distinction between $\boldsymbol{SJ}$ and $\boldsymbol{\mathcal{P}SJ\mathcal{P}}$ is important, and leads to a well-known but surprising difference between the short-time motion of a particle immersed in a compressible versus an incompressible fluid [68]. When there are many particles present, the effective inertia tensor is generalized straightforwardly by summing the added inertia over the particles,

$$\boldsymbol{\rho}_{\text{eff}} = \rho\boldsymbol{I} + \sum_i \left(m_e\right)_i \boldsymbol{\mathcal{P}S}_i\boldsymbol{J}_i\boldsymbol{\mathcal{P}}. \tag{20}$$

Equation (18) together with the no-slip condition $\dot{\boldsymbol{q}} = \boldsymbol{Jv}$ gives a closed set of equations for $\boldsymbol{v}$ and $\boldsymbol{q}$ without any constraints. We can use this unconstrained formulation to simplify analysis of the properties of the coupled fluid-particle problem. In particular, in Appendix B, the constraint-free form is used for showing fluctuation-dissipation balance in the stochastic setting. The formal solution of (18),

$$\partial_t \boldsymbol{v} = \boldsymbol{\rho}_{\text{eff}}^{-1}\boldsymbol{\mathcal{P}}\left(\boldsymbol{f} + \boldsymbol{SF}\right),$$

involves the cumbersome operator $\boldsymbol{\mathcal{P}SJ\mathcal{P}}$ via the inverse of the effective inertia $\boldsymbol{\rho}_{\text{eff}}$. In principle this makes the pressure-free formulation difficult to use in numerical methods. However, as we explain in Appendix A, with periodic boundary conditions it is possible to efficiently invert the operator $\boldsymbol{\rho}_{\text{eff}}$ using Fourier transforms and thus obtain a closed-form equation (A6) for $\partial_t\boldsymbol{v}$ suitable for numerical implementations. This relies on the fact that for a large d-dimensional periodic system $\boldsymbol{J\mathcal{P}S}$ is a constant multiple of the $d \times d$ identity matrix.

D. Energy Conservation

A crucial property of no-slip coupling (unlike frictional-coupling) between the particles and the flow is that it is non-dissipative, and therefore all dissipation (drag) comes from viscous dissipation. Specifically, in the absence of viscous dissipation, the equations of motion (9,10,11) conserve a coarse-grained Hamiltonian [69, 70] given by the sum of potential energy and the kinetic energy of the particle and the fluid,

$$H\left(\boldsymbol{v}, \boldsymbol{u}, \boldsymbol{q}\right) = \rho\int \frac{v^2}{2}\, d\boldsymbol{r} + m_e\frac{u^2}{2} + U\left(\boldsymbol{q}\right), \tag{21}$$

where $U\left(\boldsymbol{q}\right)$ is the interaction potential of the particle with external sources and other particles, with an associated conservative force

$$\boldsymbol{F}\left(\boldsymbol{q}\right) = -\frac{\partial U}{\partial \boldsymbol{q}} = -\frac{\partial H}{\partial \boldsymbol{q}}.$$

For compressible flow one needs to include the (density-dependent) internal energy of the fluid in the Hamiltonian as well [71].

To demonstrate energy conservation, we calculate the rate of change

$$\frac{dH}{dt} = -\boldsymbol{F}\cdot\boldsymbol{u} + m_e\boldsymbol{u}\cdot\dot{\boldsymbol{u}} + \int \rho\boldsymbol{v}\cdot\left(\partial_t\boldsymbol{v}\right)\, d\boldsymbol{r}$$

in the absence of viscous and stochastic fluxes. Using the equations of motion (9,10) we get

$$\begin{aligned}
\frac{dH}{dt} &= -\boldsymbol{F}\cdot\boldsymbol{u} + \boldsymbol{u}\cdot\left(\boldsymbol{F}+\boldsymbol{\lambda}\right) - \int \boldsymbol{v}\cdot\left(\boldsymbol{S\lambda}\right)\, d\boldsymbol{r} \\
&\quad - \int \boldsymbol{v}\cdot\boldsymbol{\nabla}\pi\, d\boldsymbol{r} - \rho\int \boldsymbol{v}\cdot\left(\boldsymbol{v}\cdot\boldsymbol{\nabla}\boldsymbol{v}\right)\, d\boldsymbol{r} \\
&= \left(\boldsymbol{u}-\boldsymbol{Jv}\right)\cdot\boldsymbol{\lambda} + \int \pi\left(\boldsymbol{\nabla}\cdot\boldsymbol{v}\right)\, d\boldsymbol{r} \\
&\quad -\frac{\rho}{2}\int \boldsymbol{v}\cdot\boldsymbol{\nabla}\left(v^2\right)\, d\boldsymbol{r} + \rho\int \boldsymbol{v}\cdot\left[\left(\boldsymbol{\nabla}\times\boldsymbol{v}\right)\times\boldsymbol{v}\right]\, d\boldsymbol{r},
\end{aligned}$$

where integration by parts and the adjoint property (3) were used for the first two terms, and a vector identity was used to express $\boldsymbol{v}\cdot\boldsymbol{\nabla}\boldsymbol{v}$ in terms of the vorticity $\boldsymbol{\nabla}\times\boldsymbol{v}$. The first term vanishes due to the no-slip constraint $\boldsymbol{u} = \boldsymbol{Jv}$. The second and third terms vanish for incompressible flow $\boldsymbol{\nabla}\cdot\boldsymbol{v}$, and the last term vanishes because of the basic properties of the cross product. This demonstrates that $dH/dt = 0$ in the absence of viscous dissipation, that is, the non-dissipative terms in the equation strictly conserve the coarse-grained free energy.

E. Fluctuation-Dissipation Balance

So far, we considered the equations of motion for the fluid-particle system ignoring thermal fluctuations. In Appendix B we formally demonstrate that in order to account for thermal fluctuations in a manner that preserves fluctuation-dissipation balance it is sufficient to add the usual Landau-Lifshitz stochastic stress $(k_B T\eta)^{1/2}\left(\boldsymbol{\mathcal{W}}+\boldsymbol{\mathcal{W}}^T\right)$ to the viscous stress tensor in $\boldsymbol{\sigma}$, without adding any stochastic forces on the particle. The key physical insight is that the fluid-particle coupling is non-dissipative, as demonstrated in Section II D, and the only dissipation comes from the viscous terms.

Fluctuation-dissipation balance here means that at thermodynamic equilibrium the particle-fluid system is ergodic and time-reversible with respect to the Gibbs-Boltzmann distribution $Z^{-1}\exp\left(-H/k_B T\right)$, where the "Hamiltonian" H given in (21) is to be interpreted as a coarse-grained free energy. Since $\boldsymbol{u}=\boldsymbol{J}\boldsymbol{v}$ is not an independent degree of freedom, we can formally write the Hamiltonian in terms of the degrees of freedom of the system as a sum of potential and kinetic energy. The total kinetic energy includes, in addition to the kinetic energy of the fluid $\int(\rho/2)\,v^2\,d\boldsymbol{r}$, a kinetic energy contribution due to the motion of the particle,

$$E_p = m_e\frac{u^2}{2} = m_e\frac{(\boldsymbol{J}\boldsymbol{v})\cdot\boldsymbol{u}}{2} = \frac{m_e}{2}\int \boldsymbol{v}\cdot(\boldsymbol{S}\boldsymbol{u})\;d\boldsymbol{r} = \frac{m_e}{2}\int \boldsymbol{v}^T\boldsymbol{S}\boldsymbol{J}\boldsymbol{v}\,d\boldsymbol{r},$$

where use was made of the adjoint condition (3). This leads to the coarse-grained Hamiltonian

$$H(\boldsymbol{v},\boldsymbol{q}) = m_e\int\frac{\boldsymbol{v}^T\boldsymbol{S}\boldsymbol{J}\boldsymbol{v}}{2}\,d\boldsymbol{r}+\rho\int\frac{v^2}{2}\,d\boldsymbol{r}+U(\boldsymbol{q}) = \frac{1}{2}\int\boldsymbol{v}^T\boldsymbol{\rho}_{\text{eff}}\boldsymbol{v}\,d\boldsymbol{r}+U(\boldsymbol{q})\,. \tag{22}$$

Note that it is not necessary here to include an entropic contribution to the coarse-grained free energy because our formulation is isothermal, and we assume that the particles do not have internal structure.

We emphasize that the form of $H(\boldsymbol{q},\boldsymbol{v})$ in (22) is postulated based on physical reasoning rather than derived from a more refined model. Atzberger *et al.* [27, 46] provide a careful and detailed discussion of how one might eliminate the particle velocity $\boldsymbol{u}$ by performing an adiabatic elimination, starting from a frictional coupling model in which the particle velocity is independent from the fluid velocity. The starting frictional coupling model, however, as we discussed earlier, involves an arbitrary frictional force parameter that is yet to be given a microscopic interpretation. We believe that the consistency of our inertial coupling model with general thermodynamic principles and deterministic hydrodynamics is sufficient to adopt the inertial coupling model as a consistent coarse-grained model without having to justify it from "first principles" [72]. We also note that the equations obtained in Ref. [46] are similar in structure to the ones we have presented here.

The fact that fluid-particle coupling conserves the Hamiltonian (see Section II D) and is therefore non-dissipative is a crucial component of fluctuation-dissipation balance. However, this is not sufficient on its own. An important additional requirement is that the phase space dynamics should be incompressible, which means that the dynamics preserves not just phase-space functions of H, but also preserves phase-space measures such as the Gibbs-Boltzmann distribution. As discussed by Atzberger [27], even for the case of a neutrally-buoyant particle an additional "Ito" or "thermal" drift term needs to be added to the velocity equation to ensure fluctuation-dissipation balance. This term has the form of an additional contribution to the stress tensor

$$\boldsymbol{\sigma}_{\text{th}} = -\left[\boldsymbol{S}(k_B T)\right]\boldsymbol{I}. \tag{23}$$

The physical origin of this term is the Kirkwood kinetic stress due to the thermal motion of the particle lost when eliminating $\boldsymbol{u}$ as a degree of freedom. Another way to interpret this term is that it adds a particle contribution of $\boldsymbol{S}(k_B T)$ to the pressure. For incompressible flow, this simply changes the pressure but does not change the dynamics of the velocity field since the projection $\boldsymbol{\mathcal{P}}$ eliminates the scalar gradient term $\boldsymbol{\nabla}\cdot\boldsymbol{\sigma}_{\text{th}} = -\boldsymbol{\nabla}\boldsymbol{S}(k_B T)$. In Appendix B we argue that for non-neutrally buoyant particles there is also no need to include an additional thermal drift term for periodic boundary conditions. Note, however, that the above calculations rely on continuum identities that fail to be strictly obeyed discretely. As we explain in Section B 2 of the Appendix in more detail, ensuring strict *discrete* fluctuation-dissipation balance requires keeping the contribution $-\boldsymbol{\nabla}\boldsymbol{S}(k_B T)$ in the momentum equation in *both* the compressible and incompressible settings.

F. Equipartition of Energy

For a single particle immersed in a periodic incompressible fluid in d dimensions, translational invariance implies that there is no dependence of expectation values on $\boldsymbol{q}$, and therefore we can keep $\boldsymbol{q}$ fixed when calculating expectation values. The fact that the Hamiltonian (22) is quadratic in $\boldsymbol{v}$ means that the fluctuations of velocity are Gaussian with covariance $\langle \boldsymbol{v}\boldsymbol{v}^\star \rangle = (k_B T)\, \boldsymbol{\rho}_{\text{eff}}^{-1}$. The fluctuations of the particle velocity have variance

$$\langle u^2 \rangle = \text{Trace}\left[\boldsymbol{J}\, \langle \boldsymbol{v}\boldsymbol{v}^\star \rangle\, \boldsymbol{S}\right] = (k_B T)\, \text{Trace}\left[\boldsymbol{J}\boldsymbol{\rho}_{\text{eff}}^{-1}\boldsymbol{S}\right].$$

Using the relations (A5) and (A3) derived in the Appendix, we can simplify

$$\langle u^2 \rangle = \frac{k_B T}{\rho}\text{Trace}\left[\boldsymbol{J}\left(\boldsymbol{I} - \frac{m_e \Delta\widetilde{V}}{\tilde{m}}\boldsymbol{\mathcal{P}}\boldsymbol{S}\boldsymbol{J}\boldsymbol{\mathcal{P}}\right)\boldsymbol{S}\right] = d\frac{k_B T}{\tilde{m}}, \tag{24}$$

where $\Delta\widetilde{V} = d\,\Delta V/(d-1)$ and $\tilde{m} = m_e + d\,m_f/(d-1)$.

The result $\langle u^2 \rangle = d\,(k_B T)/\tilde{m}$ should be compared to the corresponding result for a compressible fluid [43], $\langle u^2 \rangle = d\,(k_B T)/m$, which follows from the usual equipartition principle of statistical mechanics. When incompressibility is accounted for, a fraction of the equilibrium kinetic energy is carried in the unresolved sound waves, and therefore the apparent mass of the particle is $\tilde{m}$ and not $m = m_e + m_f$, as we verify numerically in Section IV B. It is reassuring that our model equations reproduce the result for rigid particles. This suggests that the model introduced here can be used to study more complicated questions such as the effect of multi-particle interactions on $\langle u^2 \rangle$ in semi-dilute to dense colloidal suspensions [73].

III. Spatio-Temporal Discretization

In this section we describe our second-order spatio-temporal discretization of the equations of motion (9,10,11). Our spatio-temporal discretization is based on the deterministic Immersed Boundary Method (IBM), and in particular, on the deterministic second-order temporal integrator presented in Ref. [**?**]. For the fluctuating fluid solver, we use the second-order staggered-grid spatial discretization of the fluctuating Navier-Stokes equations described in detail in Ref. [44]. A notable feature of the fluid solver we employ is that it handles the viscous terms semi-implicitly and is stable for large time steps. Furthermore, for the fluctuating Stokes equations, our fluid solver produces the correct spectrum of the velocity fluctuations for any time step size [44].

There are two key novel features in our incompressible inertial coupling algorithm from those previously developed. Firstly, our algorithm includes the effects of particle excess inertia in a manner that strictly conserves momentum and is second-order deterministically for smooth problems. Secondly, we focus our initial development on systems with periodic boundaries only, allowing the use of the Fast Fourier Transform (FFT) as a linear solver for the time-dependent Stokes equations. This greatly simplifies the implementation of the algorithm and allows us to use Graphics Processing Units (GPUs) for very efficient parallelization of the algorithm.

For neutrally-buoyant particles our temporal discretization is exactly that described in Ref. [**?**] with the fluid solver replaced by that described in Ref. [44]. This simplified algorithm is already implemented in the IBAMR software framework [74], an open-source library for developing fluid-structure interaction models that use the immersed boundary method. Note that IBAMR can handle non-periodic boundary conditions using a preconditioned iterative solver for the time-dependent Stokes equations [75]. In this paper we focus on describing the additional steps required to handle excess inertia and to use FFTs as a linear solver.

The majority of our presentation will focus on a single particle coupled to a fluctuating fluid. Only small changes are required to handle multiple particles by simply summing the single-particle term over the different particles. As we explain in more detail shortly, the error introduced by superposing the single particle solutions to solve the multi-particle system is small if the kernels of the different particles are not overlapping, which in practice means that there are at least 3 grid cells between the centroids of the particles.

A. Spatial Discretization

Our second-order spatial discretization of the equations (9,10,11) is based on standard techniques for incompressible flow and the immersed boundary method [32], as described in more detail in, for example, Ref. [**?**]. In the spatially-discretized equations, the same equations as for the continuum apply, but with

the interpretation that the velocity $\boldsymbol{v}$ is not a (random) vector field but rather a finite-volume discretization of that field [53, 76]. We use a uniform Cartesian staggered-grid spatial discretization of the incompressible Navier-Stokes equations, as described in more detail in Ref. [44]. In the staggered discretization the control volume grid associated to each component of velocity is shifted by half a grid spacing along the corresponding dimension relative to the pressure grid. In the discrete setting, the various continuum operators acting on vector fields become matrices. The spatial discretization of the differential operators, notably the discrete gradient, divergence and Laplacian operators, is described in detail in Ref. [44].

1. Discrete Interpolation and Spreading Operators

Application of the local averaging operator $\boldsymbol{J}$, which is a convolution operator in the continuum setting, becomes a discrete summation over the grid points that are near the particle,

$$\boldsymbol{J}\boldsymbol{v} \equiv \sum_{\boldsymbol{k}\in\text{grid}} \phi_a\left(\boldsymbol{q}-\boldsymbol{r}_k\right)\boldsymbol{v}_k,$$

where $\boldsymbol{r}_k$ denotes the center of the control volume with which $\boldsymbol{v}_k$ is associated, and ϕ_a is a function that takes the role of the kernel function δ_a. We follow the traditional choice [32] and do the local averaging independently along each dimension,

$$\phi_a\left(\boldsymbol{q}-\boldsymbol{r}_k\right) = \prod_{\alpha=1}^{d} \phi_a\left[q_\alpha - \left(r_k\right)_\alpha\right],$$

which improves the isotropy of the spatial discretization (but note that the local averaging is not rotationally invariant). As a matrix, the local spreading operator $\boldsymbol{S} = \left(\Delta V_f\right)^{-1}\boldsymbol{J}^\star$ is a weighted transpose of $\boldsymbol{J}$,

$$\left(\boldsymbol{S}\boldsymbol{F}\right)_k = \left(\Delta V_f\right)^{-1}\phi_a\left(\boldsymbol{q}-\boldsymbol{r}_k\right)\boldsymbol{F},$$

where $\Delta V_f = \Delta x\Delta y\Delta z$ is the volume of the hydrodynamic cell [2].

The discrete kernel function ϕ_a was constructed by Peskin [32] to yield translationally-invariant zeroth- and first-order moment conditions, along with a quadratic condition,

$$\begin{aligned}
\sum_{\boldsymbol{k}\in\text{grid}} \phi_a\left(\boldsymbol{q}-\boldsymbol{r}_k\right) &= 1 \\
\sum_{\boldsymbol{k}\in\text{grid}} \left(\boldsymbol{q}-\boldsymbol{r}_k\right)\phi_a\left(\boldsymbol{q}-\boldsymbol{r}_k\right) &= 0 \\
\sum_{\boldsymbol{k}\in\text{grid}} \phi_a^2\left(\boldsymbol{q}-\boldsymbol{r}_k\right) &= \Delta V^{-1} = \text{const.},
\end{aligned} \tag{25}$$

independent of the position of the particle $\boldsymbol{q}$ relative to the underlying (fixed) velocity grid. Ensuring these properties requires relating the support of the kernel function to the grid spacing, that is, making $a \sim \Delta x$ (more specifically, typically the width of the function ϕ_a has to be an integer multiple of the grid spacing). This means that the size and shape of the particles is directly tied to the discretization of the fluid equations, and the two cannot be varied independently, for example, simulating the motion of a "spherical" particle requires choosing the same grid spacing along each dimension, $\Delta x = \Delta y = \Delta z$. This is a shortcoming of our method, but, at the same time, it is physically unrealistic to resolve the fluid flow and, in particular, the fluctuations in fluid velocity, with different levels of resolution for different particles or dimensions.

The physical size of the particle (hydrodynamic radius) can be varied over a certain range independently of the grid resolution by using modified discrete kernel functions. This can be very useful when simulating polydisperse suspensions with mild polydispersity. A simple approach is to use shifted or *split kernels*,

$$\phi_{a,s}\left(\boldsymbol{q}-\boldsymbol{r}_k\right) = \frac{1}{2^d}\prod_{\alpha=1}^{d}\left\{\phi_a\left[q_\alpha - \left(r_k\right)_\alpha - \frac{s}{2}\right] + \phi_a\left[q_\alpha - \left(r_k\right)_\alpha + \frac{s}{2}\right]\right\},$$

where s denotes a shift that parametrizes the kernel. By varying s in a certain range, for example, $0 \le s \le \Delta x$, one can smoothly increase the support of the kernel and thus increase the hydrodynamic

[2] The cell volume ΔV_f is introduced here because the fluid kinetic energy appearing in the discrete Hamiltonian is $\Delta V_f \sum_{\boldsymbol{k}\in\text{grid}} \rho v_k^2/2$ and therefore $\partial H/\partial \boldsymbol{v}$ is not the functional derivative $\rho\boldsymbol{v}$ as in the continuum (see Appendix B) but rather the partial derivative $\Delta V_f\,\rho\boldsymbol{v}$.

radius of the blob by as much as a factor of two. While in principle one can increase the width of the kernels arbitrarily, if the support of the kernel grows beyond 5-6 cells it better to abandon the minimally-resolved blob approach and employ a more resolved representation of the particle [14, 39]. We do not use split kernels in this work but have found them to work as well as the unshifted kernels, while allowing increased flexibility in varying the grid spacing relative to the hydrodynamic radius of the particles.

The last condition (25), was imposed by Peskin [32] as a way of approximating independence under shifts of order of the grid spacing. This property is especially important in our context since it implies that the particle volume $\Delta V = (\boldsymbol{J}\boldsymbol{S}\,1)^{-1}$ will remain constant and independent of the position of the blob relative to the underlying grid. The function with minimal support that satisfies (25) is uniquely determined [32]. In our numerical experiments we employ this *three-point* discrete kernel function, which means that the support of ϕ extends to only three grid points along each dimension (i.e., 3^d discrete velocities are involved in the averaging and spreading operations in d dimensions), see Ref. [43] for details. This particular choice gives $\Delta V = 8\Delta V_f$ in three dimensions and $\Delta V = 4\Delta V_f$ in two dimensions. The narrow kernel improves the computational efficiency on the bandwidth-limited GPU, as detailed in Ref. [43].

2. Translational Invariance

In the continuum derivation, obtaining the closed-form pressure-free velocity equation (A6) relied sensitively on the fact that for a large continuum system

$$\boldsymbol{J}\boldsymbol{\mathcal{P}}\boldsymbol{S} = \frac{d-1}{d}\Delta V^{-1}\,\boldsymbol{I}, \tag{26}$$

see Eq. (A3) in Appendix A. Ideally, we would like the spatial discretization to have the additional property that, for periodic boundary conditions, $\boldsymbol{J}\boldsymbol{\mathcal{P}}\boldsymbol{S}$ should be invariant under translations of the particle relative to the underlying fluid grid. This is *not* ensured by the Peskin operators, which are constructed without any reference to the fluid equations and the particular form of the discrete projection or the discrete viscous dissipation. In fact, in the traditional immersed boundary method [32], a centered discretization of the velocity was used, which implies a very different form for the discrete projection operator $\boldsymbol{\mathcal{P}}$ and required the introduction of an additional "odd-even" moment condition not strictly necessary with a staggered discretization.

Numerical experiments suggest that for the staggered grid discretization that we employ, the continuum identity (26) is obeyed to within a maximal deviation of a few percent,

$$\Delta\widetilde{\boldsymbol{V}}^{-1}(\boldsymbol{q}) = \boldsymbol{J}\boldsymbol{\mathcal{P}}\boldsymbol{S} \approx \Delta\widetilde{V}^{-1}\,\boldsymbol{I} \tag{27}$$

for a sufficiently large periodic system, where $\Delta\widetilde{V} = d\Delta V/(d-1)$ is a modified volume of the blob. For a periodic three-dimensional system of $N_x \times N_y \times N_z$ hydrodynamic cells, we expect to see deviations from (27) if one of the grid dimensions becomes of the order of the kernel width (which is 3 cells for our spatial discretization). This is because a particle then becomes affected by its nearby periodic images. For a grid size $N_x \times N_y \times 1$ cells we expect to obtain two-dimensional behavior [see Eq. (A4)].

In the left panel of Fig. 1 we show numerical results for the average and maximum deviation of $\Delta\widetilde{V}(\boldsymbol{J}\boldsymbol{\mathcal{P}}\boldsymbol{S})$ from the $d \times d$ identity matrix,

$$\delta\boldsymbol{I}(\boldsymbol{q}) = \Delta\widetilde{V}^{-1}(\boldsymbol{J}\boldsymbol{\mathcal{P}}\boldsymbol{S}) - \boldsymbol{I},$$

as we vary the position of the particle $\boldsymbol{q}$ relative to a cubic Eulerian grid of size $N_x = N_y = N_z = N$ cells. Specifically, we show the average diagonal value, the maximum diagonal element, and the maximum off-diagonal element of $\delta\boldsymbol{I}(\boldsymbol{q})$. For all but the smallest systems the diagonal elements are smaller than 5% and the off-diagonal elements are on the order of 0.1%. For smaller system sizes there are visible finite-size effects due to interactions with periodic images. For comparison, we also show the corresponding two-dimensional results for a square grid of $N_x = N_x = N$ cells. The finite-size effects are more pronounced in two dimensions due to the slower decay of the Green's function for the Poisson equation, but for systems larger than $N = 32$ cells we find (26) to hold to a percent or so. In the left panel of Fig. 1 we also show the average diagonal values $(\delta\boldsymbol{I}_{xx} + \delta\boldsymbol{I}_{yy})/2$ and $\delta\boldsymbol{I}_{zz}$ for non-cubic systems, illustrating the change from three-dimensional to the two-dimensional behavior as $N_z \to 1$.

Given these numerical observations, we will make an approximation and assume that $\boldsymbol{J}\boldsymbol{\mathcal{P}}\boldsymbol{S}$ is a constant multiple of the $d \times d$ identity matrix independent of $\boldsymbol{q}$, and (27) holds as an equality. While this is, in principle, an uncontrolled approximation, we will correct for it in order to ensure strict momentum conservation and thus strict adherence to fundamental physical laws. We also

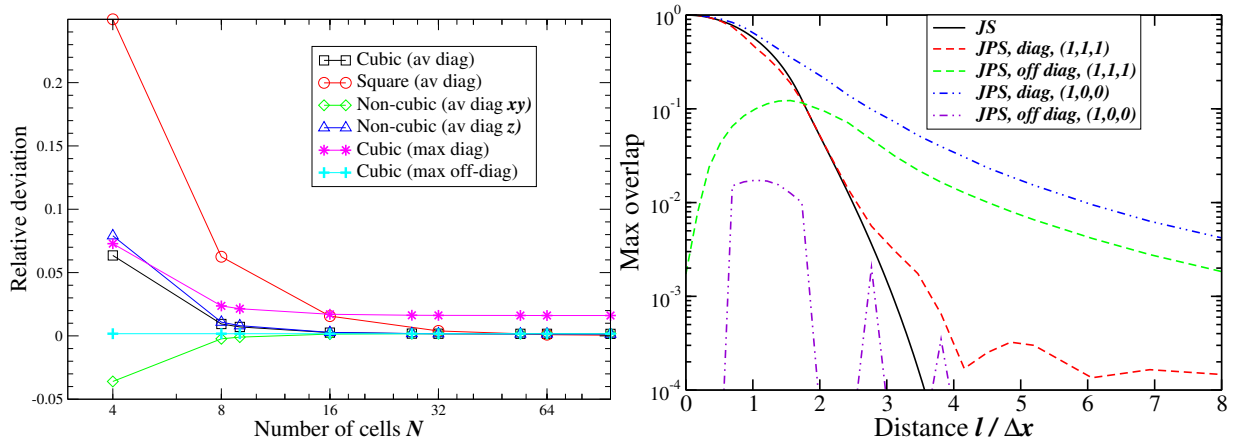

Figure 1: (*Left panel*) Translational invariance of the approximation $d\Delta V\left(\boldsymbol{J}\boldsymbol{\mathcal{P}}\boldsymbol{S}\right)/\left(d-1\right)\approx\boldsymbol{I}$ over a set of 10^3 positions of the particle relative to the underlying fluid grid. For a periodic three-dimensional system of N^3 cells, we show the average diagonal value of $\delta\boldsymbol{I}$ (black squares), the maximum diagonal element of $\delta\boldsymbol{I}$ (magenta stars), as well as the maximum off-diagonal element of $\delta\boldsymbol{I}$ (cyan pluses). We also show the average diagonal value of $\delta\boldsymbol{I}$ for a two-dimensional system of N^2 cells (red circles), as well as the average value $(\delta\boldsymbol{I}_{xx}+\delta\boldsymbol{I}_{yy})/2$ (green diamonds) and $\delta\boldsymbol{I}_{zz}$ (blue triangles) for a three-dimensional system of $32\times32\times N$ cells. (*Right panel*) Maximum diagonal and off-diagonal value of $3\Delta V\left(\boldsymbol{J}_i\boldsymbol{\mathcal{P}}\boldsymbol{S}_j\right)/2$ and the maximum diagonal value of $\Delta V\left(\boldsymbol{J}_i\boldsymbol{S}_j\right)$ for two particles i and j a distance l apart along the $(1,0,0)$ or $(1,1,1)$ direction, in a three-dimensional periodic box of 32^3 cells.

note that the small approximations we make when computing $\boldsymbol{\rho}_{\text{eff}}^{-1}$ only affect the short-time inertial dynamics and do *not* affect static properties or the long-time dynamics, as deduced from (18). The approximation (27) allows us to write the discrete equivalent of (A5)

$$\boldsymbol{\rho}_{\text{eff}}^{-1}=\left(\rho\boldsymbol{I}+m_e\boldsymbol{\mathcal{P}}\boldsymbol{S}\boldsymbol{J}\boldsymbol{\mathcal{P}}\right)^{-1}\approx\rho^{-1}\left(\boldsymbol{I}-\frac{m_e\Delta\widetilde{V}}{\tilde{m}}\boldsymbol{\mathcal{P}}\boldsymbol{S}\boldsymbol{J}\boldsymbol{\mathcal{P}}\right). \tag{28}$$

When there are multiple particles in the system we simply sum the second term on the right hand side over all particles, see Eq. (A7), which is an approximation even in the continuum setting. This approximation relies on the assumption that $\boldsymbol{J}_i\boldsymbol{\mathcal{P}}\boldsymbol{S}_j\approx\boldsymbol{0}$ for two particles i and j that are far away from each other. To test the validity of this approximation, in the right panel of Fig. 1 we show the maximum diagonal and off-diagonal value for $\boldsymbol{J}_i\boldsymbol{\mathcal{P}}\boldsymbol{S}_j$ for two particles a distance l away from each other, for both $\boldsymbol{q}_i-\boldsymbol{q}_j=l\left(1,0,0\right)$ and $\boldsymbol{q}_i-\boldsymbol{q}_j=l\left(\sqrt{3},\sqrt{3},\sqrt{3}\right)/3$, in three dimensions. If $l=0$ then $\boldsymbol{J}_i\boldsymbol{\mathcal{P}}\boldsymbol{S}_j$ approaches $2\Delta V^{-1}\boldsymbol{I}/3$, so we normalize the value of $\boldsymbol{J}_i\boldsymbol{\mathcal{P}}\boldsymbol{S}_j$ by a factor of $3\Delta V/2$. We see that for $l\gtrsim5\Delta x$, the maximum normalized value of $\boldsymbol{J}_i\boldsymbol{\mathcal{P}}\boldsymbol{S}_j$ is less than 0.01. As seen in the figure, $\boldsymbol{J}_i\boldsymbol{S}_j$ vanishes identically if the kernels of particles i and j are disjoint, which is always the case when the distance between the blobs is larger than $3\Delta x$ along at least one direction.

Another operator that is important for the long-time diffusive (Brownian) motion of the particle, as explained in Refs. [3, 27], is the discrete mobility tensor for a single particle,

$$\boldsymbol{\mu}\left(\boldsymbol{q}\right)=-\boldsymbol{J}\boldsymbol{\mathcal{L}}^{-1}\boldsymbol{S},$$

where $\boldsymbol{\mathcal{L}}^{-1}$ denotes the discrete Stokes solution operator. For a single particle in a periodic domain, we numerically find that $\boldsymbol{\mu}\left(\boldsymbol{q}\right)$ is approximately translationally-invariant to within a few percent,

$$\boldsymbol{\mu}\left(\boldsymbol{q}\right)\approx\mu\,\boldsymbol{I}=\frac{1}{6\pi\eta R_H}\boldsymbol{I},$$

where the effective hydrodynamic radius for an infinite system is numerically extrapolated to be $R_H\approx0.91\Delta x$ for a uniform grid spacing Δx and the three-point kernel [43]. This is consistent with previous results for a cell-centered discretization of the Navier-Stokes equations [77] and also with the results obtained using a Lattice-Boltzmann fluid solver in [3]. By using the Peskin four-point kernel [32] instead of the three-point discrete kernel function the translational invariance of the spatial discretization can be improved, however, at a potentially significant increase in computational cost. A more systematic investigation of different choices for the discrete kernel functions has been performed by Mori [78], however, these types of investigations have yet to be carried out within

the "blob" particle approach. We have found the inexpensive three-point function to perform quite well in our tests and use it exclusively in this work.

For non-periodic systems, one must generalize the definition of $\boldsymbol{J}$ and $\boldsymbol{S}$ in the case when the particle overlaps a physical boundary [79]. Even if a particle does not overlap a boundary, however, it will feel the boundary hydrodynamically and therefore both $\boldsymbol{J}\boldsymbol{\mathcal{P}}\boldsymbol{S}$ and $\boldsymbol{J}\boldsymbol{\mathcal{L}}^{-1}\boldsymbol{S}$ will depend on the proximity of the particle to physical boundaries. Implementing our algorithm in such cases may require first pre-tabulating the values of these $d \times d$ matrices for different positions of the particle relative to the boundaries.

B. Temporal Discretization

In this section, we describe how to integrate the spatially-discretized equations in time and update the fluid and particle velocities and particle position from time $n\Delta t$ to time $(n+1)\,\Delta t$, where Δt is the time step size, which can in principle be adjusted dynamically but we will assume it is kept fixed. We will use a superscript to denote the time level at which a given quantity is evaluated, for example, $\boldsymbol{q}^{n+\frac{1}{2}}$ will denote a mid-point estimate for the position of the particle at time $\left(n+\frac{1}{2}\right)\Delta t$. Similarly, $\boldsymbol{F}^{n+\frac{1}{2}} = \boldsymbol{F}\left(\boldsymbol{q}^{n+\frac{1}{2}}\right)$ will denote the force on the particle (due to external sources or other particles) evaluated at the position $\boldsymbol{q}^{n+\frac{1}{2}}$ and, implicitly, at time $\left(n+\frac{1}{2}\right)\Delta t$ in the case of a time dependent force.

For a neutrally-buoyant particle, $m_e = 0$, the fluid momentum equation is coupled to the particle position only through the forcing term $\boldsymbol{SF}$. In the deterministic setting, Griffith and Luo [**?**] have developed a second-order splitting scheme for integrating the spatially-discretized equations in time. The fluid solver in this scheme is very similar to the predictor-corrector scheme employed in the stochastic setting in Ref. [44]. The temporal discretization that we present next is based on replacing the fluid solver in Ref. [**?**] with that in Ref. [44], at least for the simpler case $m_e = 0$. The main difficulty is including the additional inertia from the particles in the fluid momentum update in a computationally-efficient manner.

In Ref. [43] a first-order splitting algorithm was developed for the case of a compressible fluid. This type of algorithm is similar to the original projection algorithm of Chorin [80] for incompressible flow and can be summarized as follows. Update the fluid first without accounting for the force $\boldsymbol{\lambda}$ exerted by the particle. Then, solve for the value of $\boldsymbol{\lambda}$ that, when applied as a correction to the fluid update, exactly imposes the no-slip condition. Extending this type of approach to be higher than first order accurate is known to be difficult from the literature on incompressible flow [48], due to the fact that the splitting introduces a commutator error. Here we follow a different, though related approach, which allows to construct a more accurate algorithm for viscous-dominated flows.

Our temporal scheme will be based on the following approach, which can be shown to be second-order deterministically by a Taylor series expansion of the temporal local truncation-error:

1. Estimate the position of the particle at the midpoint to leading order,
$$\boldsymbol{q}^{n+\frac{1}{2}} = \boldsymbol{q}^{n} + \frac{\Delta t}{2}\boldsymbol{J}^{n}\boldsymbol{v}^{n}. \tag{29}$$
2. Update the fluid velocity based on (13) using a second-order algorithm, while keeping the particle positions fixed at the midpoint estimates,
$$\begin{aligned}\left(\rho\boldsymbol{I} + m_e\boldsymbol{S}^{n+\frac{1}{2}}\boldsymbol{J}^{n+\frac{1}{2}}\right)\frac{\boldsymbol{v}^{n+1}-\boldsymbol{v}^{n}}{\Delta t} + \boldsymbol{\nabla}\pi^{n+\frac{1}{2}} \;&=\; -\boldsymbol{\nabla}\cdot\left(\rho\boldsymbol{v}\boldsymbol{v}^{T} - \boldsymbol{\sigma}\right)^{n+\frac{1}{2}} + \boldsymbol{S}^{n+\frac{1}{2}}\boldsymbol{F}^{n+\frac{1}{2}},\\ &-\; \left[m_e\boldsymbol{S}\boldsymbol{J}\left(\boldsymbol{v}\cdot\frac{\partial}{\partial\boldsymbol{q}}\boldsymbol{J}\right)\boldsymbol{v}\right]^{n+\frac{1}{2}}\end{aligned} \tag{30}$$
subject to $\boldsymbol{\nabla}\cdot\boldsymbol{v}^{n+1} = 0$. Here a second order Runge-Kutta [52] or Adams-Bashforth [**?**] scheme can be used to evaluate the fluid momentum fluxes to at least second-order accuracy, denoted generically here by superscript $n+\frac{1}{2}$.
3. Update the particle position using a second-order midpoint estimate of the velocity,
$$\boldsymbol{q}^{n+1} = \boldsymbol{q}^{n} + \frac{\Delta t}{2}\boldsymbol{J}^{n+\frac{1}{2}}\left(\boldsymbol{v}^{n+1}+\boldsymbol{v}^{n}\right). \tag{31}$$

Observe that the above scheme never actually uses the particle velocity $\boldsymbol{u}$, although one can and should keep track of the particle excess momentum $m_e\boldsymbol{u}$ and update it whenever the fluid momentum is updated, to ensure strict conservation of momentum. Also observe that during the fluid update we fix the particle at its midpoint position $\boldsymbol{q}^{n+\frac{1}{2}}$.

1. Velocity Update

The most difficult step in the time stepping algorithm summarized above is the momentum (velocity) update, step 2. In order to update the velocity of the fluid we need to calculate the fluid momentum change due to viscosity and thermal fluctuations and also the momentum exchange with the particle, all to second order in time. Our scheme is based on solving for the values of the Lagrange multipliers $\pi^{n+\frac{1}{2}}$ and $\boldsymbol{\lambda}^{n+\frac{1}{2}}$ such that at the end of the time step both the incompressibility and the no-slip constraints are satisfied,

$$\begin{aligned}
\rho\frac{\boldsymbol{v}^{n+1}-\boldsymbol{v}^n}{\Delta t}+\boldsymbol{\nabla}\pi^{n+\frac{1}{2}} &= -\boldsymbol{\nabla}\cdot\left(\rho\boldsymbol{v}\boldsymbol{v}^T-\boldsymbol{\sigma}\right)^{n+\frac{1}{2}}-\boldsymbol{S}^{n+\frac{1}{2}}\boldsymbol{\lambda}^{n+\frac{1}{2}}\\
m_e\boldsymbol{u}^{n+1} &= m_e\boldsymbol{u}^n+\Delta t\,\boldsymbol{F}^{n+\frac{1}{2}}+\Delta t\,\boldsymbol{\lambda}^{n+\frac{1}{2}}\\
\boldsymbol{\nabla}\cdot\boldsymbol{v}^{n+1} &= 0\\
\boldsymbol{u}^{n+1} &= \boldsymbol{J}^{n+\frac{1}{2}}\boldsymbol{v}^{n+1}+\Delta\boldsymbol{u}^{n+\frac{1}{2}}.
\end{aligned}\tag{32}$$

A correction $\Delta\boldsymbol{u}^{n+\frac{1}{2}}$ is included to account for the fact that the no-slip condition is not correctly centered since $\boldsymbol{J}$ and $\boldsymbol{v}$ are evaluated at different points in time, as we explain shortly.

If we multiply the particle velocity update by $\boldsymbol{S}^{n+\frac{1}{2}}$ and add it to the fluid equation, and use the no-slip constraint including $\boldsymbol{u}^n=\boldsymbol{J}^{n-\frac{1}{2}}\boldsymbol{v}^n+\Delta\boldsymbol{u}^{n-\frac{1}{2}}$, we get (30) with the kinetic term in the last line approximated as

$$\left[m_e\boldsymbol{S}\left(\boldsymbol{u}\cdot\frac{\partial}{\partial\boldsymbol{q}}\boldsymbol{J}\right)\boldsymbol{v}\right]^{n+\frac{1}{2}}=m_e\boldsymbol{S}^{n+\frac{1}{2}}\left[\left(\frac{\boldsymbol{J}^{n+\frac{1}{2}}-\boldsymbol{J}^{n-\frac{1}{2}}}{\Delta t}\right)\boldsymbol{v}^n+\frac{\Delta\boldsymbol{u}^{n+\frac{1}{2}}-\Delta\boldsymbol{u}^{n-\frac{1}{2}}}{\Delta t}\right].\tag{33}$$

The penultimate term in the above equation can be seen as a discretization of the kinetic term

$$\left(\frac{\boldsymbol{J}^{n+\frac{1}{2}}-\boldsymbol{J}^{n-\frac{1}{2}}}{\Delta t}\right)\boldsymbol{v}^n\approx\left[\left(\boldsymbol{u}\cdot\frac{\partial}{\partial\boldsymbol{q}}\boldsymbol{J}\right)\boldsymbol{v}\right]^n.$$

This is consistent with the continuum equations but it yields only first-order and *not* second-order accuracy in Δt because it is not centered at time level $n+\frac{1}{2}$. This reduction of the accuracy comes because the no-slip constraint in (32) uses the midpoint instead of the endpoint position of the particle.

The above discussion shows that setting $\Delta\boldsymbol{u}^{n+\frac{1}{2}}=\boldsymbol{0}$ results in a first-order scheme if $m_e\neq 0$. To get second-order accuracy, we need to apply a nonzero correction to the no-slip constraint. Imposing the no-slip constraint at the end of the time step, $\boldsymbol{u}^{n+1}=\boldsymbol{J}^{n+1}\boldsymbol{v}^{n+1}$, leads to a formulation that is implicit in both $\boldsymbol{q}^{n+1}$ and $\boldsymbol{v}^{n+1}$, which is difficult to implement in practice. Instead, we can center the no-slip constraint as

$$\frac{1}{2}\boldsymbol{J}^{n+\frac{1}{2}}\left(\boldsymbol{v}^{n+1}+\boldsymbol{v}^n\right)=\frac{1}{2}\left(\boldsymbol{u}^{n+1}+\boldsymbol{J}^n\boldsymbol{v}^n\right)=\frac{1}{2}\left(\boldsymbol{J}^{n+\frac{1}{2}}\boldsymbol{v}^{n+1}+\Delta\boldsymbol{u}^{n+\frac{1}{2}}+\boldsymbol{J}^n\boldsymbol{v}^n\right),$$

which gives the no-slip centering correction

$$\Delta\boldsymbol{u}^{n+\frac{1}{2}}=\left(\boldsymbol{J}^{n+\frac{1}{2}}-\boldsymbol{J}^n\right)\boldsymbol{v}^n.\tag{34}$$

This correction for the no-slip constraint is simple to implement with only one additional local averaging operation to evaluate $\boldsymbol{J}^n\boldsymbol{v}^n$. Note that we purposely used $\boldsymbol{J}^n\boldsymbol{v}^n$ instead of $\boldsymbol{u}^n$ here since in our formulation, and also in our algorithm, $\boldsymbol{u}^{n+1}$ is only used as an intermediate variable. A Taylor series analysis shows that using (34) makes (33) a centered second-order approximation of the kinetic term

$$\left(\frac{\boldsymbol{J}^{n+\frac{1}{2}}-\boldsymbol{J}^{n-\frac{1}{2}}}{\Delta t}\right)\boldsymbol{v}^n+\frac{\Delta\boldsymbol{u}^{n+\frac{1}{2}}-\Delta\boldsymbol{u}^{n-\frac{1}{2}}}{\Delta t}=\left[\left(\boldsymbol{u}\cdot\frac{\partial}{\partial\boldsymbol{q}}\boldsymbol{J}\right)\boldsymbol{v}\right]^{n+\frac{1}{2}}+O\left(\Delta t^3\right).$$

A Taylor series analysis confirms that using the no-slip correction (34) leads to a second-order algorithm for updating the position of the particle and the velocity of the fluid.

In order to avoid one more additional local averaging operation (which requires an irregular memory access pattern and is thus costly, especially in a GPU-based implementation) we can set $\Delta\boldsymbol{u}^{n+\frac{1}{2}}=\boldsymbol{0}$. We are primarily concerned with viscous-dominated (low Reynolds number) flows, for which the kinetic term $(\boldsymbol{u}\cdot\partial\boldsymbol{J}/\partial\boldsymbol{q})\,\boldsymbol{v}$ is small (quadratic in $\boldsymbol{v}$, just like the advective term $\boldsymbol{v}\cdot\boldsymbol{\nabla}\boldsymbol{v}$) and can be approximated to first order without a significant reduction in the overall accuracy of the method. As we explain in Appendix C, in our scheme $\Delta\boldsymbol{u}^{n+\frac{1}{2}}$ contains an additional higher-order correction ($O(\Delta t^3)$ for smooth flows) that arises solely due to the implicit handling of viscosity.

2. Semi-Implicit Discretization of Viscous Terms

During the fluid update the particle position remains fixed at $\boldsymbol{q}^{n+\frac{1}{2}}$. For notational simplicity, in the remainder of this paper we will sometimes drop the time step index from $\boldsymbol{J}$ and $\boldsymbol{S}$; unless otherwise indicated, they are always evaluated at $\boldsymbol{q}^{n+\frac{1}{2}}$.

Following Ref. [44], our second-order implementation of the velocity update (30) treats the viscous term semi-implicitly and the remaining terms explicitly,

$$-\boldsymbol{\nabla}\cdot\left(\rho\boldsymbol{v}\boldsymbol{v}^T-\boldsymbol{\sigma}\right)^{n+\frac{1}{2}}=\frac{\eta}{2}\boldsymbol{\nabla}^2\left(\boldsymbol{v}^{n+1}+\boldsymbol{v}^n\right)+\boldsymbol{\nabla}\cdot\boldsymbol{\Sigma}^n-\boldsymbol{\nabla}\cdot\left(\rho\boldsymbol{v}\boldsymbol{v}^T\right)^{n+\frac{1}{2}}.$$

The spatial discretization of the stochastic flux is [44]

$$\boldsymbol{\Sigma}^n=\left(\frac{k_BT\eta}{\Delta V_f\,\Delta t}\right)^{\frac{1}{2}}\left[\boldsymbol{W}^n+(\boldsymbol{W}^n)^T\right],$$

where $\Delta V_f=\Delta x\Delta y\Delta z$ is the volume of the hydrodynamic cells, and $\boldsymbol{W}^n$ is a collection of i.i.d. unit normal variates, generated independently at each time step on the faces of the staggered momentum grid. To approximate the advective fluxes to second order in Δt, one can use either the predictor-corrector method described in Ref. [44] or, more efficiently, one can use the classical (time lagged) Adams-Bashforth method [**?**]

$$\boldsymbol{\nabla}\cdot\left(\rho\boldsymbol{v}\boldsymbol{v}^T\right)^{n+\frac{1}{2}}=\frac{3}{2}\boldsymbol{\nabla}\cdot\left(\rho\boldsymbol{v}\boldsymbol{v}^T\right)^{n}-\frac{1}{2}\boldsymbol{\nabla}\cdot\left(\rho\boldsymbol{v}\boldsymbol{v}^T\right)^{n-1}. \tag{35}$$

For viscous-dominated (small Reynolds number) flows, one can also approximate the advective terms to first order only without a significant reduction of the overall accuracy for reasonably large time steps.

Referring back to Eq. (30), we see that updating the fluid momentum semi-implicitly requires solving the linear system

$$\begin{aligned}\left[\left(\frac{\rho\boldsymbol{I}+m_e\boldsymbol{S}^{n+\frac{1}{2}}\boldsymbol{J}^{n+\frac{1}{2}}}{\Delta t}\right)-\frac{\eta}{2}\boldsymbol{\nabla}^2\right]\boldsymbol{v}^{n+1}+\boldsymbol{\nabla}\pi^{n+\frac{1}{2}} &= \left[\left(\frac{\rho\boldsymbol{I}+m_e\boldsymbol{S}^{n+\frac{1}{2}}\boldsymbol{J}^{n+\frac{1}{2}}}{\Delta t}\right)+\frac{\eta}{2}\boldsymbol{\nabla}^2\right]\boldsymbol{v}^n\\ &\quad-\boldsymbol{\nabla}\cdot\left(\rho\boldsymbol{v}\boldsymbol{v}^T\right)^{n+\frac{1}{2}}+\boldsymbol{\nabla}\cdot\boldsymbol{\Sigma}^n+\boldsymbol{S}^{n+\frac{1}{2}}\boldsymbol{F}^{n+\frac{1}{2}}-\left[m_e\boldsymbol{S}\boldsymbol{J}\left(\boldsymbol{v}\cdot\frac{\partial}{\partial\boldsymbol{q}}\boldsymbol{J}\right)\boldsymbol{v}\right]^{n+\frac{1}{2}}.\end{aligned} \tag{36}$$

If $m_e=0$, we can solve the linear system (36) for the unknowns $\boldsymbol{v}^{n+1}$ and $\pi^{n+\frac{1}{2}}$ using a preconditioned iterative solver [75], as explained in more detail in Ref. [44]. For periodic systems the system (36) can be solved easily by using a projection method together with FFT-based velocity and pressure linear solvers. For non-neutrally-buoyant particles, however, solving (36) requires developing a specialized preconditioned Krylov method. Here we develop an approximate solver for (32) that only requires a few FFTs, and will be shown in Section III D to give nearly second-order accuracy for a wide range of relevant time step sizes.

Our approach consists of splitting the velocity solver into two steps. In the first step, we ignore the inertia of the particle, i.e., delete the $m_e\boldsymbol{S}\boldsymbol{J}$ term in (36), and solve for a provisional velocity $\tilde{\boldsymbol{v}}^{n+1}$and pressure $\tilde{\pi}^{n+\frac{1}{2}}$,

$$\left(\frac{\rho}{\Delta t}\boldsymbol{I}-\frac{\eta}{2}\boldsymbol{\nabla}^2\right)\tilde{\boldsymbol{v}}^{n+1}+\boldsymbol{\nabla}\tilde{\pi}^{n+\frac{1}{2}}=\left(\frac{\rho}{\Delta t}\boldsymbol{I}+\frac{\eta}{2}\boldsymbol{\nabla}^2\right)\boldsymbol{v}^n+\boldsymbol{\nabla}\cdot\boldsymbol{\Sigma}^n+\boldsymbol{S}\boldsymbol{F}^{n+\frac{1}{2}}-\boldsymbol{\nabla}\cdot\left(\rho\boldsymbol{v}\boldsymbol{v}^T\right)^{n+\frac{1}{2}}, \tag{37}$$

subject to $\boldsymbol{\nabla}\cdot\tilde{\boldsymbol{v}}^{n+1}=0$. If $m_e=0$, this completes the fluid solve and setting $\boldsymbol{v}^{n+1}=\tilde{\boldsymbol{v}}^{n+1}$ gives us second-order accuracy for the viscous and stochastic terms [52]. If $m_e\neq0$, we need to find a velocity correction $\Delta\boldsymbol{v}^{n+\frac{1}{2}}=\boldsymbol{v}^{n+1}-\tilde{\boldsymbol{v}}^{n+1}$ and pressure correction $\Delta\pi^{n+\frac{1}{2}}=\pi^{n+\frac{1}{2}}-\tilde{\pi}^{n+\frac{1}{2}}$ that takes into account the inertia of the particle. We do this by splitting the linear system (32) into two equations, (37) for the unperturbed velocity field, and

$$\left(\frac{\rho}{\Delta t}\boldsymbol{I}-\frac{\eta}{2}\boldsymbol{\nabla}^2\right)\Delta\boldsymbol{v}^{n+\frac{1}{2}}+\boldsymbol{\nabla}\left(\Delta\pi^{n+\frac{1}{2}}\right) = -\boldsymbol{S}\left(\boldsymbol{\lambda}^{n+\frac{1}{2}}+\boldsymbol{F}^{n+\frac{1}{2}}\right) \tag{38}$$

$$\frac{m_e}{\Delta t}\boldsymbol{u}^{n+1} = \frac{m_e}{\Delta t}\boldsymbol{u}^n+\left(\boldsymbol{\lambda}^{n+\frac{1}{2}}+\boldsymbol{F}^{n+\frac{1}{2}}\right) \tag{39}$$

$$\boldsymbol{u}^{n+1} = \boldsymbol{J}\left(\tilde{\boldsymbol{v}}^{n+1}+\Delta\boldsymbol{v}^{n+\frac{1}{2}}\right)+\Delta\boldsymbol{u}^{n+\frac{1}{2}} \tag{40}$$

$$\boldsymbol{\nabla}\cdot\left(\Delta\boldsymbol{v}^{n+\frac{1}{2}}\right) = 0,$$

for the perturbed field. This gives a linear system of equations for the unknowns $\tilde{\boldsymbol{v}}^{n+1}$, $\boldsymbol{u}^{n+1}$, $\Delta\boldsymbol{v}^{n+\frac{1}{2}}$, $\tilde{\pi}^{n+\frac{1}{2}}$, $\Delta\pi^{n+\frac{1}{2}}$, and $\boldsymbol{\lambda}^{n+\frac{1}{2}}$. We explain how we solve this linear system of equations in Appendix C

for periodic boundaries using Fourier Transform techniques. Here we simply summarize the resulting algorithm, as implemented in our code. In Appendix D we give a summary of a similar algorithm for compressible flow, which our code also implements.

C. Summary of Algorithm

1. Estimate the position of the particle at the midpoint,
$$\boldsymbol{q}^{n+\frac{1}{2}} = \boldsymbol{q}^n + \frac{\Delta t}{2}\boldsymbol{J}^n\boldsymbol{v}^n, \tag{41}$$
and evaluate the external or interparticle forces $\boldsymbol{F}^{n+\frac{1}{2}}\left(\boldsymbol{q}^{n+\frac{1}{2}}\right)$.
2. Solve the unperturbed fluid equation
$$\begin{aligned}\rho\frac{\tilde{\boldsymbol{v}}^{n+1}-\boldsymbol{v}^n}{\Delta t} + \boldsymbol{\nabla}\tilde{\pi}^{n+\frac{1}{2}} &= \frac{\eta}{2}\boldsymbol{\nabla}^2\left(\tilde{\boldsymbol{v}}^{n+1}+\boldsymbol{v}^n\right) + \boldsymbol{\nabla}\cdot\left[\left(\frac{k_BT\eta}{\Delta V_f\,\Delta t}\right)^{1/2}\left(\boldsymbol{W}^n+(\boldsymbol{W}^n)^T\right)\right] \\ &- \left[\frac{3}{2}\boldsymbol{\nabla}\cdot\left(\rho\boldsymbol{v}\boldsymbol{v}^T\right)^n - \frac{1}{2}\boldsymbol{\nabla}\cdot\left(\rho\boldsymbol{v}\boldsymbol{v}^T\right)^{n-1}\right] + \boldsymbol{S}^{n+\frac{1}{2}}\boldsymbol{F}^{n+\frac{1}{2}},\end{aligned} \tag{42}$$
$$\boldsymbol{\nabla}\cdot\tilde{\boldsymbol{v}}^{n+1} = 0, \tag{43}$$
using a projection algorithm and FFTs to diagonalize the Laplacian operator.
3. If $m_e = 0$, set $\boldsymbol{v}^{n+1} = \tilde{\boldsymbol{v}}^{n+1}$ and skip to step 9.
4. Evaluate the slip correction
$$\delta\boldsymbol{u}^{n+\frac{1}{2}} = \left(\boldsymbol{J}^{n+\frac{1}{2}}-\boldsymbol{J}^n\right)\boldsymbol{v}^n + \frac{\nu\Delta t}{2}\boldsymbol{J}^{n-\frac{1}{2}}\boldsymbol{\nabla}^2\left(\Delta\boldsymbol{v}^{n-\frac{1}{2}}\right) \tag{44}$$
and the change of the particle excess momentum
$$\Delta\boldsymbol{p} = m_e\left(\boldsymbol{u}^n - \boldsymbol{J}^{n+\frac{1}{2}}\tilde{\boldsymbol{v}}^{n+1} - \delta\boldsymbol{u}^{n+\frac{1}{2}}\right).$$
5. Calculate the fluid velocity perturbation due to the excess inertia of the particle
$$\Delta\tilde{\boldsymbol{v}} = \frac{\tilde{m}_f}{\rho\left(\tilde{m}_f+m_e\right)}\boldsymbol{\mathcal{P}}\boldsymbol{S}\Delta\boldsymbol{p},$$
using FFTs to implement the discrete projection $\boldsymbol{\mathcal{P}}$, where $\tilde{m}_f = d\,\rho\Delta V/(d-1)$.
6. Account for the viscous contribution to the velocity perturbation by solving the system
$$\left(\rho\boldsymbol{I} - \frac{\Delta t}{2}\eta\boldsymbol{\nabla}^2\right)\Delta\boldsymbol{v}^{n+\frac{1}{2}} + \Delta t\,\boldsymbol{\nabla}\left(\Delta\pi^{n+\frac{1}{2}}\right) = \boldsymbol{S}^{n+\frac{1}{2}}\left(\Delta\boldsymbol{p} - m_e\boldsymbol{J}^{n+\frac{1}{2}}\Delta\tilde{\boldsymbol{v}}\right), \tag{45}$$
$$\boldsymbol{\nabla}\cdot\left(\Delta\boldsymbol{v}^{n+\frac{1}{2}}\right) = 0 \tag{46}$$
using a projection algorithm and FFTs to diagonalize the Laplacian operator.
7. Update the fluid velocity
$$\boldsymbol{v}^{n+1} = \tilde{\boldsymbol{v}}^{n+1} + \Delta\boldsymbol{v}^{n+\frac{1}{2}}. \tag{47}$$
8. Update the particle velocity in a momentum-conserving manner,
$$\boldsymbol{u}^{n+1} = \boldsymbol{J}^{n+\frac{1}{2}}\left(\tilde{\boldsymbol{v}}^{n+1}+\Delta\tilde{\boldsymbol{v}}\right) + \delta\boldsymbol{u}^{n+\frac{1}{2}}. \tag{48}$$
9. Update the particle position,
$$\boldsymbol{q}^{n+1} = \boldsymbol{q}^n + \frac{\Delta t}{2}\boldsymbol{J}^{n+\frac{1}{2}}\left(\boldsymbol{v}^{n+1}+\boldsymbol{v}^n\right). \tag{49}$$

We note that the full slip correction (44) is only required if m_e/m_f is large and the Reynolds number is large. For sufficiently small Reynolds numbers (viscous-dominated flows) we can neglect the quadratic advective term and only keep the linear term, and set
$$\delta\boldsymbol{u}^{n+\frac{1}{2}} = \frac{\nu\Delta t}{2}\boldsymbol{J}^{n-\frac{1}{2}}\boldsymbol{\nabla}^2\left(\Delta\boldsymbol{v}^{n-\frac{1}{2}}\right). \tag{50}$$
We can also set $\delta\boldsymbol{u}^{n+\frac{1}{2}} = \boldsymbol{0}$, and obtain a first-order algorithm that does not require any time lagging and has improved stability for very large time step sizes. We compare the three options (44), (50), and $\delta\boldsymbol{u}^{n+\frac{1}{2}} = \boldsymbol{0}$ numerically in Section III D. The remainder of the algorithm is not affected by the choice of the slip correction $\delta\boldsymbol{u}^{n+\frac{1}{2}}$.

D. Efficiency, Stability and Accuracy

With periodic boundary conditions the velocity and the pressure linear systems in the incompressible formulation decouple and Fast Fourier Transforms can be used to solve the system (42) efficiently, see Ref. [81] for additional details. We first solve the velocity equation (42) without the gradient of pressure term (this is a Helmholtz equation) using a Fourier transform to diagonalize the discrete Laplacian. Then, we project the solution onto the space of divergence free vector fields by subtracting a pressure gradient term. The pressure is a solution of a discrete Poisson equation, which can also efficiently be computed using Fourier transforms. Note that it is possible to generalize our algorithm to non-periodic systems by using the fluid solver developed by one of us [75] and employed in Ref. [44], at least for the case of neutrally buoyant particles, $m_e = 0$. For $m_e \neq 0$ new iterative solvers for the Stokes subproblem need to be developed.

We have parallelized the algorithm to run efficiently on Graphics Processing Units (GPUs), as explained in more detail in Ref. [43]. Our public domain implementation [43] is written in the CUDA programming environment, and is three-dimensional with the special case of $N_z = 1$ cell along the z axes corresponding to a quasi two-dimensional system. In our implementation we create one thread per cell, and each thread only writes to the memory address associated with its cell and only accesses the memory associated with its own and neighboring cells. This avoids concurrent writes and costly synchronizations between threads, facilitating efficient execution on the GPU. For incompressible flow, our present GPU implementation is specific to periodic systems, and uses the NVIDIA FFT library as a Poisson/Helmholtz solver.

The stability and accuracy of our spatio-temporal discretization is controlled by the dimensionless advective and viscous CFL numbers

$$\alpha = \frac{V\Delta t}{\Delta x}, \quad \beta = \frac{\nu\Delta t}{\Delta x^2}, \tag{51}$$

where V is a typical advection speed, which may be dominated by the thermal velocity fluctuations or by a deterministic background flow. Here we always use the same grid spacing along all dimensions, $\Delta x = \Delta y = \Delta z$. The strength of advection relative to dissipation is measured by the cell Reynolds number $r = \alpha/\beta = V\Delta x/\nu$. Note that for compressible flow (see Ref. [44] and Appendix D) there is a sonic CFL number $\alpha_s = c\Delta t/\Delta x$, where c is the speed of sound.

The explicit handling of the advective terms places a stability condition $\alpha \lesssim 1$, in fact, for $\alpha > 1$ a particle can move more than a hydrodynamic cell during a single time step and this causes not only stability but also implementation difficulties. It is not hard to see that in the absence of advection our semi-implicit discretization of viscosity is stable for any value of β, however, it is only by keeping $\beta \lesssim 1$ that we can ensure the dynamics of all or at least most fluid modes is resolved [52]. We consider a temporal integrator to be "good" if it produces reasonably-accurate results with a time step for which at least one of α or β is close to 1/2. Typically, flows at small scales are viscous dominated ($r \ll 1$) so that the time step is primarily limited by β and not by α.

Next we numerically check the deterministic order of accuracy of the temporal integrator. Based on local truncation error analysis we expect that the temporal integrator summarized in Section III C is formally second-order accurate. For small Reynolds numbers, we expect to see nearly second-order accuracy in practice even if we use the slip correction (50) instead of (44). We also recall that for $m_e \neq 0$ we made an uncontrolled approximation in assuming that $\boldsymbol{\mathcal{JPS}}$ is translationally-invariant, which is only accurate to about a percent for the three-point Peskin local averaging and spreading operators. This approximation leads to another error in imposing the no-slip condition, which we expect to lead to first-order accuracy for very small time step sizes.

As a test of the temporal accuracy, we study the deterministic motion of a particle in an centrally-symmetric harmonic potential $V(r) = kr^2/2$, where r is the distance from the origin and k is a spring constant. In these tests we keep the spatial discretization (and thus the blob particle shape) fixed and only change the time step size Δt. We start the particle from rest at a certain distance r_0 from the origin and then release it. The particle will perform damped oscillations under the influence of the spring and viscous friction. We look at the error in the position of the particle $\boldsymbol{q}(t)$ defined as the average of the difference between the position for time steps Δt and $\Delta t/2$ over

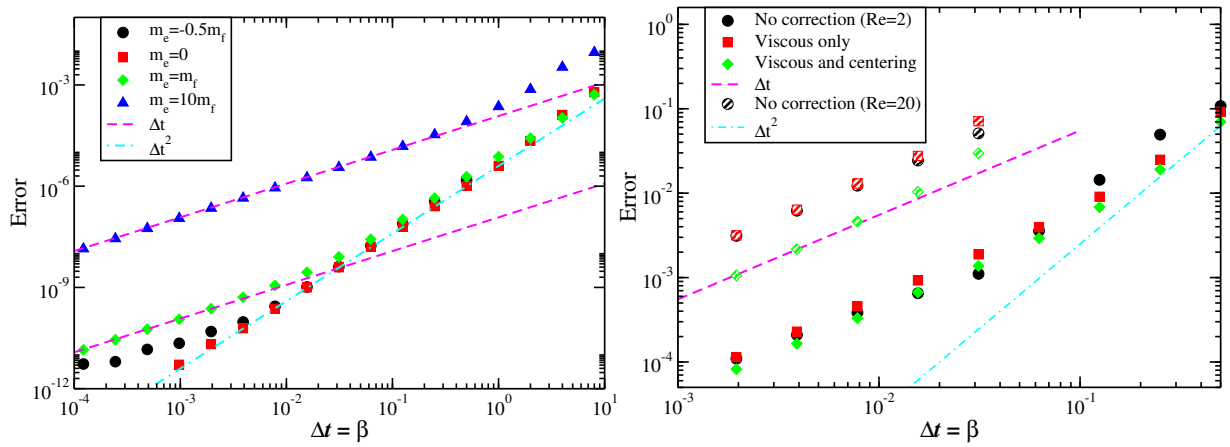

Figure 2: Average error (52) over a short deterministic trajectory of a particle initially at rest and subsequently moving under the action of a harmonic potential. (*Left panel*) Small Reynolds number, Re ≈ 0.02, with no-slip correction (50). Several values of the excess particle mass m_e relative to the mass of the dragged fluid m_f (symbols, see legend) are shown, including neutrally-buoyant particles ($m_e = 0$). Expected error decay rates for a first and second-order scheme are shown with lines. (*Right panel*) Comparison of the three choices for the slip correction, (44) (viscous and centering corrections), (50) (only viscous correction), and (44) (no correction), for Reynolds numbers Re ≈ 2 (full symbols) and Re ≈ 20 (shaded symbols). The excess mass is $m_e = 10m_f$.

a certain number of time steps N_s, from the initial time to a time $T = N_S \Delta t$,

$$E(\Delta t) = \frac{1}{N_s} \sum_{n=1}^{N_s} \left\| \mathbf{q}_{\Delta t}\left(n\Delta t\right) - \mathbf{q}_{\Delta t/2}\left(2n\frac{\Delta t}{2}\right) \right\|. \tag{52}$$

For a numerical scheme with order of accuracy p this error should behave as $E = O\left(\Delta t^p\right)$ for sufficiently small Δt.

We perform this test for several choices of the density of the particle relative to the fluid, as measured via the ratio m_e/m_f, including particles less dense than the surrounding fluid (negative excess mass m_e). The tests are performed for a periodic system of 16^3 hydrodynamic cells of size $\Delta x = 1$, with fluid density $\rho = 1$ and shear viscosity $\eta = 1$ (in arbitrary units), with the particle started at position $x_0 = r_0 = 10$, and follow the motion of the particles to a time $T = 80$. In the left panel of Fig. 2 we show the error (52) for four different values of the excess mass for a spring constant $k = 0.01$, which implies a small Reynolds number $Re \approx r = u_{\max}\Delta x/\nu \approx 0.02$, where $u_{\max}$ is the maximal speed of the particle. As expected, the figure shows clear second-order convergence for neutrally-buoyant particles ($m_e = 0$), and a transition from essentially second-order (at larger Δt) to first-order (at smaller Δt) accuracy in the other cases. Notably, for the larger time steps, which are of more practical interest, we see second-order convergence. As expected, the transition from second to first order of accuracy occurs at a larger Δt for particles that are far from neutrally-buoyant, and for $m_e/m_f = 10$ we see first-order accuracy over a broader range of time step sizes.

In order to compare the three choices for the the slip correction, (44), (50) and $\delta \boldsymbol{u}^{n+\frac{1}{2}} = \mathbf{0}$, in the right panel of Fig. 2 we compare the error for the three choices for the case of a large particle excess mass $m_e = 10m_f$ and a larger Reynolds number, Re ≈ 2 for $k = 1$, and Re ≈ 20 for $k = 500$. We see qualitatively similar behavior as for the small Reynolds number case $k = 0.01$. Compared to the simple choice $\delta \boldsymbol{u}^{n+\frac{1}{2}} = \mathbf{0}$, we see a modest improvement in the error when we use (44) for the largest Reynolds number, and we see a small improvement when we use (50) for intermediate time step sizes. Note that for our choice of parameters the viscous CFL number is $\beta = \Delta t$ and the advective CFL is $\alpha \approx \beta\,\mathrm{Re}$. For large Reynolds numbers the time step is limited by the requirement $\alpha \lesssim 1$.

Since the majority of the tests presented here are at small Reynolds numbers, we use (50) instead of the more expensive (44). For several of the tests we have also tried $\delta \boldsymbol{u}^{n+\frac{1}{2}} = \mathbf{0}$ and observed similar results.

grid spacing Δx	1
grid size	32^3
fluid density ρ	1
shear viscosity η	1
time step size Δt	1
temperature k_BT	0.001
LJ strength ϵ	0.001
LJ diameter σ	2
number of particles N	1000

grid spacing Δx	1
grid size	128^3
fluid density	1
viscosity	1
advective CFL α	0.01, 0.1, or 0.25
viscous CFL β	9.2, 92, or 230
excess mass m_e	m_f

Table I: (*Left*) Parameters used in the RDF simulations shown in the left panel of Fig. 3. (*Right*) Simulation parameters for the hydrodynamic interaction simulations presented in the right panel of Fig. 3.

IV. Results

In this section we validate and test the performance of the algorithm summarized in Section III C on a variety of standard problems from soft-condensed matter applications. We also examine the behavior of the minimally-resolved blob in a large Reynolds number flow.

In the first few tests we examine the performance of the algorithm for thermal systems. For a system at thermodynamic equilibrium at rest the typical value of the advection velocity to be used in the definition of the advective CFL number α (51) is the equilibrium magnitude of the thermal velocity fluctuations,

$$V \approx \sqrt{\frac{k_BT}{\rho\Delta x^3}}.$$

The temporal integrator we employ here is designed to accurately resolve the short time dynamics of the blobs when the time step size Δt is reasonably small. We consider a time step size Δt large if at least one of the advective (α) or the viscous (β) CFL numbers defined in (51) becomes comparable to unity. It is important to emphasize that because of the semi-implicit second-order nature of the temporal integrator, the algorithm is robust over a broad range of time step sizes, which would be well beyond the stability limit of explicit integrators for compressible flow.

A. Equilibrium Properties

One of the most important requirements on any scheme that couples fluctuating hydrodynamics to immersed particles is to reproduce the Gibbs-Boltzmann distribution at thermodynamic equilibrium. In particular, the probability distribution of the positions $\boldsymbol{Q} = \{\boldsymbol{q}_1, \boldsymbol{q}_2, \ldots, \boldsymbol{q}_N\}$ of a collection of N particles interacting with a conservative potential $U(\boldsymbol{Q})$ should be

$$P(\boldsymbol{Q}) \sim \exp\left[-\frac{U(\boldsymbol{Q})}{k_BT}\right], \tag{53}$$

independent of any dynamical parameters such as viscosity or particle inertia. This follows from the balance between the dissipative and stochastic forcing terms and requires consistently including thermal fluctuations in the momentum equation.

We verify that our incompressible inertial coupling algorithm gives the correct equilibrium distribution $P(\boldsymbol{Q})$ by computing the radial (pair) distribution function (RDF) $g(r)$ for a collection of colloidal particles interacting with a pairwise potential $V(r)$, $U(\boldsymbol{Q}) = \sum_{i,j=1}^{N} V\left(\left\|\boldsymbol{q}_i - \boldsymbol{q}_j\right\|\right)$. We use the purely repulsive truncated Lennard-Jones (WCA) potential

$$V(r) \;=\; \begin{cases} 4\epsilon\left(\left(\frac{\sigma}{r}\right)^{12} - \left(\frac{\sigma}{r}\right)^{6}\right) + \epsilon, & r < 2^{1/6}\sigma \\ 0, & r > 2^{1/6}\sigma \end{cases} \tag{54}$$

In Fig. 3 we compare $g(r)$ between a simulation where the particles are immersed in an incompressible viscous solvent, and a standard computation of the equilibrium RDF using a Monte Carlo algorithm to sample the equilibrium distribution (53). The parameters for these simulations are given in Table I. For both $m_e = 0$ and $m_e = m_f$ we obtain excellent agreement with the Monte Carlo calculations, even for the rather large time step size $\beta = 1$.

It is important to observe that the correct equilibrium structure for the positional degrees of

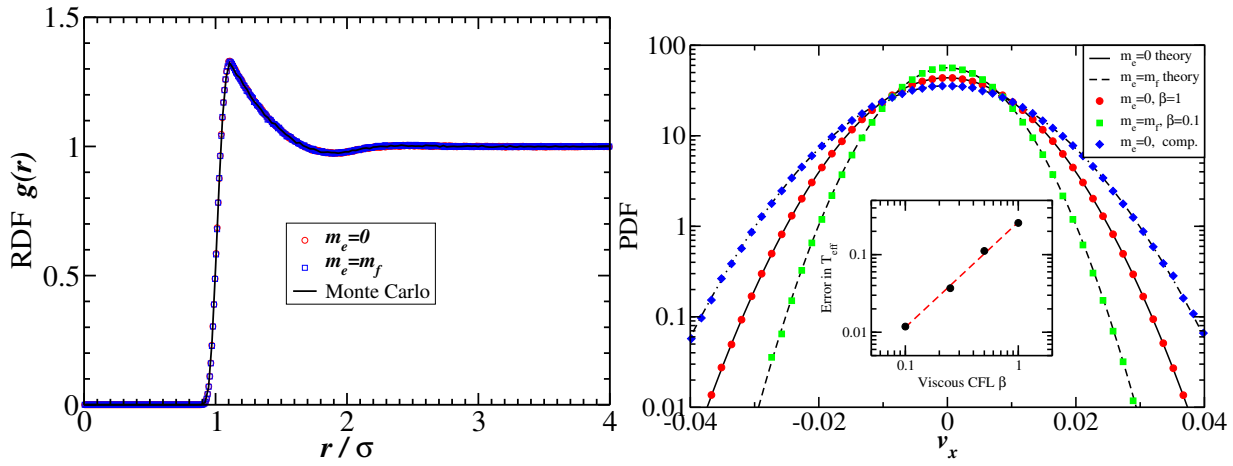

Figure 3: (*Left panel*) The equilibrium radial distribution function $g(r)$ for a suspension of particles interacting with WCA potential (54) with cutoff radius σ. The results for two different particle inertias are compared to a Monte Carlo sampling of the Gibbs-Boltzmann distribution. (*Right panel*) The equilibrium probability distribution for the x component of the velocity of the particles in the suspension for different excess masses, for both a compressible and an incompressible fluid (symbols). The Gaussian distribution dictated by the equipartition principle is shown for comparison (lines). The inset in the figure shows the error ϵ_T in the effective temperature of the particles as a function of the time step size, for $m_e = m_f$.

freedom is obtained for both the compressible and the incompressible formulations, even though the blob velocities have very different distributions. Specifically, as explained in Section II F, in the incompressible case the variance of a component of the velocity is not k_BT/m as for a compressible fluid, but rather, $k_BT/\tilde{m}$, where the effective blob mass $\tilde{m} > m$ includes an added mass due to the incompressible fluid dragged with the blob. In the right panel of Fig. 3 we show the probability distribution for the velocities of the blobs in the colloidal suspension, and compare them to the theoretical predictions. The numerical variance of the velocity is sensitive to the time step size for $m_e \neq 0$, and a small time step size ($\alpha, \beta < 0.25$) is required to obtain a reasonably accurate variance. This is shown in the inset of the right panel of Fig. 3, where the relative error in the effective "temperature" of the blob $\epsilon_T = (\langle u^2 \rangle - k_BT/\tilde{m}) / (k_BT/\tilde{m})$ is shown as a function of the viscous CFL number β.

B. Velocity Autocorrelation Function

In this section we apply our scheme to a standard test for the coupling of spherical particle of hydrodynamic radius R_H to a compressible [43, 60, 82–84] or incompressible [36, 38, 85, 86] fluid solver. The velocity autocorrelation function (VACF)

$$C(t) = \langle v_x(0) v_x(t) \rangle = \frac{1}{d} \langle \boldsymbol{v}(0) \cdot \boldsymbol{v}(t) \rangle , \tag{55}$$

of a single free Brownian particle diffusing through a periodic fluid is a non-trivial quantity that contains crucial information at both short and long times. The integral of the VACF determines the diffusion coefficient and gets contributions from three distinct stages. Firstly, at molecular times equipartition dictates that $C(0) = k_BT/m$, an important signature of fluctuation-dissipation balance that has proven challenging for several fluid-particle coupling methods [60, 82, 85, 86]. We recall that for our particle the effective particle mass $m = m_e + m_f$ includes the mass of the fluid dragged with the particle m_f, as well as the excess mass m_e. The compressible inertial coupling method is able to reproduce the intercept k_BT/m very accurately even for relatively large sound CFL numbers [43].

On the time scale of sound waves, $t < t_c = 2R_H/c$, the major effect of compressibility is that sound waves carry away a fraction of the particle momentum with the sound speed c. The VACF quickly decays from its initial value to $C(t_c) \approx k_BT/\tilde{m}$, where $\tilde{m} = m_e + d\, m_f/(d-1)$ includes an "added mass" $m_f/(d-1)$ that comes from the fluid around the particle that has to move with the particle because of incompressibility [68, 87, 88]. The initial decay of the VACF due to sound waves will appear to be instantaneous (discontinuous) if one increases the speed of sound to infinity. The incompressible

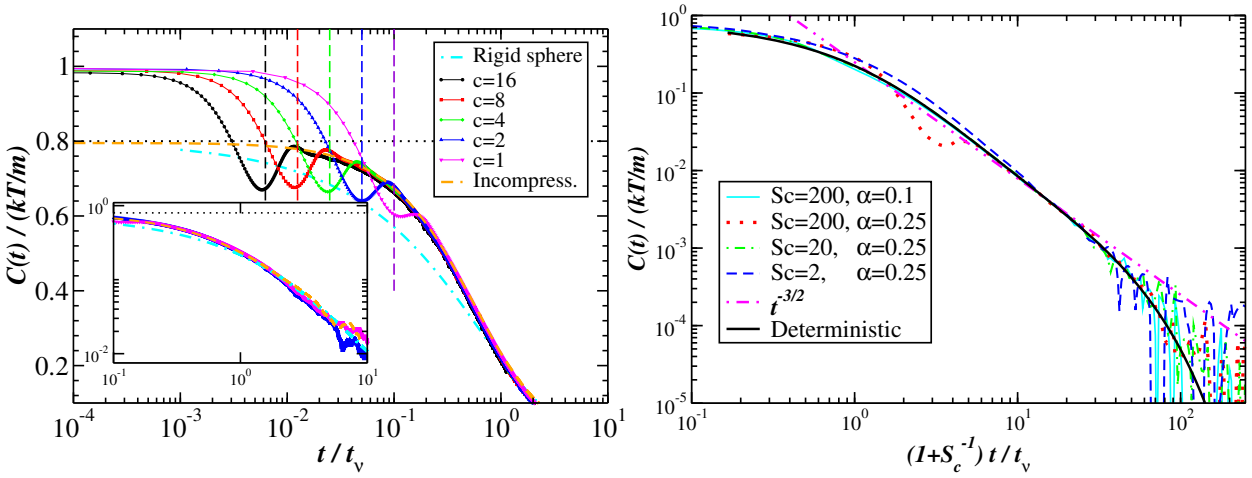

Figure 4: Velocity autocorrelation function (VACF) (55) of a single particle with excess mass $m_e = m_f$, normalized by k_BT/m so that it should be unity at the origin for a compressible fluid. (*Left panel*) Comparison between a compressible fluid for several different speeds of sound c (compressibilities), as well as an incompressible fluid ($c \to \infty$). Vertical lines indicate the sound time scale $t_c = 2R_H/c$, and the asymptotic power-law tail $(t/t_\nu)^{-3/2}$ is emphasized in the inset, where $t_\nu = R_H^2/\nu$ is the viscous time scale. The tail matches the theoretical predictions for a rigid sphere with the same effective mass immersed in an incompressible fluid [87]. All runs use a small time step size so the dynamics is well-resolved at short times. (*Right panel*) Comparison between different Schmidt numbers $S_c = \nu/\chi$ for an incompressible fluid. A deterministic calculation, corresponding to the limit $S_c \to \infty$, is also shown. In the legend, the time step size is expressed in terms of the advective CFL number α. The scaling of the time axes is adjusted to overlap the power-law tails (see text).

Parameter	Fixed S_c runs	Variable S_c runs
excess mass m_e	$m_e = m_f$	$m_e = m_f$
grid spacing Δx	10.543	1
grid size N	41^3	32^3
fluid density ρ	1	1
shear viscosity η	0.5	variable
bulk viscosity ζ	0.5	not relevant
speed of sound c	$1 - \infty$	∞
Temperature k_BT	1	0.1
viscous CFL β	$10^{-5} - 10^{-3}$	$\beta = \alpha\sqrt{S_c/(6\pi)}$
Sound CFL α_s	$0.05 - 0.1$	not relevant
Schmidt number S_c	48.2	variable

Table II: Parameters used in the compressible and incompressible simulations shown in the left panel of Fig. 4 (middle column), as well as the incompressible simulations shown in the right panel of Fig. 4 (right column).

inertial coupling method should produce an intercept $C(0^+) = k_BT/\tilde{m}$ in agreement with (24), rather than the equipartition result valid for a compressible fluid. For example, for $m_e = m_f$ and $d = 3$, we expect $C(0^+) = 0.8\, k_BT/m$. This illustrates the subtle complexity of coupling a fluctuating fluid solver to immersed particles, even in the absence of external and interparticle forces. In the left panel of Fig. 4 we show numerical results for the VACF for several different speeds of sound, obtained using the algorithm summarized in Appendix D. The approach to the incompressible limit $c \to \infty$ is evident in the figure, and it is clear that our incompressible inertia coupling method correctly reproduces the limiting behavior (without however suffering from the severe time step limitation of compressible flow solvers). The parameters for these simulations are shown in Table II.

At the viscous time scale, $t > t_\nu = \rho R_H^2/\eta$, conservation of momentum (hydrodynamics) in the

fluid introduces a memory in the motion of the particle and the VACF decays with a well-known asymptotic power-law tail $\sim (t/t_\nu)^{-d/2}$ [89]. Any numerical method that solves the time-dependent Navier-Stokes equation (as opposed to the steady or time-independent Stokes equation, as Brownian or Stokesian dynamics do) ought to reproduce this power-law decay. The amplitude of the decay depends on the shape of the particle and is well-known for the case of a rigid sphere with stick boundaries [87]. We expect that the VACF for our blob particle will have the same value at the origin and the same power-law tail as a rigid sphere of radius R_H and the same ratio m_e/m_f. For the "equivalent" rigid sphere we take $m_f = \rho V_s$ where V_s is the volume of the sphere, and $m_e = (\rho_s - \rho)\, V_s$, where ρ_s is the density of the sphere. The exact shape of $C(t)$ will in general be different between a blob and a rigid sphere. In the inset of the left panel of Fig. 4 we compare the long-time behavior of the VACF of a blob particle to that of a rigid sphere, and see the same power-law behavior. Note that the rigid sphere theory shown here does not account for finite-size effects. In these tests we use a small time step size in order to study the properties of the spatial discretization in the absence of temporal truncation errors.

At long times, the motion of the particle is diffusive with a diffusion coefficient predicted by the Stokes-Einstein relation to be

$$\chi \approx \chi_{\text{SE}} = \frac{k_B T}{6\pi\eta R_H} \tag{56}$$

in three dimensions, where we recall that for the particular spatial discretization we employ $R_H \sim \Delta x$ in three dimensions. The exact coefficient depends on the system size, and for the system size we use here an excellent approximation is $R_H \approx \Delta x$. We define here a dimensionless Schmidt number S_c based on the Stokes-Einstein diffusion coefficient (56),

$$\frac{\nu}{\chi} \approx S_c = \frac{\nu}{\chi_{\text{SE}}} = \frac{6\pi\eta^2 R_H}{\rho k_B T} \approx \frac{6\pi\beta^2}{\alpha^2}.$$

The Schmidt number is an important quantity that measures how fast momentum diffuses relative to the particles. In many cases of interest $S_c \gg 1$, which means that the dynamics of the particles approaches the Brownian (overdamped) limit [27]. Note that the limit $S_c \to \infty$ is the same as the deterministic limit $k_B T \to 0$, in which fluctuations become a very weak perturbation to the deterministic dynamics. Another important dimensionless number is the "thermal" Peclet number

$$\text{Pe} = \frac{V\, R_H}{\chi_{SE}} \approx \sqrt{6\pi S_c},$$

which measures the relative importance of advection by the thermal velocity fluctuations to diffusion. We see that $\text{Pe} \sim S_c^{1/2}$ is directly related to the Schmidt number. A similar calculation also shows that the "thermal" Reynolds number is $\text{Re} \sim S_c^{-1/2}$. Therefore S_c is the only relevant dimensionless number for a particle diffusing in a fluid at thermodynamic equilibrium.

It is important to test how well our algorithm works for a range of Schmidt numbers. We expect the case of small S_c to be the most difficult in terms of accuracy, since the particle can move a substantial distance (compared to the grid spacing) during a single time step, $\alpha = O(1)$, and the thermal and cell Reynolds number is also $r = O(1)$. The case of large S_c, on the other hand, is the most demanding in terms of computational effort since particles barely move during a single time step, $\alpha \ll 1$, and $O(S_c)$ fluid time steps may be required to reach the diffusive time scale for $\beta = O(1)$. In order to investigate the long-time behavior we try to maximize the time step, but always keeping $\alpha < 1$, specifically, here we set $\alpha = 0.25$. For the largest Schmidt number we investigate, $S_c \approx 200$, this value of α corresponds to a relatively large $\beta \approx 0.81$, and therefore we also try a smaller time step, corresponding to $\alpha = 0.1$ and $\beta \approx 0.33$. The temporal integrator developed here cannot be used for $\beta \gtrsim 1$ because the Crank-Nicolson temporal integrator we use for the velocity equation does not accurately resolve the dynamics of the small wavelength fluid modes [28, 52].

In the right panel of Fig. 4 we show the VACF for several viscosities and thus Schmidt numbers. The parameters for the runs are given in Table II. The standard theory for the tail of the VACF (long-time behavior) [87] implicitly assumes that $S_c \gg 1$, and leads to the conclusion that for an isolated particle in an infinite fluid asymptotically $C(t) \approx (t/t_\nu)^{-d/2} \sim (\nu t)^{-d/2}$. A more complete self-consistent mode coupling theory [90] corrects this to account for the fact that while momentum diffuses around the particle the particle itself diffuses, and predicts that $C(t) \sim [(\chi + \nu)\, t]^{-d/2}$ [89]. This means that we

expect the tails of the VACFs for different S_c values to collapse on one master curve if we plot them as a function not of (t/t_ν) but rather of $(1 + S_c^{-1})\,(t/t_\nu)$. This is confirmed in the right panel of Fig. 4.

It is evident in Fig. 4 that these fluctuating calculations lead to noisy results for the tail of the VACF, making it difficult to see the behavior of the long-time behavior. Many researchers have chosen to calculate the VACF by performing a *deterministic* calculation, in which the particle is given a small initial kick in velocity, and then the deterministic algorithm is used to track the subsequent decay of the velocity. This is sometimes done because thermal fluctuations are not consistently included in the algorithm, or because the deterministic calculation is much faster and more accurate, not requiring as much statistical averaging. In the right panel of Fig. 4 we show the VACF obtained from a deterministic test, which can be thought of as the VACF in the limit of vanishing fluctuations, $k_B T \to 0$ (equivalently, $S_c \to \infty$). Note that in the deterministic test the magnitude of the initial velocity $\boldsymbol{u}_0$ of the particle has to be chosen to match the thermal kick, $\|\boldsymbol{u}_0\|^2 = d\,k_B T/\tilde{m}$, which in practice means that the deterministic VACF has to be scaled so that it agrees with statistical mechanics at the origin. Due to the slight anisotropy and imperfect translational invariance of the spatial discretization, in principle even the deterministic result should be averaged over many initial positions and velocities of the particle. The VACF in the limit $S_c \to \infty$ shown in Fig. 4 matches the fluctuating runs for the larger Schmidt numbers. Due to the lack of noise, it also clearly shows the long-time exponential decay in the VACF at times $t \sim L^2/\nu$ [86], where finite-size effects become important.

C. Small Reynolds Number

In this section we focus on the ability of the blob model to reproduce some important features of the interaction of rigid spheres with deterministic low Reynolds number flow. Maxey and collaborators have performed detailed investigations of the low Reynolds number hydrodynamics for Gaussian blobs [16, 29, 30] in the context of the Force Coupling Method (FCM). They have already demonstrated that a blob model can model the behavior of hard sphere colloidal suspensions with remarkable fidelity given its minimal resolution. Because our blob is not Gaussian and our spatial discretization is very different from that employed in the FCM, we examine briefly the flow around a single particle and the hydrodynamic interactions between a pair of particles at very small Reynolds number.

1. Stokes flow around a blob

An important property of the blob is its hydrodynamic radius, which is defined in three dimensions from Stokes law for the drag force $\boldsymbol{F}_d = 6\pi\eta R_H \boldsymbol{v}$ experienced by the particle as it moves slowly through an unbounded fluid at rest far away from the particle. One can also compare the steady Stokes flow around the blob with the well-known analytical solution for the flow around a rigid sphere. These types of calculations were performed in detail for a compressible fluid in Ref. [43], and since the Mach number was kept small, very similar results are to be expected for an incompressible fluid. Here we briefly examine the steady Stokes flow around a particle as a validation of the incompressible formulation and implementation.

We exert a constant force density (pressure gradient) on the fluid in a periodic domain of 128^2 hydrodynamic cells, and attach a single blob to a tether point via a stiff elastic spring. After an initial transient, a steady state is reached in which the drag force on the particle is balanced by the spring force. The forcing is chosen so that the Reynolds number is kept small, $Re = 2\Delta x\, v_{max}/\nu < 0.002$. Because of the long-ranged r^{-1} decay of the flow away from the particle, there are strong finite-size corrections that are well-known [43]. Taking into account these corrections we estimate $R_H \approx 0.91\Delta x$, consistent with the more careful estimates obtained in Ref. [43] using a compressible flow simulation and a non-periodic domain. We emphasize again that R_H is not perfectly translational invariant and changes by a couple of percent as the particle moves relative to the underlying fluid solver grid. This is illustrated in the inset in the left panel of Fig.5. If the Peskin four-point interpolation function [32] is used instead of the three-point function, a smaller variance in the hydrodynamic radius (i.e., improved translational invariance) would be observed [3].

In the left panel of Fig. 5 we compare the radial component of the fluid velocity $u_r(r)$ along $\theta = 0$ (direction of motion of the incoming flow) and along $\theta = \pi/4$ with the analytical solution for a solid sphere with no-slip surface in a infinite system, as a function of the distance d from the particle center. A surprisingly good agreement is observed even for distances as small as $d = 2R_H$, in agreement with previous investigations for Gaussian blobs [16, 29, 30]. Note, however, that there is flow penetrating

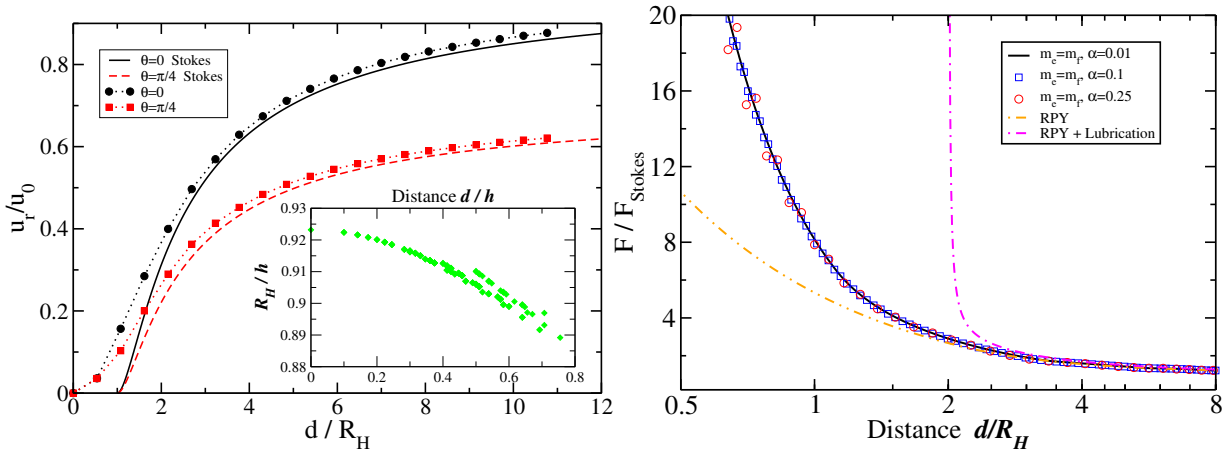

Figure 5: (*Left panel*) Comparison of the radial component of the velocity in steady flow around a fixed blob to the Stokes solution for a rigid sphere in an unbounded domain. The velocity at a given point $\boldsymbol{q}'$ is obtained by using the local averaging operator, $\boldsymbol{v}(\boldsymbol{q}') \approx \boldsymbol{J}(\boldsymbol{q}')\,\boldsymbol{v}$. The inset shows the hydrodynamic radius of the particle for flow along the x axes, as a function of the distance between it and the center of the nearest x face of the grid (location of the x component of velocity). (*Right panel*) Hydrodynamic interaction force between two particles as a function of the interparticle distance for several time step sizes. For large distances the Stokes mobility is recovered, and for moderate distances the next-order correction (Rotne-Prager mobility) is recovered, independent of dynamical parameters. At close distances a large increase in the hydrodynamic force is observed, as in the Rotne-Prager-Yamakawa (RPY) mobility (also shown). Note that the sharp increase in the hydrodynamic interaction at short distances is qualitatively similar but distinct from the divergent lubrication force observed between two rigid spheres (dashed-dotted line).

the blob at distances $d < R_H$, unlike a rigid sphere. Also note that similar calculations performed using a frictional coupling [3] in Ref. [43] clearly reveal much larger penetration of the flow into the blob, unless a rather large friction constant is employed.

2. Hydrodynamic Interactions

In this section we investigate the hydrodynamic interaction force between two particles in the deterministic setting, as done for a compressible fluid in Ref. [43]. In this test, we apply a force $\boldsymbol{F}_0$ on one particle toward the other particle, and the opposite force $-\boldsymbol{F}_0$ on the other particle, so that the center of mass remains at rest. The applied force is weak so that the Reynolds number $Re \leq 10^{-3}$ and the flow is in the Stokes (steady-state) limit. As the particles approach, we measure the relative speed of approach v_r and compare it to the prediction of Stokes's law,

$$\frac{F}{F_{\text{Stokes}}} = -\frac{2F_0}{6\pi\eta R_H v_r}. \tag{57}$$

In the right panel of Fig. 3 we show results for F/F_{Stokes} as a function of the interparticle distance l. The simulation parameters are reported in Table I. We performed several simulation to verify that the results presented here are independent of the excess mass and time step size. In the figure we compare the results of our calculations to a theoretical calculation based on the Rotne-Prager-Yamakawa (RPY) tensor for the mobility of a pair of particles in a periodic system [91]. We see that the RPY correction correctly captures the behavior for the blobs for distances $d \gtrsim 3$. At short distances there is a strong repulsive force between the blobs similar to the well-known "lubrication force" that develops as an incompressible fluid is squeezed out between two approaching rigid spheres. However, unlike the lubrication force between two rigid spheres, the hydrodynamic interaction force between two blobs does not diverge like $(d - 2R_H)^{-1}$. This is expected because blobs do not have a well-defined surface; however, they are not point particles either, and they do squeeze the incompressible fluid in-between them as they approach each other.

In fact, at short distances the hydrodynamic interaction between blobs is similar to that for the RPY tensor (see Eq. (8) in Ref. [92]), also shown in the right panel of Fig. 3. An examination of the derivation of the RPY mobility (see Eqs. (6) and (7) in Ref. [92]) reveals that the RPY correction

arises due the Faxen term in (6). Therefore, the fact that we get such good agreement between RPY and the numerical results for blobs at larger interparticle distances shows that the Faxen radius of our blob is very close to its hydrodynamic radius, $a_F \approx R_H \approx 0.9\Delta x$. Numerical investigations show that the Faxen radius of the blob (which can be expressed in terms of the second moment of the discrete kernel function) is translationally-invariant to within about 5%. Note that using kernel functions that try to approximate a Dirac delta function to higher accuracy will not give the correct Faxen term since the Faxen radius of a true point particle is zero. In particular, for the Peskin 6pt function [32] the second moment and thus the Faxen radius is identically zero. By contrast, the 3pt function used here gives an excellent approximation to a rigid sphere at intermediate and large distances.

D. Large Reynolds Number

As discussed in the Introduction, we believe that the blob particle model can be successfully used for simulations of particle-laden flows at moderate and large particle Reynolds number. In the $\mathrm{Re}_P \to 0$ limit, we showed that the perturbative flow created by the blob particle agrees with the Stokes solution for the steady flow around a no-slip sphere of hydrodynamic radius $R_H = 0.91\,h$ for the three point kernel and a staggered grid solver [43]. This is consistent with other calculations utilizing a four-point kernel and a non-staggered grid [77]. However it is not *a priori* clear how consistently the blob hydrodynamic behavior will be at large Re_P. The drag force provides a non-trivial test because it captures the average effect of the perturbative flow. In a previous study [43] we observed that the drag on the blob is consistent with that of a rigid sphere up to $\mathrm{Re}_P < 10$. This study was performed using a compressible flow solver, and ensuring a low Mach number was prohibitively expensive at larger Re_P. Using the incompressible (zero Mach number) solver developed here we now study larger Re_P.

In the top left panel of Fig. 6 we show the drag force on a blob particle in a periodic domain as a function of Re, normalized by the Stokes limit drag ($Re \to 0$). We estimate the drag coefficient by dragging the blob with a constant applied force and measuring the average velocity $v_0 = \langle u \rangle$ along the direction of the applied force at long times. The particle interacts with its periodic images approximately after time $\tau_L = L/v_0$ (where L is the box length in flow direction), while viscous transport around the blob requires a time longer than $\tau_\nu = R_H^2/\nu$ to settle down. In order to mimic the behavior of an isolated blob in an infinite medium, we must have $\tau_L/\tau_\nu > 1$, or, equivalently, $L/R_H > \mathrm{Re}_p$. We therefore performed the calculations in a box of $2^n \times 32 \times 32$ cells, with $2^n > 3\mathrm{Re}_P$; the size of the simulation is indicated in the legend in the figure. In the top left panel of Fig. 6 we compare the numerical results with the empirical law for the drag on a rigid sphere with a no-slip surface [93]. The agreement is remarkably good over the studied range $0 \leq \mathrm{Re}_P \leq 324$.

Encouraged by this result we now briefly analyze the structure of the perturbative flow around the blob, and defer a more detailed study for future work. In the remaining panels in Fig. 6 we show iso-contours of the vorticity perpendicular to the snapshots' plane at a few values of Re_P. For $\mathrm{Re}_P \simeq 0$ the fully symmetric pattern observed around the blob agrees with the Stokes solution at distances larger than $2R_H$ [43]. The fore-and-aft symmetry of the Stokes flow becomes distorted by advection at moderate Reynolds $\mathrm{Re}_P \sim 1$ leading to the so-called Oseen flow [16], as observed in the top right panel of Fig.6 for $\mathrm{Re}_P = 1.5$. For $20 < \mathrm{Re}_P < 270$ a transition with symmetry breaking leads to a steady axisymmetric "double-thread" structure [94, 95]. Although we have not studied the transition in detail, our blob model qualitatively reproduces a steady axisymmetric wake with a bifid vortex trail, as illustrated in the bottom left panel of Fig. 6 for $\mathrm{Re}_P = 137$. This type of wake is observed above $\mathrm{Re}_P > 10$ and its topology is maintained at least up to $\mathrm{Re}_P \simeq 300$. For flow around a rigid sphere a Hopf transition to oscillatory flow leading to vortex shedding takes place at $\mathrm{Re}_P \approx 270$ [94, 95]. For the blob we observe small oscillations of the wake, without vortex shedding at $\mathrm{Re}_P > 300$, indicating that some possible transition to unsteady flow could be induced by a small perturbation. In fact, in particle-laden flows typically $\mathrm{Re}_P < 100$ [34], and the induction of vortex shedding due to perturbations from the wake of other particle is a more relevant vorticity source. Using smaller boxes, where the particle interacts with its image, we frequently observed induced vortex shedding for $\mathrm{Re}_P > 70$, as illustrated in the bottom right panel of Fig. 6.

It is a remarkable fact that the ICM blob minimal-resolution model can produce wakes containing many of the features of realistic flows. The thickness of the viscous (Oseen) layer around a sphere decreases like R/Re_P [16]; therefore, this layer is completely unresolved for $\mathrm{Re}_P > 1$ in our model. However, the "local" no-slip constraint is able to capture the non-linear velocity-pressure coupling

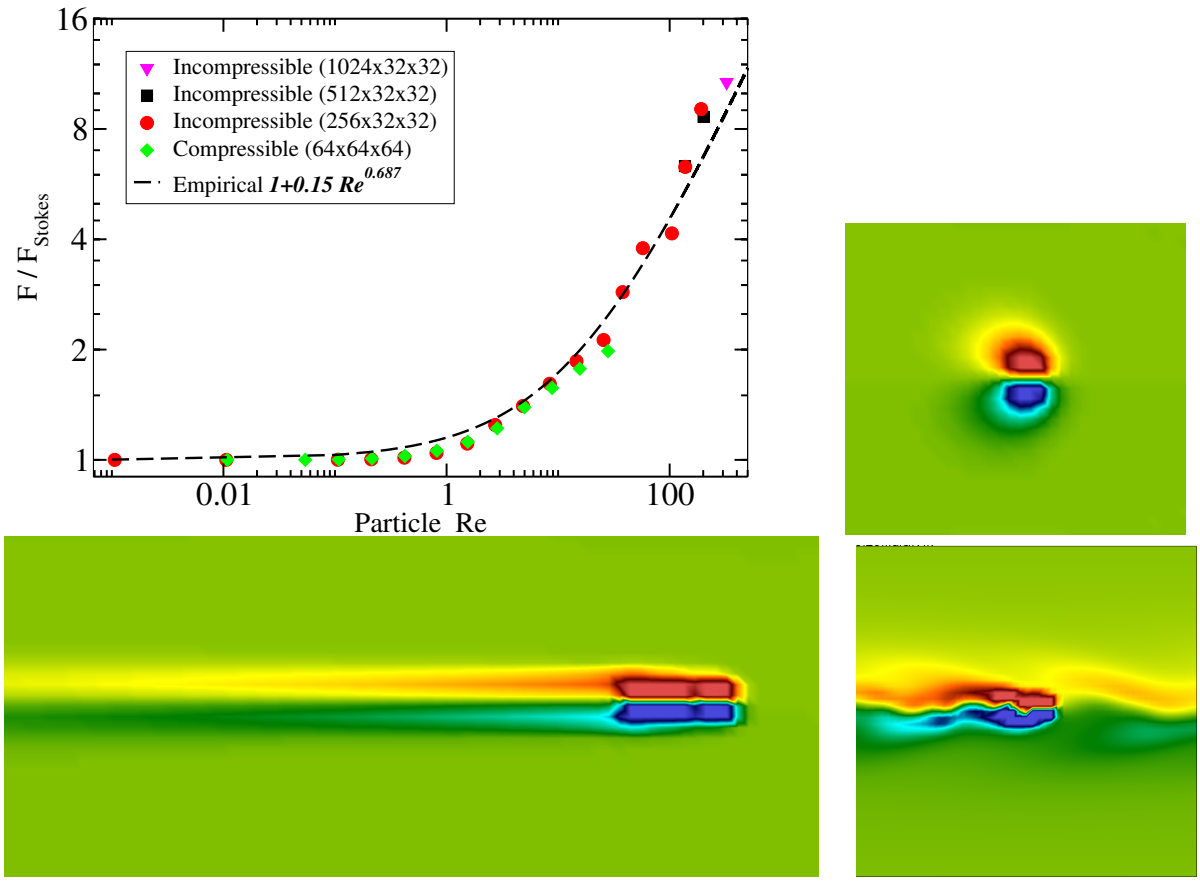

Figure 6: (*Top left*) Drag force on a blob particle in a periodic domain as a function of the particle Reynolds number $\mathrm{Re}_P = 2R_H \langle u \rangle / \nu$, normalized by the Stokes drag ($\mathrm{Re}_P \to 0$) (symbols). Results for the incompressible and compressible solvers (see Appendix D and Ref. [43]) are compared with the empirical law for rigid no-slip spheres (line) [93]. The size of the domain box in cells is indicated in the legend. (*Top right*) Out-of-plane vorticity iso-contours at $\mathrm{Re}_P = 1.5$ for a box of 32^3 cells, and (*Bottom left*) at $\mathrm{Re}_P = 137$ in a long box of $512 \times 32 \times 32$ cells. (*Bottom right*) A snapshot of an unsteady (nearly oscillatory) flow in a box of 32^3 cells at $\mathrm{Re}_P = 70$ where the particle sheds vortices due to the interaction with the wake from its image. The simulation parameters are as in Table I, but $m_e = 0$ and advective CFL $\alpha = 0.2$. Error bars are estimated to be less than 5%.

which dominates the drag at large Re_P. A somewhat similar scenario was observed in simulations of ultrasound-particle interaction [43]. In the inviscid regime (sound frequency faster than ν/R_H^2), an excellent agreement with the theory was observed even in cases where the sound-viscous boundary layer around the blob was unresolved.

V. Conclusions

In this work we described a bidirectional coupling between a point-like "blob" particle and an incompressible fluctuating fluid, building on prior work by some of us in the compressible setting [43]. At the continuum level, the proposed model includes inertial and stochastic effects in a consistent manner, ensuring fluctuation-dissipation balance and independence of equilibrium thermodynamic properties on dynamical parameters. We constructed a second-order spatio-temporal discretization that tries to preserve the properties of the continuum model as well as possible.

Through numerical experiments, we showed that the proposed *inertial coupling method* (ICM) can consistently describe the dynamics of "blob" particles in fluid flow over a broad range of particle and fluid Reynolds numbers, from Brownian motion to convective regimes. We demonstrated that the method reproduces well-known non-trivial effects of particle inertia on the short-time dynamics of Brownian particles, while also capturing the long-time (Brownian) diffusive dynamics and the associated equilibrium distribution correctly. Remarkably, we found the minimally-resolved blob model to reproduce

non-trivial features of turbulent flow around a rigid sphere, including the non-Stokes drag at moderately large Reynolds numbers $\mathrm{Re}_P < 300$ and interactions between particle wakes. As such we believe that the method presented here can be applied to model the dynamics of dilute and semi-dilute colloidal suspensions and polymeric fluids over a broader range of conditions than existing methods [3, 21, 24, 25, 27].

The algorithm we described here was specifically optimized for periodic boundary conditions. In particular, we constructed a semi-implicit temporal integrator that only requires a few applications of the FFT algorithm per time step. This enabled us to implement the algorithm on GPU platforms, achieving excellent performance with little development effort. However, this simplicity was not without a cost. Firstly, we had to make several approximations in order to avoid iterative linear solvers and use a fixed number of FFTs per time step. Notably, we had to assume that particles were far away from other particles compared to their size (dilute conditions). Secondly, in many problems of interest the fluid is confined in non-periodic geometries such as channels and periodic boundary conditions are not appropriate. It is not difficult, at least in principle, to adopt our algorithm to non-periodic geometries and to dense collections of blobs. This requires developing specialized preconditioned Krylov linear solvers for solving the "inertial" Stokes problem (36) in non-periodic geometries, similar to the fluid-only solver developed by one of us [75] and implemented in the IBAMR framework [74]. In future work we will explore effective preconditioners for solving (36) using a Krylov method.

The temporal integrator algorithm we developed in this work can accurately resolve the inertial, viscous, and fluctuating dynamics of particles immersed in an incompressible fluid if the viscous and advective CFL numbers are less than unity. Some further work is required to tackle the large *separation of time scales* present in many situations of practical interest. For example, in the Brownian dynamics limit there is an increasing separation of scales between the particle movement, the viscous damping, and the inertial dynamics. For the second-order temporal integrator that we presented here, the case of large Schmidt number $S_c \gg 1$ is the most demanding in terms of computational effort. This is because particles barely move during a single time step, $\alpha \ll 1$, and $O(S_c)$ fluid time steps may be required to reach the diffusive time scale for $\beta = O(1)$. For the case of a neutrally-buoyant particle ($m_e = 0$) and periodic boundary conditions Atzberger *et al.* [28] have developed a specialized exponential integrator that tackles the large separation of time scales that arise when $S_c \gg 1$. Future work will consider extending such techniques to non-periodic systems and to the case $m_e \neq 0$.

In ICM the fluid-particle coupling is expressed as a no-slip constraint equating the translational particle velocity with the local fluid velocity. This implies that only the *monopole* term (*stokeslet*) is included in the fluid-particle force. As a consequence the present approach can only accurately resolve the fluid flow at distances larger than the typical size of the particles. However, it is a remarkable fact that with such minimal resolution permits to capture so many hydrodynamic effects in a qualitatively, and, in some cases, quantitatively correct way (see Sec. IV and also Ref. [43]). Even though lubrication flows (at low Re_P) or viscous boundary layers (at large Re_P and/or high forcing rates) are unresolved, the locally averaged no-slip constraint proves to be remarkably robust in capturing their essential hydrodynamics and permits to go beyond the Stokes limit. Crucial to this success is the fact that we employ a carefully-constructed spatial discretization, combining nearly grid-invariant Lagrangian-Eulerian transformations, local energy and momentum conservation, and a staggered discretization of the incompressibility condition and the fluctuating stresses.

It is not difficult to extend our approach to also include the anti-symmetric component of the dipole (*rotlet*) stress [29]. Firstly, particle rotational degrees of freedom would need to be added to the blob description, along with an angular velocity $\boldsymbol{\omega}$ and an associated excess moment of inertia. We would need to impose an additional *rotational no-slip* constraint, requiring that the particle rotate with the locally-averaged angular velocity of the fluid, $\boldsymbol{\omega} = \boldsymbol{J}\left(\boldsymbol{\nabla} \times \boldsymbol{v}\right)/2$, and distribute the fluid-particle torque $\boldsymbol{\tau}$ (Lagrange multiplier) as a force density $\boldsymbol{f}_{\boldsymbol{\tau}} = -\boldsymbol{\nabla} \times \left(\boldsymbol{S}\boldsymbol{\tau}\right)/2$ in the fluid momentum equation. This type of approach has already been employed in the deterministic context to model suspensions of neutrally-buoyant semi-rigid rods [96, 97]. To our knowledge, such an approach has not yet been applied to fluctuating hydrodynamics and the resulting rotational diffusion has not been investigated.

An additional *rigidity constraint* would be required to also constrain the locally-averaged deformation tensor, and thus consistently include the *symmetric* components of the dipole (*stresslet*) force terms, as proposed by Maxey and collaborators in the context of the deterministic Force Coupling Method (FCM) [29]. Unlike the smooth Gaussian kernels used in the FCM [29], the existing Peskin

kernels do not have well-behaved derivatives. Therefore, it appears necessary to generalize the IBM kernels to enable the local averaging ("interpolation") of spatial derivatives in a reasonably translationally-invariant manner. The inclusion of thermal fluctuations of the stresslet requires careful consideration even at the continuum level, and will be the subject of future research.

In a different spirit, the approach used here for single "blobs" can be extended to account for the finite extent and shape of arbitrarily-shaped rigid bodies immersed in fluid flow. One approach that has been successfully employed in the deterministic setting is to construct the immersed body out of a collection of blobs constrained to move rigidly [98]. Future work will consider the inclusion of thermal fluctuations in this type of approach, and the minimal amount of resolution required to capture the geometry of immersed particles.

At the other extreme, our method can be used to provide a coarse-grained model for hydrodynamics at very small scales, for example, for the Brownian motion of a small molecule suspended in a simple solvent. At such small scales, the suspended particles do not have a fixed, or even a well-defined shape, and there are many competing and sometimes canceling effects: normal and tangential slip [99], breakdown of Navier-Stokes hydrodynamics, non-Gaussian fluctuations, non-Markovian effects, etc. Our simple fluid-particle coupling model can be used to isolate hydrodynamic from non-hydrodynamic effects and study basic physics questions about the importance of inertia and fluctuations on Brownian motion, going beyond the uncontrolled approximations required by existing theoretical approaches. Notably, preliminary investigations have shown that for small Schmidt numbers nonlinear effects become important and lead to a non-trivial contribution of the thermal fluctuations to the mean fluid-particle force. Specifically, the mobility of a particle in a fluctuating fluid was found to differ from that in a deterministic (Stokes) fluid, thus leading to a deviation of the diffusion coefficient from the standard Stokes-Einstein prediction. More careful investigations are required to assess how well self-consistent mode-coupling theories can model this effect, as well as to study the influence of particle inertia (density contrast), random slip (see Appendix B), and spatial dimensionality.

Acknowledgments

We thank Alejandro Garcia, Charles Peskin, Paul Atzberger, Eric Vanden-Eijnden, Martin Maxey, Tony Ladd and Burkhard Dünweg for informative and inspiring discussions. A. Donev was supported in part by the Air Force Office of Scientific Research under grant number FA9550-12-1-0356. B. Griffith acknowledges research support from the National Science Foundation under awards OCI 1047734 and DMS 1016554. R. Delgado-Buscalioni and F. Balboa acknowledge funding from the Spanish government FIS2010-22047-C05 and from the Comunidad de Madrid MODELICO-CM (S2009/ESP-1691). Collaboration between A. Donev and R. Delgado-Buscalioni was fostered at the Kavli Institute for Theoretical Physics in Santa Barbara, California, and supported in part by the National Science Foundation under Grant No. NSF PHY05-51164.

Appendix

A. Periodic Boundary Conditions

The pressure-free fluid-only equation (18) can formally be written in a form suitable for direct application of standard numerical methods for integrating initial value problems,

$$\partial_t \boldsymbol{v} = \boldsymbol{\rho}_{\text{eff}}^{-1} \left(\boldsymbol{\mathcal{P}} \boldsymbol{f} + \boldsymbol{\mathcal{P}} \boldsymbol{S} \boldsymbol{F} \right), \tag{A1}$$

although this form is only useful if one can actually compute the action of the inverse inertia tensor $\boldsymbol{\rho}_{\text{eff}}^{-1}$. It turns out that this is possible for periodic systems.

To see this, we expand $\boldsymbol{\rho}_{\text{eff}}^{-1}$ into a formal series,

$$\boldsymbol{\rho}_{\text{eff}}^{-1} = \rho^{-1} \left[\boldsymbol{I} - m_e \rho^{-1} \boldsymbol{\mathcal{P}} \boldsymbol{S} \boldsymbol{J} \boldsymbol{\mathcal{P}} + \left(m_e \rho^{-1} \boldsymbol{\mathcal{P}} \boldsymbol{S} \boldsymbol{J} \boldsymbol{\mathcal{P}} \right)^2 - \left(m_e \rho^{-1} \boldsymbol{\mathcal{P}} \boldsymbol{S} \boldsymbol{J} \boldsymbol{\mathcal{P}} \right)^3 + \dots \right]. \tag{A2}$$

Observe that $\left(\boldsymbol{\mathcal{P}} \boldsymbol{S} \boldsymbol{J} \boldsymbol{\mathcal{P}} \right)^n = \boldsymbol{\mathcal{P}} \boldsymbol{S} \left(\boldsymbol{J} \boldsymbol{\mathcal{P}} \boldsymbol{S} \right)^{n-1} \boldsymbol{J} \boldsymbol{\mathcal{P}}$ involves the $(n-1)$-st power of the $d \times d$ matrix $\Delta \widetilde{\boldsymbol{V}}^{-1} = \boldsymbol{J} \boldsymbol{\mathcal{P}} \boldsymbol{S}$, where we made use of the fact that $\boldsymbol{\mathcal{P}}^2 = \boldsymbol{\mathcal{P}}$. Recall from (2) that $\boldsymbol{J} \boldsymbol{S} = \Delta V^{-1} \boldsymbol{I}$ is a related to the inverse of the particle volume, which is independent of the position of the particle for the particular kernel function used herein. In principle, the matrix $\Delta \widetilde{\boldsymbol{V}}^{-1}$ could depend on the position

of the particle because of the appearance of $\boldsymbol{\mathcal{P}}$, which implicitly encodes the boundary conditions. However, periodic systems are translationally invariant and therefore $\Delta\widetilde{\boldsymbol{V}}^{-1}$ cannot depend on the position of the particle and is simply a constant $d \times d$ matrix.

By performing a relatively-straightforward calculation in Fourier space it is possible to show that for periodic systems much larger than the kernel extent $\Delta\widetilde{\boldsymbol{V}}^{-1}$ is simply a multiple of the identity matrix,

$$\Delta\widetilde{\boldsymbol{V}}^{-1} = \boldsymbol{J}\boldsymbol{\mathcal{P}}\boldsymbol{S} = \frac{d-1}{d}\Delta V^{-1}\boldsymbol{I}. \tag{A3}$$

The prefactor $(d-1)/d$ in the relation accounts for the elimination of the longitudinal (compressible) velocity mode by the projection operator. We can therefore define the *effective* particle volume accounting for incompressibility to be

$$\Delta\widetilde{V} = \frac{d}{d-1}\Delta V,$$

and accordingly, define the effective mass of the fluid dragged with the particle to be $\tilde{m}_f = \rho\Delta\widetilde{V} = d\,m_f/(d-1)$, and the effective particle inertia to be $\tilde{m} = \tilde{m}_f + m_e$. In three dimensions, the added fluid mass due to incompressibility is $m_f/2$, which is a well-known result for rigid spheres immersed in an incompressible fluid [87, 88].

Note that for periodic systems where some of the dimensions of the unit cell are comparable to the kernel width (A3) is only an approximation. Notably, for a three-dimensional periodic box of shape $L_x \times L_y \times L_z$, if $L_z \ll a$ the value of $\Delta\widetilde{\boldsymbol{V}}^{-1}$ converges to the two-dimensional result,

$$\Delta\widetilde{\boldsymbol{V}}_{\text{2D}}^{-1} = \Delta V^{-1}\begin{bmatrix} \frac{1}{2} & 0 & 0 \\ 0 & \frac{1}{2} & 0 \\ 0 & 0 & 1 \end{bmatrix}, \tag{A4}$$

rather than the three-dimensional $\Delta\widetilde{\boldsymbol{V}}_{\text{3D}}^{-1} = (2\Delta V^{-1}/3)\,\boldsymbol{I}$.

Using (A3) it is possible to simplify the infinite series (A2) and obtain the important result for a single particle immersed in a periodic fluid,

$$\boldsymbol{\rho}_{\text{eff}}^{-1} = (\rho\boldsymbol{I} + m_e\boldsymbol{\mathcal{P}}\boldsymbol{S}\boldsymbol{J}\boldsymbol{\mathcal{P}})^{-1} = \rho^{-1}\left(\boldsymbol{I} - \frac{m_e\Delta\widetilde{V}}{\tilde{m}}\boldsymbol{\mathcal{P}}\boldsymbol{S}\boldsymbol{J}\boldsymbol{\mathcal{P}}\right). \tag{A5}$$

Using this result we can rewrite (A1) in the simple form

$$\rho\partial_t\boldsymbol{v} = \boldsymbol{\mathcal{P}}\left(\boldsymbol{I} - \frac{m_e\Delta\widetilde{V}}{\tilde{m}}\boldsymbol{S}\boldsymbol{J}\right)\boldsymbol{\mathcal{P}}\boldsymbol{f} + \frac{\tilde{m}_f}{\tilde{m}}\boldsymbol{\mathcal{P}}\boldsymbol{S}\boldsymbol{F}, \tag{A6}$$

which is useful for analysis and for numerical approximations.

It is important to point out, however, that when there are many particles present, there is no simple formula for $\boldsymbol{\rho}_{\text{eff}}^{-1}$. That is, we cannot just add a summation over all particles in (A5). This is because there are cross-terms between two particles i and j in the infinite series (A2) involving the operator $\boldsymbol{J}_i\boldsymbol{\mathcal{P}}\boldsymbol{S}_j$. If the particles are *not* overlapping, meaning that the kernel functions of the different particles have disjoint support, then $\boldsymbol{J}_i\boldsymbol{S}_j = \boldsymbol{0}$, however, this is not true for incompressible flow because the projection $\boldsymbol{\mathcal{P}}$ is a non-local operator involving the inverse Laplacian. Nevertheless, the cross terms decay fast as particles become well-separated from each other. Specifically, theoretical calculations suggest that to leading order $\boldsymbol{J}_i\boldsymbol{\mathcal{P}}\boldsymbol{S}_j$ decays with the distance between the two particles like dipole and quadrupole terms, and is thus expected to be very small in semi-dilute suspensions [73]. In many problems of practical interest there are repulsive forces between the particles that will keep them from coming close to each other, and the approximation $\boldsymbol{J}_i\boldsymbol{\mathcal{P}}\boldsymbol{S}_j \approx \left(\Delta\widetilde{V}\right)^{-1}\delta_{ij}\,\boldsymbol{I}$,

$$\boldsymbol{\rho}_{\text{eff}}^{-1} \approx \rho^{-1}\left[\boldsymbol{I} - \sum_i \frac{(m_e)_i\,\Delta\widetilde{V}_i}{\tilde{m}_i}\boldsymbol{\mathcal{P}}\boldsymbol{S}_i\boldsymbol{J}_i\boldsymbol{\mathcal{P}}\right] \tag{A7}$$

will be quite accurate even for multiparticle systems. We investigate the accuracy of the approximation $\boldsymbol{J}_i\boldsymbol{\mathcal{P}}\boldsymbol{S}_j = \boldsymbol{0}$ for $i \neq j$ numerically in Section III A 2.

B. Fluctuation-Dissipation Balance

In this Appendix we demonstrate that the coupled fluid-particle equations written in the form (18)

$$\begin{aligned}\partial_t \boldsymbol{v} &= \boldsymbol{\rho}_{\text{eff}}^{-1}\boldsymbol{\mathcal{P}}\left\{-\left[\rho\left(\boldsymbol{v}\cdot\boldsymbol{\nabla}\right)+m_e\boldsymbol{SJ}\left(\boldsymbol{v}\cdot\frac{\partial}{\partial\boldsymbol{q}}\boldsymbol{J}\right)\right]\boldsymbol{v}+\boldsymbol{SF}\right.\\ &+ \left.\eta\boldsymbol{\nabla}^2\boldsymbol{v}+\boldsymbol{\nabla}\cdot\left[\left(k_BT\eta\right)^{\frac{1}{2}}\left(\boldsymbol{\mathcal{W}}+\boldsymbol{\mathcal{W}}^T\right)\right]\right\} \qquad \text{(B1)}\\ \frac{d\boldsymbol{q}}{dt} &= \boldsymbol{J}\boldsymbol{v}, \qquad \text{(B2)}\end{aligned}$$

obey fluctuation-dissipation balance with respect to the Gibbs-Boltzmann distribution with coarse-grained free energy

$$H\left(\boldsymbol{v},\boldsymbol{q}\right)=\frac{1}{2}\int\boldsymbol{v}^T\boldsymbol{\rho}_{\text{eff}}\boldsymbol{v}\,d\boldsymbol{r}+U\left(\boldsymbol{q}\right),$$

where the effective fluid inertia operator $\boldsymbol{\rho}_{\text{eff}}$ is given in (19). The calculations here will follow the techniques described in detail in Ref. [72] (ignoring boundary terms), and are purely formal in the continuum (infinite-dimensional) setting. More precisely, the equations we write should really be interpreted as a short-hand notation for a spatially-discretized system in which there is a finite number of degrees of freedom [27].

It is well-known [69, 72] that within an isothermal and Markovian approximation the generic form of evolution equations for a set of macroscopic variables $\boldsymbol{x}$ has the Ito Langevin form

$$\partial_t\boldsymbol{x}=-\boldsymbol{N}\left(\boldsymbol{x}\right)\frac{\partial H}{\partial\boldsymbol{x}}+\left(2k_BT\right)^{\frac{1}{2}}\boldsymbol{B}\left(\boldsymbol{x}\right)\boldsymbol{W}(t)+\left(k_BT\right)\frac{\partial}{\partial\boldsymbol{x}}\cdot\boldsymbol{N}^{\star}\left(\boldsymbol{x}\right), \qquad \text{(B3)}$$

where $\boldsymbol{W}(t)$ denotes a collection of independent white noise processes, star denotes an adjoint, and $\left(\partial_{\boldsymbol{x}}\cdot\boldsymbol{N}^{\star}\right)_k=\partial N_{kj}/\partial x_j$ in indicial notation. We will suppress the explicit dependence on $\boldsymbol{x}$ and usually write the *mobility operator* as $\boldsymbol{N}\equiv\boldsymbol{N}\left(\boldsymbol{x}\right)$. The fluctuation dissipation balance is contained in the relation

$$\boldsymbol{B}\boldsymbol{B}^{\star}=\boldsymbol{M}=\frac{1}{2}\left(\boldsymbol{N}+\boldsymbol{N}^{\star}\right),$$

and the last term in (B3) is a "thermal" drift which ensures that the dynamics obeys detailed-balance (time-reversibility) at equilibrium with respect the Gibbs-Boltzmann distribution $Z^{-1}\exp\left[-H\left(\boldsymbol{x}\right)/k_BT\right]$.

In our case, the coarse-grained variables are $\boldsymbol{x}=(\boldsymbol{v},\boldsymbol{q})$ and the thermodynamic driving forces are given by the functional and partial derivatives

$$\begin{aligned}\frac{\partial H}{\partial\boldsymbol{v}} &= \boldsymbol{\rho}_{\text{eff}}\boldsymbol{v} \qquad \text{(B4)}\\ \frac{\partial H}{\partial\boldsymbol{q}} &= -\boldsymbol{F}\left(\boldsymbol{q}\right)+m_e\boldsymbol{J}\left(\boldsymbol{v}\cdot\frac{\partial}{\partial\boldsymbol{q}}\boldsymbol{J}\right)\boldsymbol{v}. \qquad \text{(B5)}\end{aligned}$$

Note that the kinetic term $m_e\boldsymbol{SJ}\left(\boldsymbol{v}\cdot\frac{\partial}{\partial\boldsymbol{q}}\boldsymbol{J}\right)\boldsymbol{v}$ appearing in (B1) term follows from $\partial H/\partial\boldsymbol{q}$, and therefore the approximation $\boldsymbol{a}_{\boldsymbol{J}}\approx\boldsymbol{0}$ in (8) is not consistent in general.

1. Mobility Operator

The mobility operator that gives (B1,B2) from the thermodynamic driving forces (B4,B5) can be symbolically written in the block form

$$\boldsymbol{N}=\begin{bmatrix}\boldsymbol{\rho}_{\text{eff}}^{-1}\boldsymbol{N}_{\text{NS}}\boldsymbol{\rho}_{\text{eff}}^{-1} & \boldsymbol{\rho}_{\text{eff}}^{-1}\boldsymbol{\mathcal{P}}\boldsymbol{S}\\ -\boldsymbol{J}\boldsymbol{\mathcal{P}}\boldsymbol{\rho}_{\text{eff}}^{-1} & \boldsymbol{N}_{\text{BD}}\end{bmatrix},\quad \boldsymbol{B}=\begin{bmatrix}\boldsymbol{\rho}_{\text{eff}}^{-1}\boldsymbol{B}_{\text{NS}}\\ \boldsymbol{B}_{\text{BD}}\end{bmatrix}$$

where the subscript NS denotes the corresponding operators for the fluctuating Navier-Stokes equations without any immersed particles [71]. In the form of the equations that we employ $\boldsymbol{N}_{\text{BD}}=\boldsymbol{B}_{\text{BD}}=\boldsymbol{0}$. Fluctuation-dissipation balance then follows from the corresponding property of the pure fluid equations, $\boldsymbol{B}_{\text{NS}}\boldsymbol{B}_{\text{NS}}^{\star}=\left(\boldsymbol{N}_{\text{NS}}+\boldsymbol{N}_{\text{NS}}^{\star}\right)/2$, since there is no dissipative terms in the particle equation. The fluid-particle coupling contribution is non-dissipative or skew-adjoint, which follows from the antisymmetry between to the lower left and upper right blocks in $\boldsymbol{N}$.

It is, however, consistent with the general framework of augmented Langevin descriptions and fluctuation-dissipation balance to allow for a nonzero $\boldsymbol{B}_{\mathrm{BD}}$ with $\boldsymbol{N}_{\mathrm{BD}} = \boldsymbol{B}_{\mathrm{BD}}\boldsymbol{B}^{\star}_{\mathrm{BD}}$. This does not affect the fluid equation but changes the particle equation. For example, the simple choice

$$\boldsymbol{B}_{\mathrm{BD}} = \sqrt{\zeta}\,\boldsymbol{I}, \quad \boldsymbol{N}_{\mathrm{BD}} = \zeta\,\boldsymbol{I},$$

leads to a modified equation of motion for the particle,

$$\frac{d\boldsymbol{q}}{dt} = \boldsymbol{J}\left[\boldsymbol{I} - \zeta m_e\left(\boldsymbol{v}\cdot\frac{\partial}{\partial\boldsymbol{q}}\boldsymbol{J}\right)\right]\boldsymbol{v} + \zeta\boldsymbol{F}\left(\boldsymbol{q}\right) + \sqrt{2\zeta k_B T}\,\boldsymbol{\mathcal{W}}_{\boldsymbol{q}},$$

where $\boldsymbol{\mathcal{W}}_{\boldsymbol{q}}\left(t\right)$ is a white-noise "random slip". If $\boldsymbol{v} = \boldsymbol{0}$ the above equation is recognized as the usual equation of Brownian dynamics with friction coefficient ζ^{-1} and mobility ζ. It is therefore expected that adding a random slip component would increase the diffusion coefficient of the particle by $\zeta k_B T$. This can be used to tune the diffusion coefficient to some target (e.g., experimental) value. In this work, we fix $\zeta = 0$, and in future work we will consider second-order algorithms for $\zeta > 0$ for the case of a neutrally buoyant particle, $m_e = 0$.

2. Thermal Drift

Complications arise because of the presence of the last term in (B3) (proportional to $k_B T$), which is non-zero because $\boldsymbol{J}$ and $\boldsymbol{S}$ depend on $\boldsymbol{q}$. This "thermal drift" term arises because $\boldsymbol{u}$ is eliminated from the description as a variable adiabatically-slaved to $\boldsymbol{v}$, neglecting the fact that the particle position fluctuates rapidly due to the fluctuations in $\boldsymbol{u}$. This term was identified in Ref. [27] as missing from the formulation of the Stochastic Immersed Boundary Method [28], which is equivalent to our formulation with $m_e = 0$. The missing contribution is an extra term in the velocity equation (B1) of the form

$$\boldsymbol{f}_{\mathrm{th}} = \left(k_B T\right)\frac{\partial}{\partial\boldsymbol{q}}\cdot\left(\boldsymbol{J}\boldsymbol{\mathcal{P}}\boldsymbol{\rho}_{\mathrm{eff}}^{-1}\right).$$

One can, in principle, numerically evaluate this term without requiring any derivatives by using "random finite-differences" [100] via the identity

$$\frac{\partial}{\partial\boldsymbol{q}}\cdot\left(\boldsymbol{J}\boldsymbol{\mathcal{P}}\boldsymbol{\rho}_{\mathrm{eff}}^{-1}\right) = \lim_{\epsilon\to 0}\epsilon^{-2}\left\langle\left[\boldsymbol{\rho}_{\mathrm{eff}}^{-1}\left(\boldsymbol{q}+\Delta\boldsymbol{q}\right)\boldsymbol{\mathcal{P}}\boldsymbol{S}\left(\boldsymbol{q}+\Delta\boldsymbol{q}\right)\right]\Delta\boldsymbol{q} - \left[\boldsymbol{\rho}_{\mathrm{eff}}^{-1}\left(\boldsymbol{q}\right)\boldsymbol{\mathcal{P}}\boldsymbol{S}\left(\boldsymbol{q}\right)\right]\Delta\boldsymbol{q}\right\rangle_{\Delta\boldsymbol{q}}.$$

where the expectation value is with respect to the small random particle displacement $\Delta\boldsymbol{q} = \epsilon\boldsymbol{W}_{\boldsymbol{q}}$, and the components of $\boldsymbol{W}_{\boldsymbol{q}}$ are independent standard normal variates.

For periodic boundaries, using (A5) we can simplify, for any $\boldsymbol{q}$,

$$\left[\boldsymbol{\rho}_{\mathrm{eff}}^{-1}\left(\boldsymbol{q}\right)\boldsymbol{\mathcal{P}}\boldsymbol{S}\left(\boldsymbol{q}\right)\right]\Delta\boldsymbol{q} = \frac{\tilde{m}_f}{\tilde{m}}\boldsymbol{\mathcal{P}}\left[\boldsymbol{S}\left(\boldsymbol{q}\right)\right]\Delta\boldsymbol{q},$$

and therefore the thermal drift

$$\boldsymbol{f}_{\mathrm{th}} = \left(k_B T\right)\frac{\tilde{m}_f}{\tilde{m}}\boldsymbol{\mathcal{P}}\lim_{\epsilon\to 0}\epsilon^{-2}\left\langle\left[\boldsymbol{S}\left(\boldsymbol{q}+\Delta\boldsymbol{q}\right) - \boldsymbol{S}\left(\boldsymbol{q}\right)\right]\Delta\boldsymbol{q}\right\rangle_{\Delta\boldsymbol{q}} = -\frac{\tilde{m}_f}{\tilde{m}}\boldsymbol{\mathcal{P}}\boldsymbol{\nabla}\boldsymbol{S}\left(k_B T\right) = \boldsymbol{0} \tag{B6}$$

vanishes as the projection of a gradient of a scalar for incompressible flow. Observe that we can interpret $\boldsymbol{f}_{\mathrm{th}}$ as the divergence of an additional "thermal" contribution to the kinetic stress,

$$\boldsymbol{\sigma}_{\mathrm{kin}} + \boldsymbol{\sigma}_{\mathrm{th}} = -\rho\boldsymbol{v}\boldsymbol{v}^T - \boldsymbol{S}\left(m_e\boldsymbol{u}\boldsymbol{u}^T + \frac{\tilde{m}_f}{\tilde{m}}k_B T\right)\boldsymbol{I}. \tag{B7}$$

In the limit of large Schmidt number, the particle positions evolves much slower than the momenta (and thus also $\boldsymbol{u}$), and therefore the temporal average of the second term gives an average "themal" contribution to the pressure

$$\pi_{\mathrm{th}} = \boldsymbol{S}\left(\frac{m_e}{d}\left\langle\boldsymbol{u}^T\boldsymbol{u}\right\rangle + \frac{\tilde{m}_f}{\tilde{m}}k_B T\right) = \left(\frac{m_e}{\tilde{m}} + \frac{\tilde{m}_f}{\tilde{m}}\right)\boldsymbol{S}\left(k_B T\right) = \boldsymbol{S}\left(k_B T\right),$$

where we used the equipartition result (24). We see that π_{th} has the physical interpretation as an osmotic pressure contribution arising from the thermal motion of the suspended particle, even when excess particle inertia is present.

Note that some of the mathematical manipulations used above rely on continuum identities which are not necessarily obeyed by the discrete operators. In particular, for the spatial discretization that we employ, the discrete vector field $\boldsymbol{w} = \frac{\partial}{\partial\boldsymbol{q}}\boldsymbol{S}$ is not equal to a discrete gradient of a scalar. However, numerical observations suggest that it is very close to a discrete gradient of a scalar field, in the sense that application of the discrete projection reduces the magnitude by several orders, $\|\boldsymbol{\mathcal{P}}\boldsymbol{w}\| \ll \|\boldsymbol{w}\|$, regardless of the position of the particle relative to the fluid grid. This is one more case in which we find that the Peskin discrete local averaging and local spreading operators closely mimic the properties

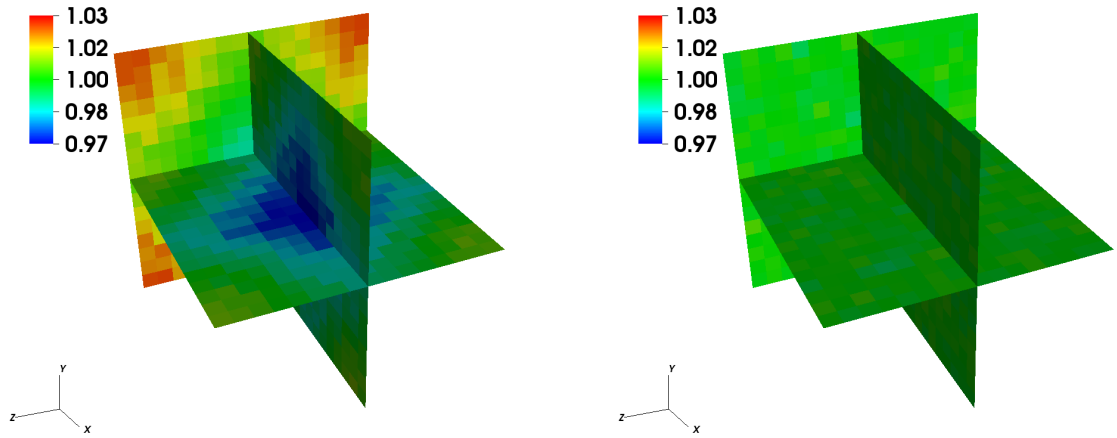

Figure 7: Histogram of the equilibrium distribution of the position a particle freely diffusing at thermodynamic equilibrium in a translationally-invariant system, without including the drift $\boldsymbol{f}_{\text{th}}$ (left) and including it (right). The distribution is projected on one grid cell by averaging over all cells in a periodic grid of 16^3 cells, and averaging over 100 non-interacting neutrally-buyuoant particles ($m_e = 0$). A small time step size corresponding to viscous CFL number $\beta = 0.05$ is used in order to minimize time discretization artifacts.

of the continuum operators even though they were not specifically designed with the staggered-grid discrete projection $\boldsymbol{\mathcal{P}}$ in mind.

Nevertheless, maintaining strict discrete fluctuation-dissipation balance in the semi-discrete equations requires keeping $\boldsymbol{f}_{\text{th}}$ in the momentum balance equation. Note that the gradient in $\boldsymbol{\nabla S}$ is a *continuum* rather than a discrete gradient. There are two ways to implement this gradient numerically. The first is to differentiate the Peskin kernel and spread a force $-(k_B T)\,\tilde{m}_f/\tilde{m}$ using the derivative of the kernel instead of the kernel itself, as we do in the calculations reported below. The second is to use a random finite difference [100] and numerically obtain the required derivative *in expectation* by adding a term

$$\frac{\tilde{m}_f}{\tilde{m}}\,\frac{k_B T}{\epsilon}\left[\boldsymbol{S}\left(\boldsymbol{q}+\epsilon\boldsymbol{W}_{\boldsymbol{q}}\right)-\boldsymbol{S}\left(\boldsymbol{q}\right)\right]\boldsymbol{W}_{\boldsymbol{q}} \quad \text{or} \quad -(k_B T)\,\frac{\tilde{m}_f}{\tilde{m}}\boldsymbol{\partial}\delta_a\left(\boldsymbol{q}-\boldsymbol{r}\right) \tag{B8}$$

to the momentum balance equation, where $\boldsymbol{W}_{\boldsymbol{q}}$ is generated independently at each time step, and $\boldsymbol{\partial}\delta_a$ is the gradient of the kernel.

To demonstrate that it is necessary to include (B8) in the discrete setting, we investigate the equilibrium distribution of the positions of a number of non-interacting particles (i.e., an ideal gas of particles) in a periodic domain. Due to the translational invariance of the problem we know that the particle positions should be uniformly distributed through space. However, the presence of the fluid grid breaks translational invariance and if discrete fluctuation-dissipation balance is not strictly obeyed grid artifacts can appear in the solution. In Fig. 7 we show a histogram of the positions of a freely-diffusing particle relative to one grid cell (by grid translational invariance all cells and particles are equivalent). We clearly see that when $\boldsymbol{f}_{\text{th}}$ is omitted the particle spends more time near the corners of the grid cell (nodes of the grid) than near the center of the grid cell, by about 6% for these parameters, Including the thermal drift forcing in the momentum equation eliminates these artifacts to within statistical and temporal integration errors. In this test we use neutrally-buyuoant particles because the particles do not interact via a steric repulsion so the approximation (A7) fails and the no-slip condition is not enforced to sufficient accuracy by the algorithm developed here. We note that similar results are obtained for the compressible fluid equations, solved using the algorithm described in Appendix D.

C. Semi-Implicit Inertial Velocity Correction

In this Appendix we explain how we solve the linear system (37,38,39,40) using only FFT-based linear subsolvers.

First, we solve (37) for $\tilde{\boldsymbol{v}}^{n+1}$ and $\tilde{\pi}^{n+\frac{1}{2}}$ using a projection algorithm and FFTs as explained, for

example, in Ref. [81]. Then, we eliminate $\boldsymbol{u}^{n+1}$ in (39) using the no-slip constraint (40) to obtain

$$\Delta t\,\left(\boldsymbol{\lambda}^{n+\frac{1}{2}}+\boldsymbol{F}^{n+\frac{1}{2}}\right)=m_e\left[\boldsymbol{J}\left(\tilde{\boldsymbol{v}}^{n+1}+\Delta\boldsymbol{v}^{n+\frac{1}{2}}\right)+\Delta\boldsymbol{u}^{n+\frac{1}{2}}-\boldsymbol{u}^{n}\right]. \tag{C1}$$

Substituting this expression into (38), we get the fluid correction equation

$$\left[\left(\rho\boldsymbol{I}-\frac{\Delta t}{2}\eta\boldsymbol{L}\right)+m_e\boldsymbol{S}\boldsymbol{J}\right]\Delta\boldsymbol{v}^{n+\frac{1}{2}}+\Delta t\,\boldsymbol{\nabla}\left(\Delta\pi^{n+\frac{1}{2}}\right)=m_e\boldsymbol{S}\left(\boldsymbol{u}^{n}-\boldsymbol{J}\tilde{\boldsymbol{v}}^{n+1}-\Delta\boldsymbol{u}^{n+\frac{1}{2}}\right), \tag{C2}$$

where $\boldsymbol{L}\equiv\boldsymbol{\nabla}^2$ denotes the discrete Laplacian operator. This is nothing but a rewriting of (36) and just as difficult to solve, so we need to make further approximations.

We begin solving (C2) by writing

$$\rho\boldsymbol{I}-\frac{\Delta t}{2}\eta\boldsymbol{L}+m_e\boldsymbol{S}\boldsymbol{J}=\left(\rho\boldsymbol{I}+m_e\boldsymbol{S}\boldsymbol{J}\right)\left(\boldsymbol{I}-\frac{\Delta t}{2}\nu\boldsymbol{L}\right)+\frac{m_e\nu\Delta t}{2}\boldsymbol{S}\boldsymbol{J}\boldsymbol{L}. \tag{C3}$$

If we ignore the last term we obtain an approximation of order $O\left(m_e\Delta t\right)$. We can improve the accuracy for viscous-dominated flows if we time lag the last term, that is, evaluate it during the previous time step. Using (C3) in (C2) along with time lagging we get the modified fluid correction equation

$$\left(\rho\boldsymbol{I}+m_e\boldsymbol{S}\boldsymbol{J}\right)\left(\boldsymbol{I}-\frac{\Delta t}{2}\nu\boldsymbol{L}\right)\Delta\boldsymbol{v}^{n+\frac{1}{2}}+\Delta t\,\boldsymbol{\nabla}\left(\Delta\pi^{n+\frac{1}{2}}\right)=\boldsymbol{S}\left[m_e\left(\boldsymbol{u}^{n}-\boldsymbol{J}\tilde{\boldsymbol{v}}^{n+1}-\delta\boldsymbol{u}^{n+\frac{1}{2}}\right)\right]=\boldsymbol{S}\Delta\boldsymbol{p}, \tag{C4}$$

where we have time-lagged the last term in (C3) and denoted

$$\delta\boldsymbol{u}^{n+\frac{1}{2}}=\Delta\boldsymbol{u}^{n+\frac{1}{2}}+\frac{\nu\Delta t}{2}\boldsymbol{J}^{n-\frac{1}{2}}\boldsymbol{L}\Delta\boldsymbol{v}^{n-\frac{1}{2}}.$$

Note that if one approximates $\delta\boldsymbol{u}^{n+\frac{1}{2}}=\boldsymbol{0}$ (lowering the order of accuracy), then as $\Delta t\to\infty$ we get $\Delta\boldsymbol{v}^{n+\frac{1}{2}}\to\boldsymbol{0}$ since both $\rho\boldsymbol{I}+m_e\boldsymbol{S}\boldsymbol{J}$ and $\boldsymbol{I}-\Delta t\nu\boldsymbol{L}/2$ are positive semi-definite matrices. This is consistent with physical intuition that inertia should play a negligible role at long time scales, and implies that the algorithm will be unconditionally stable in the absence of advection.

Denote $\Delta\tilde{\boldsymbol{v}}=\left(\boldsymbol{I}-\Delta t\nu\boldsymbol{L}/2\right)\Delta\boldsymbol{v}^{n+\frac{1}{2}}$, and note that $\boldsymbol{\nabla}\cdot\left(\Delta\tilde{\boldsymbol{v}}\right)=\boldsymbol{\nabla}\cdot\left(\Delta\boldsymbol{v}^{n+\frac{1}{2}}\right)=0$ with periodic BCs since the divergence and the Laplacian commute. Therefore, we can solve (C4) by first solving

$$\left(\rho\boldsymbol{I}+m_e\boldsymbol{S}\boldsymbol{J}\right)\Delta\tilde{\boldsymbol{v}}+\Delta t\,\boldsymbol{\nabla}\left(\Delta\pi^{n+\frac{1}{2}}\right)=\boldsymbol{S}\Delta\boldsymbol{p}, \tag{C5}$$

subject to $\boldsymbol{\nabla}\cdot\left(\Delta\tilde{\boldsymbol{v}}\right)=0$, and then solving

$$\left(\boldsymbol{I}-\frac{\Delta t}{2}\nu\boldsymbol{L}\right)\Delta\boldsymbol{v}^{n+\frac{1}{2}}=\Delta\tilde{\boldsymbol{v}}, \tag{C6}$$

using FFTs to diagonalize the Laplacian [81].

We solve (C5) using the pressure-elimination procedure described in Appendix A, to obtain the equivalent pressure-free equation

$$\left(\rho\boldsymbol{I}+m_e\boldsymbol{\mathcal{P}}\boldsymbol{S}\boldsymbol{J}\boldsymbol{\mathcal{P}}\right)\Delta\tilde{\boldsymbol{v}}=\boldsymbol{\mathcal{P}}\boldsymbol{S}\Delta\boldsymbol{p},$$

where $\boldsymbol{\mathcal{P}}$ denotes the discrete projection operator [44]. By assuming the particles are sufficiently far away from each other, we can employ the approximation (28) to approximately solve (C5) [compare to (A6) with $\boldsymbol{f}=\boldsymbol{0}$],

$$\Delta\tilde{\boldsymbol{v}}=\frac{\tilde{m}_f}{\rho\tilde{m}}\boldsymbol{\mathcal{P}}\boldsymbol{S}\Delta\boldsymbol{p}=\frac{m_e\tilde{m}_f}{\rho\left(\tilde{m}_f+m_e\right)}\boldsymbol{\mathcal{P}}\boldsymbol{S}\left(\boldsymbol{u}^{n}-\boldsymbol{J}\tilde{\boldsymbol{v}}^{n+1}-\delta\boldsymbol{u}^{n+\frac{1}{2}}\right), \tag{C7}$$

where we recall that $\tilde{m}_f=d\,m_f/\left(d-1\right)$.

In order to ensure strict momentum conservation even in the presence of the approximation (28), we rewrite (C6) using (C5) in the form

$$\left(\rho\boldsymbol{I}-\frac{\Delta t}{2}\eta\boldsymbol{L}\right)\Delta\boldsymbol{v}^{n+\frac{1}{2}}+\Delta t\,\boldsymbol{\nabla}\left(\Delta\pi^{n+\frac{1}{2}}\right)=\boldsymbol{S}\left(\Delta\boldsymbol{p}-m_e\boldsymbol{J}\Delta\tilde{\boldsymbol{v}}\right), \tag{C8}$$

subject to $\boldsymbol{\nabla}\cdot\left(\Delta\boldsymbol{v}^{n+\frac{1}{2}}\right)=0$. This is consistent with (38) if we set

$$\Delta t\,\left(\boldsymbol{\lambda}^{n+\frac{1}{2}}+\boldsymbol{F}^{n+\frac{1}{2}}\right)=m_e\left[\boldsymbol{J}\left(\tilde{\boldsymbol{v}}^{n+1}+\Delta\tilde{\boldsymbol{v}}\right)+\delta\boldsymbol{u}^{n+\frac{1}{2}}-\boldsymbol{u}^{n}\right], \tag{C9}$$

which is to be compared to (C1). The linear system (C8) can be solved using the same techniques as used to solve (37).

Having determined $\Delta\boldsymbol{v}^{n+\frac{1}{2}}$ we can update the fluid velocity (momentum), $\boldsymbol{v}^{n+1}=\tilde{\boldsymbol{v}}^{n+1}+\Delta\boldsymbol{v}^{n+\frac{1}{2}}$. If desired, we can update the particle momentum in a conservative manner by substituting (C9)

into (39) to get

$$\boldsymbol{u}^{n+1} = \boldsymbol{J}\left(\tilde{\boldsymbol{v}}^{n+1} + \Delta\tilde{\boldsymbol{v}}\right) + \delta\boldsymbol{u}^{n+\frac{1}{2}}. \tag{C10}$$

From (C10) we see that the approximations we used lead to a violation of the no-slip constraint (40)

$$\Delta\boldsymbol{u}^{n+1} = \boldsymbol{u}^{n+1} - \boldsymbol{J}\left(\tilde{\boldsymbol{v}}^{n+1} + \Delta\boldsymbol{v}^{n+\frac{1}{2}}\right) - \Delta\boldsymbol{u}^{n+\frac{1}{2}} \approx -\frac{\nu\Delta t}{2}\left(\boldsymbol{J}^{n+\frac{1}{2}}\boldsymbol{L}\Delta\boldsymbol{v}^{n+\frac{1}{2}} - \boldsymbol{J}^{n-\frac{1}{2}}\boldsymbol{L}\Delta\boldsymbol{v}^{n-\frac{1}{2}}\right),$$

which is of $O\left(m_e\Delta t^2\right)$ if $\Delta\boldsymbol{v}$ is smooth in time, $\Delta\boldsymbol{v}^{n+\frac{1}{2}} = \Delta\boldsymbol{v}^{n-\frac{1}{2}} + O\left(\Delta t\right)$. This means that our procedure approximates the solution of (37,38,39,40) with slip $\Delta\boldsymbol{u}^{n+1} = O\left(m_e\Delta t^2\right)$. Therefore, in the deterministic setting, for sufficiently smooth flows, the temporal integrator summarized in Section III C is expected to be second-order accurate.

D. Compressible Algorithm

In this Appendix we propose a modification of the first-order temporal integrator developed in Ref. [43], following a similar approach to the incompressible algorithm summarized in Section III C. The algorithm summarized below is still only first-order accurate if $m_e \neq 0$, however, it is second-order accurate for $m_e = 0$ unlike the algorithm used in Ref. [43]. Numerical results show that the modified algorithm below is a substantial improvement over that used previously [43].

1. Estimate the position of the particle at the midpoint,
$$\boldsymbol{q}^{n+\frac{1}{2}} = \boldsymbol{q}^n + \frac{\Delta t}{2}\boldsymbol{J}^n\boldsymbol{v}^n, \tag{D1}$$
and evaluate the external or interparticle forces $\boldsymbol{F}^{n+\frac{1}{2}}\left(\boldsymbol{q}^{n+\frac{1}{2}}\right)$.
2. Solve the coupled density and unperturbed momentum equations [101]
$$\begin{aligned} D_t\rho &= -\rho\left(\boldsymbol{\nabla}\cdot\boldsymbol{v}\right) \\ \rho\left(D_t\boldsymbol{v}\right) &= -c_T^2\boldsymbol{\nabla}\rho - \boldsymbol{\nabla}\cdot\boldsymbol{\sigma} + \boldsymbol{S}^{n+\frac{1}{2}}\boldsymbol{F}^{n+\frac{1}{2}} - \left(k_BT\right)\sum_i \frac{\left(\tilde{m}_f\right)_i}{\tilde{m}_i}\boldsymbol{\partial}\delta_a\left(\boldsymbol{q}_i^{n+\frac{1}{2}} - \boldsymbol{r}\right) \end{aligned}$$
using a step of the third-order Runge-Kutta algorithm described in Ref. [44], to obtain the density ρ^{n+1} and the unperturbed fluid velocity $\tilde{\boldsymbol{v}}^{n+1}$. During this step we treat the force density $\boldsymbol{S}^{n+\frac{1}{2}}\boldsymbol{F}^{n+\frac{1}{2}}$ in the momentum equation as a constant external forcing. Here we have included the thermal drift term (B6) as the last term on the right hand side of the momentum equation, denoting the gradient of the kernel with $\boldsymbol{\partial}\delta_a$; omitting this term violates discrete fluctuation-dissipation balance and leads to small but measurable unphysical grid artifacts in the particle dynamics.
3. If $m_e = 0$, set $\boldsymbol{v}^{n+1} = \tilde{\boldsymbol{v}}^{n+1}$ and skip to step 7.
4. Calculate the momentum exchange between the particle and the fluid during the time step [43],
$$\Delta t\left(\boldsymbol{\lambda}^{n+\frac{1}{2}} + \boldsymbol{F}^{n+\frac{1}{2}}\right) = \Delta\boldsymbol{p} = \frac{m_em_f}{m_e+m_f}\left(\boldsymbol{J}^{n+\frac{1}{2}}\tilde{\boldsymbol{v}}^{n+1} - \boldsymbol{u}^n\right), \tag{D2}$$
where the mass of the dragged fluid is estimated from the local density as
$$m_f = \Delta V\,\boldsymbol{J}^{n+\frac{1}{2}}\rho^{n+1}.$$
5. Update the particle momentum,
$$\boldsymbol{u}^{n+1} = \boldsymbol{u}^n + \frac{\Delta\boldsymbol{p}}{m_e} = \boldsymbol{u}^n + \frac{m_f}{m_e+m_f}\left(\boldsymbol{J}^{n+\frac{1}{2}}\tilde{\boldsymbol{v}}^{n+1} - \boldsymbol{u}^n\right). \tag{D3}$$
6. Update the fluid velocity to enforce the no-slip condition $\boldsymbol{u}^{n+1} = \boldsymbol{J}^{n+\frac{1}{2}}\boldsymbol{v}^{n+1}$,
$$\boldsymbol{v}^{n+1} = \tilde{\boldsymbol{v}}^{n+1} - \frac{\Delta V}{m_f}\boldsymbol{S}^{n+\frac{1}{2}}\Delta\boldsymbol{p} = \tilde{\boldsymbol{v}}^{n+1} + \Delta V\,\boldsymbol{S}^{n+\frac{1}{2}}\left(\boldsymbol{u}^{n+1} - \boldsymbol{J}^{n+\frac{1}{2}}\tilde{\boldsymbol{v}}^{n+1}\right) \tag{D4}$$
Note that this conserves the total momentum since
$$\rho^{n+1}\boldsymbol{v}^{n+1} - \rho^{n+1}\tilde{\boldsymbol{v}}^{n+1} = -\left(\boldsymbol{J}^{n+\frac{1}{2}}\rho^{n+1}\right)^{-1}\int\rho^{n+1}\boldsymbol{S}^{n+\frac{1}{2}}\Delta\boldsymbol{p}\,d\boldsymbol{r} = -\Delta\boldsymbol{p}.$$
7. Update the particle position,
$$\boldsymbol{q}^{n+1} = \boldsymbol{q}^n + \frac{\Delta t}{2}\boldsymbol{J}^{n+\frac{1}{2}}\left(\boldsymbol{v}^{n+1} + \boldsymbol{v}^n\right). \tag{D5}$$

There are some differences between this algorithm and the incompressible algorithm summarized in Section III C. Notably, viscosity is handled explicitly in the compressible algorithm to avoid costly linear solvers. In the incompressible algorithm the no-slip condition is violated slightly while here the velocity

correction $\boldsymbol{v}^{n+1} - \tilde{\boldsymbol{v}}^{n+1}$ is designed to enforce the no-slip condition exactly even in the presence of density variations. Another important difference is that the compressible algorithm presented here is not second-order accurate even for small Reynolds number because viscous corrections of $O(\nu\Delta t)$ are not taken into account when computing $\boldsymbol{\lambda}^{n+\frac{1}{2}}$. Since the time step is typically strongly limited by propagation of sound, typically $\beta \ll 1$ and there is little need for higher-order handling of the viscous terms.

---

[1] S. Kim and S.J. Karrila. *Microhydrodynamics:Principles and Selected Applications*. Butterworth Heinemann, Boston, 1991.

[2] R. Kapral. Multiparticle collision dynamics: simulation of complex systems on mesoscales. *Adv. Chem. Phys*, 140:89, 2008.

[3] B. Dünweg and A.J.C. Ladd. Lattice Boltzmann simulations of soft matter systems. *Adv. Comp. Sim. for Soft Matter Sciences III*, pages 89–166, 2009.

[4] C.W. Hsu and Y.L. Chen. Migration and fractionation of deformable particles in microchannel. *The Journal of chemical physics*, 133:034906, 2010.

[5] J. K. Eaton. Two-way coupled turbulence simulations of gas-particle flows using point-particle tracking. *International Journal of Multiphase Flow*, 35(9):792–800, 2009.

[6] T. Tanaka and J. K. Eaton. Sub-Kolmogorov resolution partical image velocimetry measurements of particle-laden forced turbulence. *Journal of Fluid Mechanics*, 643:177, 2010.

[7] M. Rhodes. *Introduction to particle technology*. John Wiley and Sons, 2008.

[8] X. Bian, S. Litvinov, R. Qian, M. Ellero, and N. A. Adams. Multiscale modeling of particle in suspension with smoothed dissipative particle dynamics. *Physics of Fluids*, 24(1):012002, 2012.

[9] A. Vázquez-Quesada, M. Ellero, and P. Español. Consistent scaling of thermal fluctuations in smoothed dissipative particle dynamics. *J. Chem. Phys.*, 130:034901, 2009.

[10] J. Horbach and S. Succi. Lattice Boltzmann versus Molecular Dynamics Simulation of Nanoscale Hydrodynamic Flows. *Physical Review Letters*, 96(22):1–4, 2006.

[11] Simone Melchionna. Incorporation of smooth spherical bodies in the Lattice-Boltzmann method. *J. Comp. Phys*, 230(10):3966, 2011.

[12] H. H. Hu, Patankar N. A., and M. Y. Zhu. Direct numerical simulation of fluid-solid systems using the arbitrary Lagrangian-Eulerian technique. *J. Comp. Phys.*, 169:427, 2001.

[13] R. Yamamoto, K. Kim, and Y. Nakayama. Strict simulations of non-equilibrium dynamics of colloids. *Colloids and Surfaces A*, 311:42, 2007.

[14] X. Luo, M. R. Maxey, and G. Karniadakis. Smoothed profile method for particulate flows: Error analysis and simulations. *Journal of Computational Physics*, 228:1750, 2009.

[15] M. R. Maxey and J. J. Riley. Equation of motion for a small rigid sphere in a nonuniform flow. *Physics of Fluids*, 26:883, 1983.

[16] S. Lomholt. *Numerical Investigations of Macroscopic Particle Dynamics in Microflows*. PhD thesis, RisøNational Laboratory, Roskilde, Denmark, 2001.

[17] R. Monchaux, M. Bourgoin, and A. Cartellier. Analyzing preferential concentration and clustering of inertial particles in turbulence. *International Journal of Multiphase Flow*, 40:1–18, 2012.

[18] E. Falck, J. M. Lahtinen, I. Vattulainen, and T. Ala-Nissila. Influence of hydrodynamics on many-particle diffusion in 2D colloidal suspensions. *The European physical journal. E, Soft matter*, 13(3):267–75, 2004.

[19] M. Ripoll, K. Mussawisade, R. Winkler, and G. Gompper. Dynamic regimes of fluids simulated by multiparticle-collision dynamics. *Physical Review E*, 72(1):016701, July 2005.

[20] M. Balvin, E. Sohn, T. Iracki, G. Drazer, and J. Frechette. Directional locking and the role of irreversible interactions in deterministic hydrodynamics separations in microfluidic devices. *Phys. Rev. Lett.*, 103:078301, 2009.

[21] D. L. Ermak and J. A. McCammon. Brownian dynamics with hydrodynamic interactions. *The Journal of Chemical Physics*, 69(4):1352, 1978.

[22] R. M. Jendrejack, J. J. de Pablo, and M. D. Graham. Stochastic simulations of DNA in flow: Dynamics and the effects of hydrodynamic interactions. *J. Chem. Phys.*, 116(17):7752–7759, 2002.

[23] Adolfo J. Banchio and John F. Brady. Accelerated Stokesian dynamics: Brownian motion. *The Journal of Chemical Physics*, 118(22):10323, 2003.

[24] A. Sierou and J. F. Brady. Accelerated Stokesian Dynamics simulations. *J. Fluid Mech.*, 448:115–146, 2001.

[25] J. P. Hernandez-Ortiz, J. J. de Pablo, and M. D. Graham. Fast Computation of Many-Particle Hydrodynamic and Electrostatic Interactions in a Confined Geometry. *Phys. Rev. Lett.*, 98(14):140602, 2007.

[26] 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(15):154903, 2007.

[27] P. J. Atzberger. Stochastic Eulerian-Lagrangian Methods for Fluid-Structure Interactions with Thermal Fluctuations. *J. Comp. Phys.*, 230:2821–2837, 2011.

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

[29] S. Lomholt and M.R. Maxey. Force-coupling method for particulate two-phase flow: Stokes flow. *J. Comp. Phys.*, 184(2):381–405, 2003.

[30] M. R. Maxey and B. K. Patel. Localized force representations for particles sedimenting in Stokes flow. *International journal of multiphase flow*, 27(9):1603–1626, 2001.

[31] Eric E. Keaveny. Fluctuating force-coupling method for simulations of colloidal suspensions. *J. Comp. Phys.*, 269(0):61 – 79, 2014.

[32] C.S. Peskin. The immersed boundary method. *Acta Numerica*, 11:479–517, 2002.

[33] P. M. Lovalenti and J. F. Brady. The force on a bubble, drop, or particle in arbitrary time-dependent motion at small Reynolds number. *Physics of Fluids*, 5(5):2104–2116, 1993.

[34] L. Botto and A. Prosperetti. A fully resolved numerical simulation of turbulent flow past one or several spherical particles. *Physics of Fluids*, 24(1):013303, 2012.

[35] M. Uhlmann. Interface-resolved direct numerical simulation of vertical particulate channel flow in the turbulent regime. *Physics of Fluids*, 20(5):053305, 2008.

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

[37] Y. Chen, N. Sharma, and N. Patankar. Fluctuating Immersed Material (FIMAT) dynamics for the direct simulation of the Brownian motion of particles. In *IUTAM Symposium on Computational Approaches to Multiphase Flow*, pages 119–129. Springer, 2006.

[38] Adrian M Kopacz, Neelesh A Patankar, and Wing K Liu. The immersed molecular finite element method. *Computer Methods in Applied Mechanics and Engineering*, 233:28–39, 2012.

[39] Zhi-Gang Feng and Efstathios E Michaelides. < i> proteus</i>: a direct forcing method in the simulations of particulate flows. *Journal of Computational Physics*, 202(1):20–51, 2005.

[40] F. Lucci, A. Ferrante, and S. Elghobashi. Modulation of isotropic turbulence by particles of Taylor length-scale size. *Journal of Fluid Mechanics*, 650:5, 2010.

[41] Zhi-Gang Feng and Efstathios E Michaelides. Inclusion of heat transfer computations for particle laden flows. *Physics of Fluids*, 20:040604, 2008.

[42] A. Dejoan. DNS experiments on the settling of heavy particles in homogeneous turbulence: two-way coupling and Reynolds number effects. *Journal of Physics: Conference Series*, 333:012006, December 2011.

[43] F. Balboa Usabiaga, I. Pagonabarraga, and R. Delgado-Buscalioni. Inertial coupling for point particle fluctuating hydrodynamics. *J. Comp. Phys.*, 235:701–722, 2013.

[44] F. Balboa Usabiaga, J. B. Bell, R. Delgado-Buscalioni, A. Donev, T. G. Fai, B. E. Griffith, and C. S. Peskin. Staggered Schemes for Fluctuating Hydrodynamics. *SIAM J. Multiscale Modeling and Simulation*, 10(4):1369–1408, 2012.

[45] J. Wang and J. Dual. Theoretical and numerical calculations for the time-averaged acoustic force and torque acting on a rigid cylinder of arbitrary size in a low viscosity fluid. *The Journal of the Acoustical Society of America*, 129(6):3490–501, 2011.

[46] G. Tabak and P.J. Atzberger. Systematic stochastic reduction of inertial fluid-structure interactions subject to thermal fluctuations. *arXiv preprint arXiv:1211.3798*, 2015.

[47] M. Uhlmann. An immersed boundary method with direct forcing for the simulation of particulate flows. *J. Comp. Phys.*, 209(2):448–476, 2005.

[48] J. B. Bell, P. Colella, and H. M. Glaz. A second order projection method for the incompressible Navier-Stokes equations. *J. Comp. Phys.*, 85(2):257–283, 1989.

[49] A. Pal Singh Bhalla, B. E. Griffith, N. A. Patankar, and A. Donev. A Minimally-Resolved Immersed Boundary Model for Reaction-Diffusion Problems. *J. Chem. Phys.*, 139(21):214112, 2013.

[50] F. Balboa Usabiaga and R. Delgado-Buscalioni. A minimal model for acoustic forces on Brownian particles. *Phys. Rev. E*, 88:063304, 2013.

[51] J. M. O. De Zarate and J. V. Sengers. *Hydrodynamic fluctuations in fluids and fluid mixtures*. Elsevier Science Ltd, 2006.

[52] S. Delong, B. E. Griffith, E. Vanden-Eijnden, and A. Donev. Temporal Integrators for Fluctuating Hydrodynamics. *Phys. Rev. E*, 87(3):033302, 2013.

[53] P. Español, J.G. Anero, and I. Zúñiga. Microscopic derivation of discrete hydrodynamics. *J. Chem. Phys.*, 131:244117, 2009.

[54] C. Foias, D. D Holm, and E. S Titi. The Navier–Stokes-alpha model of fluid turbulence. *Physica D: Nonlinear Phenomena*, 152:505–519, 2001.

[55] M. van Reeuwijk, H. J. J. Jonker, and K. HanjalicÌ. Incompressibility of the Leray-α model for wall-bounded flows. *Physics of Fluids*, 18(1):018103, January 2006.

[56] A. Donev, T. G. Fai, and E. Vanden-Eijnden. Reversible Diffusion by Thermal Fluctuations. Arxiv

preprint 1306.3158, 2013.
[57] A. Donev, A. J. Nonaka, Y. Sun, T. G. Fai, A. L. Garcia, and J. B. Bell. Low Mach Number Fluctuating Hydrodynamics of Diffusively Mixing Fluids. *Communications in Applied Mathematics and Computational Science*, 9(1):47–105, 2014.
[58] G. Da Prato. *Kolmogorov equations for stochastic PDEs*. Birkhauser, 2004.
[59] A.J.C. Ladd. Short-time motion of colloidal particles: Numerical simulation via a fluctuating lattice-Boltzmann equation. *Phys. Rev. Lett.*, 70(9):1339–1342, Mar 1993.
[60] D. Nie and J. Lin. A fluctuating lattice-Boltzmann model for direct numerical simulation of particle Brownian motion. *Particuology*, 7(6):501–506, 2009.
[61] J.M. Rallison. Note on the faxen relations for a particle in stokes flow. *J. Fluid Mech*, 88(part 3):529–533, 1978.
[62] Y. Mori and C.S. Peskin. Implicit second-order immersed boundary methods with boundary mass. *Computer Methods in Applied Mechanics and Engineering*, 197(25-28):2049–2067, 2008.
[63] L. Zhu and C.S. Peskin. Simulation of a flapping flexible filament in a flowing soap film by the immersed boundary method. *J. Comp. Phys.*, 179(2):452–468, 2002.
[64] Y. Kim and C.S. Peskin. Penalty immersed boundary method for an elastic boundary with mass. *Physics of Fluids*, 19:053103, 2007.
[65] Y. Kim and C.S. Peskin. Numerical study of incompressible fluid dynamics with nonuniform density by the immersed boundary method. *Physics of Fluids*, 20:062101, 2008.
[66] R. Delgado-Buscalioni and F. B. Usabiaga. Particle hydrodynamics using hybrid models: from molecular to colloidal fluids. In E. Oñate and R. Owen, editors, *Particle-Based Methods II: Fundamentals and Applications*. International Center for Numerical Methods in Engineering, 2011.
[67] A. J. Chorin. Numerical Solution of the Navier-Stokes Equations. *J. Math. Comp.*, 22:745–762, 1968.
[68] R. Zwanzig and M. Bixon. Compressibility effects in the hydrodynamic theory of Brownian motion. *J. Fluid Mech.*, 69:21–25, 1975.
[69] H. Grabert. *Projection operator techniques in nonequilibrium statistical mechanics*. Springer-Verlag, 1982.
[70] P. Español. Statistical mechanics of coarse-graining. *Novel Methods in Soft Matter Simulations*, pages 69–115, 2004.
[71] P.J. Morrison. Hamiltonian description of the ideal fluid. *Rev. Mod. Phys.*, 70(2):467, 1998.
[72] H. C. Öttinger. *Beyond equilibrium thermodynamics*. Wiley Online Library, 2005.
[73] B. Cichocki and B. U. Felderhof. Velocity autocorrelation function of interacting Brownian particles. *Phys. Rev. E*, 51:5549–5555, 1995.
[74] B.E. Griffith, R.D. Hornung, D.M. McQueen, and C.S. Peskin. An adaptive, formally second order accurate version of the immersed boundary method. *J. Comput. Phys.*, 223(1):10–49, 2007. Software available at `http://ibamr.googlecode.com`.
[75] B.E. Griffith. An accurate and efficient method for the incompressible Navier-Stokes equations using the projection method as a preconditioner. *J. Comp. Phys.*, 228(20):7565–7595, 2009.
[76] P. Español and I. Zúñiga. On the definition of discrete hydrodynamic variables. *J. Chem. Phys*, 131:164106, 2009.
[77] T.T. Bringley and C.S. Peskin. Validation of a simple method for representing spheres and slender bodies in an immersed boundary method for Stokes flow on an unbounded domain. *J. Comp. Phys.*, 227(11):5397–5425, 2008.
[78] Yang Liu and Yoichiro Mori. Properties of discrete delta functions and local convergence of the immersed boundary method. *SIAM Journal on Numerical Analysis*, 50(6):2986–3015, 2012.
[79] B.E. Griffith, X. Luo, D.M. McQueen, and C.S. Peskin. Simulating the fluid dynamics of natural and prosthetic heart valves using the immersed boundary method. *International Journal of Applied Mechanics*, 1(01):137–177, 2009.
[80] A.J. Chorin. Numerical simulation of the Navier-Stokes equations. *Mathematics of Computation*, 22(104):745–762, 1968.
[81] D. Devendran and C. S. Peskin. An immersed boundary energy-based method for incompressible viscoelasticity. *J. Comp. Phys.*, 231(14):4613–4642, 2012.
[82] A.J.C. Ladd. Numerical simulations of particulate suspensions via a discretized Boltzmann equation. II. Numerical results. *Journal of Fluid Mechanics*, 271(1):311–339, 1994.
[83] S. Melchionna. Incorporation of smooth spherical bodies in the Lattice Boltzmann method. *J. Comp. Phys.*, 230:3966–3976, 2011.
[84] R. Tatsumi and R. Yamamoto. Direct numerical simulation of dispersed particles in a compressible fluid. *Phys. Rev. E*, 85:066704, 2012.
[85] T. Iwashita, Y. Nakayama, and R. Yamamoto. Velocity Autocorrelation Function of Fluctuating Particles in Incompressible Fluids: Toward Direct Numerical Simulation of Particle Dispersions. *Progress of Theoretical Physics Supplement*, 178:86–91, 2009.
[86] P. J. Atzberger. Velocity correlations of a thermally fluctuating Brownian particle: A novel model of

the hydrodynamic coupling. *Physics Letters A*, 351(4-5):225–230, 2006.
[87] E. J. Hinch. Application of the Langevin equation to fluid suspensions. *J. Fluid Mech.*, 72(03):499–511, 1975.
[88] L.D. Landau and E.M. Lifshitz. *Fluid Mechanics*, volume 6 of *Course of Theoretical Physics*. Pergamon Press, Oxford, England, 1959.
[89] B. J. Alder and T. E. Wainwright. Decay of the velocity autocorrelation function. *Phys. Rev. A*, 1(1):18–21, 1970.
[90] Y. Pomeau and P. Résibois. Time dependent correlation functions and mode-mode coupling theories. *Phys. Rep.*, 19:63–139, June 1975.
[91] C. W. J. Beenakker. Ewald sum of the Rotne-Prager tensor. *J. Chem. Phys.*, 85:1581, 1986.
[92] Rahul Kekre, Jason E. Butler, and Anthony J. C. Ladd. Comparison of lattice-Boltzmann and Brownian-dynamics simulations of polymer migration in confined flows. *Phys. Rev. E*, 82:011802, 2010.
[93] R. Clift, J. R. Grace, and M. E. Weber. *Bubbles, drops and Particles*. Academic Press, 1978.
[94] D. Ormières and M. Provansal. Transition to Turbulence in the Wake of a Sphere. *Physical Review Letters*, 83(1):5–8, 1999.
[95] T. A. Johnson and V. C. Patel. Flow past a sphere up to a Reynolds number of 300. *Journal of Fluid Mechanics*, 378(-1):19–70, 1999.
[96] B.E. Griffith and S. Lim. Simulating an elastic ring with bend and twist by an adaptive generalized immersed boundary method. *Commun. Comput. Phys.*, 12:433–461, 2012.
[97] S. Lim, A. Ferent, X. S. Wang, and C. S. Peskin. Dynamics of a closed rod with twist and bend in fluid. *SIAM J Sci Comput*, 31(1):273–302, 2008.
[98] A. P. S. Bhalla, R. Bale, B. E. Griffith, and N. A. Patankar. A unified mathematical framework and an adaptive numerical method for fluid-structure interaction with rigid, deforming, and elastic bodies. *Journal of Computational Physics*, 250:446–476, 2013.
[99] J. T. Hynes, Raymond K., and Michael W. Molecular theory of translational diffusion: Microscopic generalization of the normal velocity boundary condition. *J. Chem. Phys.*, 70(3):1456, February 1979.
[100] S. Delong, Y. Sun, B. E. Griffith, E. Vanden-Eijnden, and A. Donev. Multiscale temporal integrators for fluctuating hydrodynamics. *Phys. Rev. E*, 90:063312, 2014. Software available at `https://github.com/stochasticHydroTools/MixingIBAMR`.
[101] N. K. Voulgarakis and J.-W. Chu. Bridging fluctuating hydrodynamics and molecular dynamics simulations of fluids. *J. Chem. Phys.*, 130(13):134111, 2009.