

Spectral Ewald Acceleration of Stokesian Dynamics for polydisperse suspensions

Mu Wang^{a,*}, John F. Brady^a

^a*Division of Chemistry and Chemical Engineering, California Institute of Technology, Pasadena, California 91125, USA*

Abstract

In this work we develop the Spectral Ewald Accelerated Stokesian Dynamics (SEASD), a novel computational method for dynamic simulations of polydisperse colloidal suspensions with full hydrodynamic interactions. SEASD is based on the framework of Stokesian Dynamics (SD) with extension to compressible solvents, and uses the Spectral Ewald (SE) method [Lindbo & Tornberg, *J. Comput. Phys.* **229** (2010) 8994] for the wave-space mobility computation. To meet the performance requirement of dynamic simulations, we use Graphic Processing Units (GPU) to evaluate the suspension mobility, and achieve an order of magnitude speedup compared to a CPU implementation. For further speedup, we develop a novel far-field block-diagonal preconditioner to reduce the far-field evaluations in the iterative solver, and SEASD-nf, a polydisperse extension of the mean-field Brownian approximation of Banchio & Brady [*J. Chem. Phys.* **118** (2003) 10323]. We extensively discuss implementation and parameter selection strategies in SEASD, and demonstrate the spectral accuracy in the mobility evaluation and the overall $O(N \log N)$ computation scaling. We present three computational examples to further validate SEASD and SEASD-nf in monodisperse and bidisperse suspensions: the short-time transport properties, the equilibrium osmotic pressure and viscoelastic moduli, and the steady shear Brownian rheology. Our validation results show that the agreement between SEASD and SEASD-nf is satisfactory over a wide range of parameters, and also provide significant insight into the dynamics of polydisperse colloidal suspensions.

Keywords: Stokes flow, Stokesian Dynamics, Brownian Dynamics, GPU computation, Ewald summation, spectral accuracy, colloidal suspensions, polydispersity

1. Introduction

Colloidal suspensions are dispersions of small particles in a viscous solvent, and are found in almost every aspect of our life, ranging from dairy milk to printer ink. They have two distinguishing features: (i) Brownian motion of the particles due to thermal fluctuations, and (ii) the long-range, non-pairwise-additive hydrodynamic interactions (HIs) mediated by the solvent. As a result of these features, dispersions exhibit many surprising behaviors such as non-Newtonian rheology, glass transitions, phase transitions, *etc.*, and have attracted extensive scientific and engineering interests [1]. Using monodisperse colloidal suspensions as a model system, significant understanding has been achieved through theoretical, simulation, and experimental studies.

However, naturally occurring colloidal suspensions are seldom monodisperse, and particle size differences are often unavoidable. In addition, particle size disparity introduces phenomena otherwise not observed in monodisperse suspensions. For example, size polydispersity reduces suspension viscosity [2–4], softens and even melts colloidal glasses [5], and promotes particle segregation in pressure driven flows [6]. Apparently, these behaviors can only be understood by studying dynamics of polydisperse colloidal suspensions.

In this work we develop a computational method based on the framework of Stokesian Dynamics [7] (SD) for fast and realistic dynamic simulations of dense, polydisperse colloidal suspensions, with a focus on suspension rheology.

*Corresponding author.

Email addresses: mwwang@caltech.edu (Mu Wang), jfbrady@caltech.edu (John F. Brady)

Presently, theoretical and computational studies on polydisperse colloidal suspensions, even for the simplest case of neutrally buoyant hard-sphere particles, are scarce, and heavily focus on the dilute or the short-time limits [8–12]: the former restricts HIs to the two- or three-body level, and the latter ignores suspension dynamic evolution, particularly the influence of Brownian motion. Beyond these limiting cases, we are only aware of the work of Ando & Skolnick [13], who studied particle diffusion in dense polydisperse colloidal suspensions using conventional SD in the context of biological molecular crowding. Their implementation limits HIs to the force-torque level, and therefore is unsuitable for rheological investigations.

A difficulty in dynamic simulations of dense colloidal suspensions is the singular HIs due to the lubrication interactions between close particle pairs. To directly resolve HIs, a computational method must capture the flow details in the small gap between particles. For multipole expansion based methods [7, 14, 15], a large number of expansion terms are necessary to achieve convergence, and for methods based on surface or spatial discretization, such as the boundary element method [16, 17] or direct numerical simulations [18–21], very fine meshing is needed in the gap. Directly resolving lubrication interactions drastically increases the computational cost and limits many studies to low volume fractions. For example, the force coupling method study of Abbas *et al.* [22] on the dynamics of non-Brownian bidisperse suspensions is limited to particle volume fractions below 20%.

A solution to the above difficulty is the SD framework [7], which exploits the local and pairwise additive nature of lubrication interactions. In SD, the long-range, non-pairwise-additive HIs are computed from the mobility perspective using low-order multipole expansions, and for particles in close contact, lubrication corrections are added pairwise to the corresponding resistance formalism. The corrections are based on the solutions of two-body problems with the far-field contributions removed. In this way, SD avoids directly resolving the singular lubrication interactions. The idea of lubrication correction in SD is general enough for incorporation to other computational methods. For example, similar lubrication corrections has been developed for hydrodynamic multipole methods [14, 15, 23, 24], the force coupling method [25], the lattice Boltzmann method [26], and the fictitious domain method [27]. Moreover, with an appropriate fluid solver, the lubrication corrections can be improved beyond the pairwise level [28]. We feel that, by incorporating the lubrication corrections, many recent computational techniques can significantly extend their accessible parameter range without an increased computational burden. This point is demonstrated in the present work, which essentially combines the lubrication corrections and the Spectral Ewald (SE) method of Lindbo & Tornberg [29, 30] for dynamic simulations of dense polydisperse suspensions.

The Spectral Ewald (SE) method is a new particle mesh technique for computing long-range electrostatic [30] or hydrodynamic [29] interactions, and has recently been incorporated into the boundary element method for soft particles [31]. Particle mesh techniques including the Particle Mesh Ewald (PME) method [32] and the Smooth Particle Mesh Ewald (SPME) method [33] have been extensively used for calculating HIs with $O(N \log N)$ computation scaling. Note that, although algorithms based on the fast multipole method [34] can achieve a better computation scaling—down to $O(N)$, they often have significant computation overheads, and require large system sizes to justify the complexity [35]. Therefore, for many dynamic simulations, the particle mesh techniques remain the practical choice. Notable examples are Accelerated Stokesian Dynamics (ASD) [36] which uses the PME method for the far-field mobility evaluation, and the work of Saintillan *et al.* [37], where the SPME method is employed to study fiber sedimentation. Compared to other particle mesh techniques, the SE method is spectrally accurate, and can separate errors from mesh interpolation and the wave-space truncation. Both features are essential for capturing the complicated HIs in polydisperse suspensions.

Another challenge in dynamic simulations of colloidal suspensions is Brownian motion, which is configuration dependent due to the fluctuation-dissipation relation. When Euler-Maruyama time integration is used, the deterministic particle drift due to the Brownian motion must also be included [38]. As a result, computing Brownian related quantities requires the gradient and the square root of the mobility tensor. Fortunately, these quantities can be evaluated in a matrix-free manner under the framework of ASD, making dynamic studies on hundreds of colloidal particles possible [39, 40]. Moreover, the mean-field Brownian approximation, which estimates the mobility tensor based on the near-field HIs, is able to further speed up the computations [39, 41]. In this work, these developments are fully incorporated for the dynamic simulation of Brownian polydisperse suspensions. Note that a different approach to treat the Brownian motion is based on fluctuating hydrodynamics [42], where the thermal fluctuations are directly incorporated in the governing fluid equations. It has been applied to the lattice Boltzmann method [43], the force coupling method [44], and the immersed boundary method [45].

The emergence of the General Purpose Graphic Processing Unit (GPGPU) programming often brings significant,

sometimes orders of magnitude, speed improvements for many existing algorithms. Recently, Kopp & Höfling [46] implemented the conventional SD for infinite solvent using GPGPU with direct HI summation. Despite the $O(N^2)$ scaling, they achieved impressive speedup over the CPU implementation. However, to study the dynamics of homogeneous suspensions, further extension to periodic systems are necessary. On the other hand, GPU acceleration of the SPME method [47, 48] in molecular dynamics provides access to millisecond-scale dynamics on personal computers. These acceleration techniques are applicable to particle mesh techniques in general, and inspired the present work. In particular, we used GPGPU programming to compute the HIs with the SE method in homogeneous suspensions, and realized almost an order of magnitude speedup in dynamic simulations.

Furthermore, our computation method extends SD to compressible suspensions, allowing dynamic simulations of constant pressure rheology [49] without introducing geometric confinement. This is possible because the flow disturbances due to rigid particles in a compressible solvent are incompressible and satisfy the Stokes equation [50]. Another benefit of such extension is that the suspension normal stress, which is essential for particle migration in sheared suspensions [51–53], can be directly evaluated.

The remainder of the paper is arranged as follows: Sec. 2 establishes the basic formalism for HIs in compressible Stokes flow. In Sec. 3, various aspects of mobility computations with the SE method are presented. Here, we also discuss different approaches to incorporate particle size polydispersity and the GPGPU implementation. In Sec. 4, we present the Spectral Ewald Accelerated Stokesian Dynamics (SEASD) and its mean-field Brownian approximation, SEASD-nf, for dynamic simulations of Brownian polydisperse suspensions. In Sec. 5 we carefully discuss the accuracy and parameter selections for the SE method, and the computation scaling of various SEASD implementations. Sec. 6 presents a series of validation calculations for monodisperse and bidisperse suspensions with SEASD and SEASD-nf: Sec. 6.1 addresses the short-time transport properties, Sec. 6.2 evaluates the equilibrium osmotic pressure and viscoelastic moduli, and Sec. 6.3 presents various aspects of the steady shear rheology of Brownian suspensions. The results also reveal the role of particle sizes in the dynamics of bidisperse suspensions. Finally, we conclude this work with a few comments in Sec. 7.

2. Hydrodynamic interactions in (compressible) Stokes flow

2.1. The mobility and resistance formalism

We first consider a suspension of N spherical rigid particles, each with radius a_i and position \mathbf{r}_i , in an *incompressible* solvent of viscosity η_0 and density ρ_0 , occupying a volume V . For the special case of bidisperse suspensions with particle sizes a_1 and a_2 , the suspension composition is fully characterized by three dimensionless parameters,

$$\lambda = a_1/a_2, \phi = \phi_1 + \phi_2, \text{ and } y_2 = \phi_2/\phi, \quad (1)$$

where λ is the size ratio, ϕ is the total volume fraction, and y_2 is the volume ratio of species 2. The species volume fraction is $\phi_\alpha = \frac{4}{3}\pi a_\alpha^3 n_\alpha$, $\alpha \in \{1, 2\}$, and the species number density is n_α . The total number density satisfies $n = n_1 + n_2$, and the species number fraction is $x_\alpha = n_\alpha/n$. Without loss of generality, we take $a_2 > a_1$.

If the particles are sufficiently small, the particle Reynolds number $\text{Re}_{p,\alpha} = \rho_0 a_\alpha U_\alpha / \eta_0 \ll 1$, where U_α is the species characteristic velocity. In this limit, the velocity field $\mathbf{v}(\mathbf{r})$ and the pressure field $p(\mathbf{r})$ of the solvent satisfy the Stokes equation,

$$\nabla p = \eta_0 \nabla^2 \mathbf{v}, \quad \nabla \cdot \mathbf{v} = 0, \quad (2)$$

supplemented by no-slip boundary conditions on particle surfaces. Due to the linearity of Eq. (2), there is a linear relation between the velocity disturbance on the surface of a particle i , \mathbf{u}'_i , and the surface force density of another particle j , \mathbf{f}_j ,

$$\mathbf{u}'_i(\mathbf{r}) = - \int d\mathbf{r}' \sum_j \mathbf{M}_{ij}(\mathbf{r}, \mathbf{r}'; X) \cdot \mathbf{f}_j(\mathbf{r}'), \quad (3)$$

where $\mathbf{M}_{ij}(\mathbf{r}, \mathbf{r}'; X)$ is a mobility operator depending on positions \mathbf{r} and \mathbf{r}' and the suspension configuration $X = \{\mathbf{r}_1, \mathbf{r}_2, \dots\}$. The surface force density is localized on the particle surface, *i.e.*, $\mathbf{f}_j(\mathbf{r}) = \boldsymbol{\sigma}(\mathbf{r}) \cdot \mathbf{n}_j \delta(\|\mathbf{r}\| - a_j)$, where $\boldsymbol{\sigma}$ is the stress tensor, \mathbf{n}_j is the surface normal of particle j , and $\delta(x)$ is the Dirac delta function. The stress tensor $\boldsymbol{\sigma} = -p\mathbf{I} + \eta_0[\nabla\mathbf{v} + (\nabla\mathbf{v})^\dagger]$, with \dagger indicating transposition and \mathbf{I} is the idem tensor. The velocity disturbance $\mathbf{u}'_i(\mathbf{r}) =$

$\mathbf{U}_i + \boldsymbol{\Omega}_i \times (\mathbf{r} - \mathbf{r}_i) - \mathbf{v}^\infty(\mathbf{r})$, where $\mathbf{v}^\infty(\mathbf{r})$ is the ambient flow satisfying $\nabla \cdot \mathbf{v}^\infty = 0$, and \mathbf{U}_i and $\boldsymbol{\Omega}_i$ are respectively the linear and angular velocities of particle i . By stacking the force density vectors $\mathbf{f} = (f_1, f_2, \dots)^\dagger$ and the velocity disturbance vectors $\mathbf{u}' = (u'_1, u'_2, \dots)^\dagger$ the grand mobility operator \mathbf{M} is constructed from elements \mathbf{M}_{ij} in Eq. (3), such that

$$\mathbf{u}'(\mathbf{r}) = - \int d\mathbf{r}' \mathbf{M}(\mathbf{r}, \mathbf{r}'; X) \cdot \mathbf{f}(\mathbf{r}'), \quad (4)$$

for the N particles in the suspension. Eqs. (3) and (4) are known as the mobility formalism, and the inverse relation is the resistance formalism,

$$\mathbf{f}(\mathbf{r}) = - \int d\mathbf{r}' \mathbf{R}(\mathbf{r}, \mathbf{r}'; X) \cdot \mathbf{u}'(\mathbf{r}'), \quad (5)$$

where $\mathbf{R}(\mathbf{r}, \mathbf{r}'; X)$ is the grand resistance operator.

The integral representations in Eqs. (4) and (5) can be equivalently expressed as multipole expansions of $\mathbf{f}(\mathbf{r})$ and $\mathbf{u}'(\mathbf{r})$, \mathfrak{f} and \mathbf{u}' respectively, around the particle centers, *i.e.*,

$$\mathbf{f}(\mathbf{r}) \rightarrow \mathfrak{f} = \begin{bmatrix} \mathcal{F}^{\text{H}} \\ \mathbf{S}^{\text{H}} \\ \vdots \end{bmatrix} \text{ and } \mathbf{u}'(\mathbf{r}) \rightarrow \mathbf{u}' = \begin{bmatrix} \mathcal{U}' \\ -\mathbf{E}^\infty \\ \vdots \end{bmatrix}, \quad (6)$$

where \mathcal{F}^{H} is the generalized hydrodynamic force, \mathbf{S}^{H} is the hydrodynamic stresslet, \mathcal{U}' is the generalized velocity disturbance, and \mathbf{E}^∞ is the rate of strain tensor for the ambient flow. Note that $\mathcal{F}^{\text{H}} = (\mathbf{F}^{\text{H}}, \mathbf{T}^{\text{H}})^\dagger$, where \mathbf{F}^{H} and \mathbf{T}^{H} are respectively the particle hydrodynamic force and torque for all particles, and $\mathcal{U}' = (\mathbf{U} - \mathbf{U}^\infty, \boldsymbol{\Omega} - \boldsymbol{\Omega}^\infty)^\dagger$, where $\mathbf{U} - \mathbf{U}^\infty$ and $\boldsymbol{\Omega} - \boldsymbol{\Omega}^\infty$ are respectively the linear and angular velocity disturbances. The hydrodynamic force, torque, and stresslet for particle i are defined as integrals of the localized surface force density \mathbf{f}_i ,

$$\mathbf{F}_i^{\text{H}} = - \int d\mathbf{r} \mathbf{f}_i(\mathbf{r}), \quad (7)$$

$$\mathbf{T}_i^{\text{H}} = - \int d\mathbf{r} (\mathbf{r} - \mathbf{r}_i) \times \mathbf{f}_i(\mathbf{r}), \quad (8)$$

$$\mathbf{S}_i^{\text{H}} = - \int d\mathbf{r} \frac{1}{2} [(\mathbf{r} - \mathbf{r}_i) \mathbf{f}_i + \mathbf{f}_i (\mathbf{r} - \mathbf{r}_i)]. \quad (9)$$

In Eq. (6) the ambient velocities are evaluated at particle centers, *i.e.*, $\mathbf{U}_i^\infty = \mathbf{v}^\infty(\mathbf{r}_i)$, $\boldsymbol{\Omega}_i^\infty = \frac{1}{2} \nabla \times \mathbf{v}^\infty|_{\mathbf{r}_i}$, and $\mathbf{E}^\infty = \frac{1}{2} [\nabla \mathbf{v}^\infty + (\nabla \mathbf{v}^\infty)^\dagger]_{\mathbf{r}_i}$. The expansions in Eqs. (4) and (5) lead to the following infinite dimension linear relation,

$$\mathfrak{f} = -\mathfrak{M}(X) \cdot \mathbf{u}' \text{ and } \mathbf{u}' = -\mathfrak{R}(X) \cdot \mathfrak{f} \quad (10)$$

where $\mathfrak{M}(X)$ and $\mathfrak{R}(X)$ are the multipole grand mobility and resistance tensors of operators $\mathbf{M}(\mathbf{r}, \mathbf{r}'; X)$ and $\mathbf{R}(\mathbf{r}, \mathbf{r}'; X)$, respectively. Evidently, $\mathfrak{M} = \mathfrak{R}^{-1}$, and from the Lorentz reciprocal theorem [54], both are positive definite.

The infinite dimension vectors \mathfrak{f} and \mathbf{u}' can be reduced to finite dimensions by projection. To the stresslet level of \mathfrak{f} and the strain rate level of \mathbf{u}' , we introduce projection matrices \mathcal{P} and \mathcal{Q} , such that $\mathcal{P} \cdot \mathfrak{f} = (\mathcal{F}^{\text{H}}, \mathbf{S}^{\text{H}})^\dagger$ and $\mathcal{Q} \cdot \mathbf{u}' = (\mathcal{U}', -\mathbf{E}^\infty)^\dagger$. Moreover, $\mathcal{P} \cdot \mathcal{P}^\dagger = \mathcal{Q} \cdot \mathcal{Q}^\dagger = \mathcal{I}$, where \mathcal{I} is an identity matrix. The following linear relation holds:

$$\begin{bmatrix} \mathcal{U}' \\ -\mathbf{E}^\infty \end{bmatrix} = -\mathcal{M} \cdot \begin{bmatrix} \mathcal{F}^{\text{H}} \\ \mathbf{S}^{\text{H}} \end{bmatrix}, \text{ and } \mathcal{R} = \mathcal{M}^{-1}, \quad (11)$$

where $\mathcal{M} = \mathcal{Q} \mathfrak{M} \mathcal{P}^\dagger$ is the (exact) grand mobility tensor and $\mathcal{R} = \mathcal{P} \mathfrak{R} \mathcal{Q}^\dagger$ is the (exact) grand resistance tensor. For convenience, the grand resistance tensor is partitioned as

$$\mathcal{R} = \begin{bmatrix} \mathbf{R}_{\mathcal{F}\mathcal{U}} & \mathbf{R}_{\mathcal{F}\mathbf{E}} \\ \mathbf{R}_{\mathcal{S}\mathcal{U}} & \mathbf{R}_{\mathcal{S}\mathbf{E}} \end{bmatrix}, \quad (12)$$

where, for example, $\mathbf{R}_{\mathcal{F}\mathcal{U}}$ describes the coupling between the generalized force and the generalized velocity. The linear relation in Eq. (11) can also be deduced from the linearity of Eq. (2) without appealing to the multipole expansion, but here we establish a connection with other works, particularly the multipole methods of Cichocki and

coworkers [15, 55]. Note that for rigid spherical particles, external flows can only affect the first two moments of \mathbf{f} and \mathbf{u}' due to symmetry and the no-slip boundary condition.

Elements of \mathfrak{M} and \mathfrak{R} can be computed from, for example, the induced force multipole [56, 57], eigenfunction expansions [15, 24, 58], and multipole expansions [7]. To the stresslet level, \mathfrak{M} can be conveniently evaluated by combining the Faxén formulae and the multipole expansions. For a rigid particle i in an incompressible solvent, the Faxén formulae are [7],

$$U_i - U^\infty = -\frac{\mathbf{F}_i^{\text{H}}}{6\pi\eta_0 a_i} + \left(1 + \frac{1}{6}a_i^2 \nabla^2\right) \mathbf{v}'|_{r_i} \quad (13)$$

$$\Omega_i - \Omega^\infty = -\frac{\mathbf{T}_i^{\text{H}}}{8\pi\eta_0 a_i^3} + \frac{1}{2} \nabla \times \mathbf{v}'|_{r_i} \quad (14)$$

$$-\overline{\mathbf{E}^\infty} = -\frac{\overline{\mathbf{S}_i^{\text{H}}}}{\frac{20}{3}\pi\eta_0 a_i^3} + \left(1 + \frac{1}{10}a_i^2 \nabla^2\right) \frac{1}{2} [\nabla \mathbf{v}' + (\nabla \mathbf{v}')^\dagger]|_{r_i}, \quad (15)$$

where the overline indicates the traceless part of the symmetric tensor, and $\mathbf{v}'(\mathbf{r})$ is the velocity field in the absence of particle i . With the fundamental solution of Stokes equation $\mathbf{J}(\mathbf{r})$ and the force density \mathbf{f} , the velocity field $\mathbf{v}'(\mathbf{r})$ can be computed as [54],

$$\mathbf{v}(\mathbf{r}) = -\frac{1}{8\pi\eta_0} \int d\mathbf{r}' \mathbf{J}(\mathbf{r} - \mathbf{r}') \cdot \mathbf{f}(\mathbf{r}'). \quad (16)$$

Expanding the force density around particle centers, we have

$$\mathbf{v}'(\mathbf{r}) = \frac{1}{8\pi\eta_0} \sum_j' \left(1 + \frac{1}{6}a_j^2 \nabla^2\right) \mathbf{J} \cdot \mathbf{F}_j^{\text{H}} + \mathbf{R} \cdot \mathbf{T}_j^{\text{H}} - \left(1 + \frac{1}{10}a_j^2\right) \mathbf{K} : \overline{\mathbf{S}_j^{\text{H}}} + \dots, \quad (17)$$

where the prime on the summation excludes the case $i = j$, and the functions \mathbf{J} , \mathbf{R} , and \mathbf{K} are evaluated at $\mathbf{r} - \mathbf{r}_j$. In the Cartesian tensor form, $\mathbf{R} = R_{\alpha\beta} = \frac{1}{4}\epsilon_{\delta\gamma\beta}(\nabla_\gamma J_{\alpha\delta} - \nabla_\delta J_{\alpha\gamma})$ and $\mathbf{K} = K_{\alpha\beta\gamma} = \frac{1}{2}[\nabla_\gamma J_{\alpha\beta} + \nabla_\beta J_{\alpha\gamma}]$, with $\epsilon_{\alpha\beta\gamma}$ the Levi-Civita symbol. With Eqs. (13)–(15) and (17), the grand mobility tensor \mathfrak{M} for incompressible solvents can be constructed in a pairwise fashion.

2.2. The fundamental solutions

The formalism in Sec. 2.1 relies on $\mathbf{J}(\mathbf{r})$, the fundamental solution of Stokes equation. Different boundary conditions such as periodicity [59, 60], confinement [24, 61], or a combination of both [62], can be incorporated to $\mathbf{J}(\mathbf{r})$. For an infinite expanse of fluid, we have the well-known Oseen tensor,

$$\mathbf{J}(\mathbf{r}) = \frac{1}{r}(\mathbf{I} + \hat{\mathbf{r}}\hat{\mathbf{r}}), \quad (18)$$

where $r = \|\mathbf{r}\|$ and $\hat{\mathbf{r}} = \mathbf{r}/r$.

To study dynamics of homogeneous suspensions, periodic boundary conditions are necessary to assess the HIs. In this case, the proper fundamental solution $\mathbf{J}(\mathbf{r})$ describes the fluid velocity disturbance due to an array of periodic forces $\mathbf{F} \sum_p \delta(\mathbf{r} - \mathbf{R}_p)$, where $\mathbf{R}_p = \sum_{d=1}^3 p_d \mathbf{a}_d$ is the location of the periodic forcing. Here, $\mathbf{p} = (p_1, p_2, p_3) \in \mathbb{Z}^3$, $\delta(\mathbf{r})$ is the 3D Dirac delta function, and \mathbf{a}_1 , \mathbf{a}_2 , and \mathbf{a}_3 are the Bravais lattice vectors describing the spatial periodicity. From Fourier expansion of Stokes equation [Eq. (2)], we have for the periodic $\mathbf{J}(\mathbf{r})$:

$$\mathbf{J}(\mathbf{r}) = -\frac{8\pi}{V} (\mathbf{I}\nabla^2 - \nabla\nabla) \sum_{\mathbf{k} \neq 0} \frac{1}{k^4} \exp(-i\mathbf{k} \cdot \mathbf{r}), \quad (19)$$

where $\iota = \sqrt{-1}$, the unit cell volume $V = \mathbf{a}_1 \cdot (\mathbf{a}_2 \times \mathbf{a}_3)$, the wave vector $\mathbf{k} = \sum_{d=1}^3 j_d \mathbf{b}_d$ is defined by the reciprocal vectors \mathbf{b}_1 , \mathbf{b}_2 , and \mathbf{b}_3 , $\mathbf{j} = (j_1, j_2, j_3) \in \mathbb{Z}^3$, and $k^2 = \mathbf{k} \cdot \mathbf{k}$. Writing the lattice and the reciprocal vectors as column vectors and defining matrices $\mathbf{A} = [\mathbf{a}_1 \mathbf{a}_2 \mathbf{a}_3]$ and $\mathbf{B} = [\mathbf{b}_1 \mathbf{b}_2 \mathbf{b}_3]$, we have $\mathbf{B}^\dagger = 2\pi \mathbf{A}^{-1}$ and $\exp(i\mathbf{k} \cdot \mathbf{R}_p) = 1$. By requiring $\mathbf{k} \neq 0$ in Eq. (19), the external forces are balanced by the pressure gradient [59], a necessary condition for convergent HIs [63].

A difficulty associated with HIs is the long range nature of $\mathbf{J}(\mathbf{r})$, *i.e.*, Eq. (18) decays as r^{-1} in the real space and Eq. (19) as k^{-2} in the wave space. For periodic systems, however, the conditionally converging sum in Eq. (19) can be split into two exponentially fast converging series, *i.e.*,

$$\mathbf{J}(\mathbf{r}) = \mathbf{J}_R(\mathbf{r}) + \mathbf{J}_W(\mathbf{r}), \quad (20)$$

where $\mathbf{J}_R(\mathbf{r})$ is the real-space sum, $\mathbf{J}_W(\mathbf{r})$ is the wave-space sum. Although the splitting in Eq. (20) is not unique [29], a particularly efficient scheme by Hasimoto [59] utilizes the integral

$$\frac{1}{k^4} = \pi^2 \int_0^\infty \beta \exp(-\pi k^2 \beta) d\beta, \quad (k \neq 0), \quad (21)$$

and the Poisson summation formula. The result is

$$\mathbf{J}_R(\mathbf{r}) = \sum_{p \neq 0} (\mathbf{I}\nabla^2 - \nabla\nabla) \left[r \operatorname{Erfc}(r\xi) - \frac{1}{\xi\sqrt{\pi}} e^{-r^2\xi^2} \right], \quad (22)$$

$$\mathbf{J}_W(\mathbf{r}) = \frac{8\pi}{V} \sum_{k \neq 0} (\mathbf{I}\nabla^2 - \nabla\nabla) \left(-1 - \frac{k^2}{4\xi^2} \right) \frac{1}{k^4} e^{-\frac{1}{4}k^2\xi^2} e^{-i\mathbf{k}\cdot\mathbf{r}}, \quad (23)$$

where ξ is the splitting parameter and $\operatorname{Erfc}(x)$ is the complementary error function. The real-space sum \mathbf{J}_R only covers the *neighboring* periodic cells. The parameter ξ is consistent with the convention of Beenakker[60] and satisfies $4\pi\alpha\xi^2 = 1$, where α is the splitting parameter introduced by Hasimoto [59].

2.3. Extension to compressible fluid

The formalism in Sec. 2.1 is limited to an incompressible fluid, *i.e.*, the imposed flow must satisfy $\nabla \cdot \mathbf{v}^\infty = 0$. This requirement is relaxed by imposing a *uniform* rate of expansion everywhere in the fluid, such that $\nabla \cdot \mathbf{v}^\infty = E^\infty$, and the fluid is assumed compressible with a bulk viscosity κ_0 . The rigid particles, unable to expand with the compressible fluid, generate a velocity disturbance that satisfies the incompressible Stokes equation [50]. From the linearity of Stokes flow, this velocity disturbance can be superimposed with other flows in the suspension, extending the existing formalism to compressible fluids.

For a rigid particle of radius a_i located at $\mathbf{r}_i = 0$, the velocity disturbance \mathbf{v}_s due to a compressible flow with an expansion rate E^∞ is

$$\mathbf{v}_s(\mathbf{r}) = -\frac{1}{3}a_i^3 E^\infty \frac{\mathbf{r}}{r^3}. \quad (24)$$

This isotropic flow disturbance generates an isotropic stress contribution. Introducing the pressure moment as the trace of the stresslet in Eq. (9), *i.e.*,

$$S_i^H = - \int d\mathbf{r} (\mathbf{r} - \mathbf{r}_i) \cdot \mathbf{f}_i(\mathbf{r}), \quad (25)$$

we have $S_i^H = -\frac{16}{3}\pi\eta_0 a_i^3 E^\infty$ from Eq. (24). Therefore, the velocity disturbance due to a pressure moment S_i^H at the origin is

$$\mathbf{v}_s(\mathbf{r}) = \frac{1}{16\pi\eta_0} \frac{\mathbf{r}}{r^3} S_i^H = \mathbf{Q}(\mathbf{r}) S_i^H. \quad (26)$$

Adding the compressible velocity disturbances $\mathbf{v}_s(\mathbf{r})$ from other particles to the incompressible velocity disturbance $\mathbf{v}'(\mathbf{r})$ in Eq. (17), the general velocity disturbance in a compressible suspension is

$$\mathbf{v}'_c(\mathbf{r}) = \mathbf{v}'(\mathbf{r}) + \sum_j' \mathbf{Q}(\mathbf{r} - \mathbf{r}_j) S_j^H. \quad (27)$$

When applying the Faxén formulae [Eqs. (13)–(15)] in compressible suspensions, the velocity disturbance \mathbf{v}'_c , instead of \mathbf{v}' , is used.

In addition to Eqs. (13)–(15), the Faxén relation for the pressure moment in a compressible fluid is [64, 65]

$$S_i^H = -\frac{16}{3}\pi\eta_0 a_i^3 E^\infty + 4\pi a_i^3 p'(\mathbf{r}_i), \quad (28)$$

where p' is the pressure disturbance without the particle at \mathbf{r}_i . The pressure disturbance can be obtained from the pressure fundamental solution of Stokes equation,

$$\mathbf{P}(\mathbf{r}) = \frac{\mathbf{r}}{r^3}, \quad (29)$$

such that the pressure distribution due to a force density is

$$p(\mathbf{r}) = -\frac{1}{4\pi} \int d\mathbf{r}' \mathbf{P}(\mathbf{r} - \mathbf{r}') \cdot \mathbf{f}(\mathbf{r}'). \quad (30)$$

For the pressure disturbance p' in Eq. (28), expanding the surface force densities leads to

$$p'(\mathbf{r}) = \frac{1}{4\pi} \sum_j' \mathbf{P}(\mathbf{r} - \mathbf{r}_j) \cdot \mathbf{F}_j^{\text{H}} - \nabla \mathbf{P} : \mathbf{S}_j^{\text{H}}|_{(\mathbf{r}-\mathbf{r}_j)} + \dots. \quad (31)$$

Eq. (28) is different from the Faxén formulae in Eqs. (13)–(15) as it presents the pressure moment or trace of the stresslet on the left hand side. This subtle difference highlights a distinct feature of the compressible flow disturbances: in a compressible fluid, the pressure moment can cause particle movement satisfying the incompressible Stokes equation, but the incompressible force moments cannot generate compressible disturbances. As a result, the interaction part of the pressure moment can only be evaluated after \mathbf{F}_i^{H} , \mathbf{T}_i^{H} , and $\overline{\mathbf{S}}_i^{\text{H}}$ are known. Otherwise, the resulting hydrodynamic interactions contain spurious contributions due to the unphysical coupling between the incompressible force moments and the compressible flow disturbances.

To extend the above results for \mathbf{v}_s and S_i^{H} to periodic boundary conditions, we note that the divergence of \mathbf{Q} in Eq. (26) satisfies

$$\nabla \cdot \mathbf{Q} = \frac{1}{4\eta_0} \delta(\mathbf{r}), \quad (32)$$

since $\nabla^2 r^{-1} = -4\pi\delta(\mathbf{r})$. This means that, for uniform expansion in compressible suspensions, the particles act as fluid sources, each with a strength proportional to its pressure moment. In a periodic system, the velocity disturbance corresponds to an array of sources are obtained by replacing the delta function in Eq. (32) with $\sum_p \delta(\mathbf{r} - \mathbf{R}_p)$. From Fourier transform, the solution is

$$\mathbf{Q}(\mathbf{r}) = \frac{1}{4\eta_0 V} \sum_{\mathbf{k} \neq 0} \nabla \frac{1}{k^2} e^{-i\mathbf{k} \cdot \mathbf{r}}. \quad (33)$$

The above wave-space sum can be split to two exponentially converging series [30, 59]

$$\sum_{\mathbf{k} \neq 0} \frac{1}{k^2} e^{-i\mathbf{k} \cdot \mathbf{r}} = \frac{V}{4\pi} \sum_{p \neq 0} \frac{1}{r} \text{Erfc}(r\xi) + \sum_{\mathbf{k} \neq 0} \frac{1}{k^2} e^{-\frac{1}{4}k^2\xi^2} e^{-i\mathbf{k} \cdot \mathbf{r}}. \quad (34)$$

Similar to $\mathbf{Q}(\mathbf{r})$, the pressure fundamental solution $\mathbf{P}(\mathbf{r})$ in Eq. (29) can also be extended to periodic systems.

3. The mobility computation

The mobility problem seeks the action of the grand mobility tensor \mathfrak{M} on the force moments such as \mathcal{F}^{H} and \mathbf{S}^{H} . It can be constructed in a pairwise fashion using the formalism in Sec. 2 for compressible suspensions. Naïvely, this is an $O(N^2)$ operation for an N -particle system since the long-range HIs necessitate considerations of all particle pairs. However, with the Ewald summation that splits the fundamental solutions $\mathbf{J}(\mathbf{r})$, $\mathbf{Q}(\mathbf{r})$, and $\mathbf{P}(\mathbf{r})$ into exponentially fast converging wave-space and real-space series, the particle mesh techniques can improve the computation scaling to $O(N \log N)$. In the following, our implementation of the mobility computation is discussed.

3.1. Wave-space computation: the Spectral Ewald (SE) method

The wave-space computation concerns the part of grand mobility tensor associated with $\mathbf{J}_w(\mathbf{r})$ of Eq. (23) and the wave-space sum of Eq. (34) in $\mathbf{P}(\mathbf{r})$ and $\mathbf{Q}(\mathbf{r})$. Using the Fast Fourier Transform (FFT) algorithm, the computation cost can be reduced to $O(N \log N)$. To illustrate this, let us consider the wave-space velocity disturbance \mathbf{U}_i^W on particle i at the Rotne-Prager level, obtained by combining Eqs. (13), (17), and (23), *i.e.*,

$$\mathbf{U}_i^W = \frac{1}{\eta_0 V} \sum_{\mathbf{k} \neq 0} e^{-i\mathbf{k} \cdot \mathbf{r}_i} \left(1 - \frac{1}{6} a_i^2 k^2\right) \mathbf{g}_1(\mathbf{k}) \cdot \sum_j \left(1 - \frac{1}{6} a_j^2 k^2\right) e^{i\mathbf{k} \cdot \mathbf{r}_j} \mathbf{F}_j^H, \quad (35)$$

and the wave-space kernel

$$\mathbf{g}_1(\mathbf{k}) = \left(1 + \frac{1}{4} k^2 \xi^{-2}\right) k^{-4} e^{-\frac{1}{4} k^2 \xi^{-2}} (\mathbf{I} k^2 - \mathbf{k} \mathbf{k}). \quad (36)$$

Different from Eq. (17), the summation over particle j in Eq. (35) is unrestricted and includes the case of $i = j$. Therefore, the self interaction term for $i = j$, which is

$$\frac{\xi(9 - 10a_i^2 \xi^2 + 7a_i^4 \xi^4)}{18\eta_0 \pi^{3/2}} \mathbf{F}_i^H, \quad (37)$$

should be removed later. Eq. (35) exposes the basic idea behind many particle mesh techniques including the PME method and the SPME method. From an inverse Fourier transform, the real-space force distribution corresponding to the summation over j in Eq. (35) is

$$\sum_j \left(1 + \frac{1}{6} a_j^2 \nabla^2\right) \mathbf{F}_j^H \delta(\mathbf{r} - \mathbf{r}_j). \quad (38)$$

The force distribution in Eq. (38) is assigned to a regular spatial grid by approximating the delta functions by Lagrangian polynomials in the PME method [66] or Cardinal B-splines in the SPME method [33]. The interpolated forces are then transformed to the wave space by FFT and the wave-space computation in Eq. (35) is performed. The wave-space results is then brought back to the real space by inverse FFTs. Subsequently, the velocity on each particle, \mathbf{U}_i^W , is interpolated back from the grid, preferably using the same interpolation scheme for the force assignment [67]. Here, the action of the mobility tensor on the force \mathbf{F}^H , rather than the tensor itself, is computed. The kernel $\mathbf{g}_1(\mathbf{k})$ in Eq. (36) is effectively a low-pass filter that cuts off the spatial signals at high k . Computationally, for M^3 grid points the FFT scales as $O(M^3 \log M^3)$. To ensure reasonable accuracy, $M^3 \propto N$, and the wave-space computation scales as $O(N \log N)$.

There are two sources of error affecting the accuracy of particle mesh techniques. The first is associated with the truncation of the wave-space sum (k -summation) in Eq. (35). This is only affected by the number of grid points M in the simulation box. The second error is the interpolation error, and arises from polynomial approximation of the δ -functions in Eq. (38). For a simulation box of size L , this error scales as $(L/M)^p$, where p is the polynomial order of the approximation scheme. Since both errors are associated with M , we cannot separate the two error sources. Consequently, to maintain a satisfactory overall accuracy, a large M is often used in the wave-space computations to keep the interpolation error small, resulting in unnecessary FFT computations.

In addition, for polydisperse suspensions, different particle sizes introduce additional complications to traditional particle mesh techniques. If the Laplacian in Eq. (38) is computed in the real space in the SPME method, the interpolation error increases to $(L/M)^{p-2}$, which further increases the M requirement. For the PME method, real-space differentiation is unsuitable due to the discontinuity of Lagrangian polynomials, and all the computations have to be carried out in the wave space. This significantly increases the total number of FFTs. In addition, different particle sizes increase the complexity in the algorithm implementation. Therefore, a simple method with flexible error control is crucial for accurate and efficient wave-space computation in polydisperse systems.

To address these concerns, we use a new particle mesh technique, the Spectral Ewald (SE) method [29–31] for the wave-space mobility computation. The SE method decouples the k -space truncation and interpolation errors, and is accurate, efficient, and flexible for polydisperse systems. To show this, we use Eq. (35) again as an example and consider the general case of non-orthogonal lattice vectors. We first introduce the fractional coordinate $\mathbf{t} = (t_1, t_2, t_3)^\dagger \in [0, 1]^3$. For each point \mathbf{r} in the simulation box, $\mathbf{r} = t_1 \mathbf{a}_1 + t_2 \mathbf{a}_2 + t_3 \mathbf{a}_3 = \mathbf{A} \cdot \mathbf{t}$. Accordingly, defining

$\mathbf{q} = (q_1, q_2, q_3)^\dagger$ such that $\mathbf{k} = q_1\mathbf{b}_1 + q_2\mathbf{b}_2 + q_3\mathbf{b}_3 = \mathbf{B} \cdot \mathbf{q}$, $\exp(i\mathbf{k} \cdot \mathbf{r}) = \exp(2\pi i\mathbf{q} \cdot \mathbf{t})$, and $k^2 = \mathbf{q}^\dagger \cdot \mathbf{B}^\dagger \cdot \mathbf{B} \cdot \mathbf{q}$. Eq. (35) is rewritten in \mathbf{t} and \mathbf{q} as

$$\begin{aligned} U_i^W = & \frac{1}{\eta_0 V} \sum_{\mathbf{q} \neq 0} e^{-2\pi i\mathbf{q} \cdot \mathbf{t}_i - \frac{1}{8}\theta q^2 \xi^{-2}} \left(1 - \frac{1}{6}a_i^2 \mathbf{q}^\dagger \cdot \mathbf{B}^\dagger \cdot \mathbf{B} \cdot \mathbf{q}\right) e^{\frac{1}{4}\theta q^2 \xi^{-2}} \mathbf{g}_1(\mathbf{B} \cdot \mathbf{q}) \\ & \cdot \sum_j \left(1 - \frac{1}{6}a_j^2 \mathbf{q}^\dagger \cdot \mathbf{B}^\dagger \cdot \mathbf{B} \cdot \mathbf{q}\right) e^{2\pi i\mathbf{q} \cdot \mathbf{t}_j - \frac{1}{8}\theta q^2 \xi^{-2}} \mathbf{F}_j^H, \end{aligned} \quad (39)$$

with two $e^{-\frac{1}{8}\theta q^2 \xi^{-2}}$ multiplied after particle positions and one $e^{\frac{1}{4}\theta q^2 \xi^{-2}}$ before \mathbf{g}_1 , and θ is a parameter. Introducing the Fourier transform pair

$$\hat{f}_{\mathbf{q}} = \int d\mathbf{t} f(\mathbf{t}) e^{2\pi i\mathbf{q} \cdot \mathbf{t}} \text{ and } f(\mathbf{t}) = \int d\mathbf{q} \hat{f}_{\mathbf{q}} e^{-2\pi i\mathbf{q} \cdot \mathbf{t}}, \quad (40)$$

the basic idea of SE is to note that

$$h(\mathbf{t}) = \int d\mathbf{q} e^{-2\pi i\mathbf{q} \cdot \mathbf{t} - \frac{1}{8}\theta q^2 \xi^{-2}} = \left(\frac{8\pi\xi^2}{\theta}\right)^{\frac{3}{2}} \exp\left(-\frac{8\pi^2\xi^2}{\theta} \|\mathbf{t}\|_*^2\right), \quad (41)$$

i.e., the Fourier transform of a Gaussian remains a Gaussian, and the shape of the Gaussian is controlled by θ . Here, $\|\cdot\|_*$ indicates distance computation using the minimum image convention for periodic systems. The inverse Fourier transform of the second line of Eq. (39) with respect to \mathbf{q} is

$$\mathbf{H}(\mathbf{t}) = \sum_j \left(1 + \frac{1}{24}a_j^2 \pi^{-2} \nabla_j^\dagger \cdot \mathbf{B}^\dagger \cdot \mathbf{B} \cdot \nabla_j\right) h|_{(\mathbf{t}-\mathbf{t}_j)} \mathbf{F}_j^H, \quad (42)$$

where $\nabla_j = (\partial/\partial t_1, \partial/\partial t_2, \partial/\partial t_3)^\dagger$. Eq. (42) facilitates interpolation of a discrete force distribution onto a uniform grid of coordinate \mathbf{t} via the Gaussian shape function $h(\mathbf{t})$ in Eq. (41). The effect of particle size is automatically incorporated in the grid assignment scheme in the real space. After converting the real-space $\mathbf{H}(\mathbf{t})$ to the wave-space $\hat{\mathbf{H}}_{\mathbf{q}}$ using FFTs, the wave-space computation produces

$$\hat{\mathbf{G}}_{\mathbf{q}} = \begin{cases} e^{\frac{1}{4}\theta q^2 \xi^{-2}} \mathbf{g}_1(\mathbf{B} \cdot \mathbf{q}) \cdot \hat{\mathbf{H}}_{\mathbf{q}}, & \mathbf{q} \neq 0 \\ 0 & \text{otherwise.} \end{cases} \quad (43)$$

From Parseval's theorem,

$$\int_{\mathbf{T}} d\mathbf{t} f(\mathbf{t}) g^*(\mathbf{t}) = \sum_{\mathbf{q}} \hat{f}_{\mathbf{q}} \hat{g}_{\mathbf{q}}^*, \quad (44)$$

where \mathbf{T} is a periodic lattice and $(\cdot)^*$ indicates complex conjugation, Eq. (39) becomes a convolution integral with the Gaussian shape function,

$$U_i^W = \frac{1}{\eta_0 V} \int_{\mathbf{T}} d\mathbf{t} \mathbf{G}(\mathbf{t}) \left(1 + \frac{1}{24}a_i^2 \pi^{-2} \nabla_i^\dagger \cdot \mathbf{B}^\dagger \cdot \mathbf{B} \cdot \nabla_i\right) h|_{(\mathbf{t}-\mathbf{t}_i)}, \quad (45)$$

where $\mathbf{G}(\mathbf{t})$ is the inverse Fourier transform of $\hat{\mathbf{G}}_{\mathbf{q}}$. Extending the SE method to couplings beyond Rotne-Prager level is straightforward, with adjusted $\mathbf{H}(\mathbf{t})$ and $\mathbf{G}(\mathbf{t})$ based on the Faxén laws and multipole expansions in Sec. 2. In this work, we have implemented the mobility computation to the stresslet and the strain rate level.

Unlike other particle mesh techniques, the SE formulation in Eqs. (39)–(45) is exact and therefore the errors are entirely from the numerical implementations. Since the FFT algorithm is accurate to machine precision, the sources of error include the discretization and truncation of the shape function [Eq. (41)], and the numerical integration in Eq. (45). Practically, the evaluation of each shape function is limited to P^3 points ($P \leq M$) around the particle. Due to the exponential decay of $h(\mathbf{t})$, the truncation error decreases exponentially with increasing P . Meanwhile, the integral in Eq. (45) is evaluated using trapezoidal quadrature [29, 30], which also exhibits exponential error decay with increasing P . Therefore, the interpolation error in SE method depends exclusively on P for sufficiently large M ,

and can be separately controlled from the k -space truncation error. The rapid, exponential error decay is known as spectral accuracy [29, 30], and this is the namesake of the SE method.

The computation cost of the SE method also becomes apparent with the truncation of $h(\mathbf{t})$. The grid assignment in Eq. (42) and the convolution Eq. (45) are $O(NP^3)$ for an N -particle system, and the FFTs to and from the wave space are $O[M^3 \log(M^3)]$. With $M^3 \propto N$, the time limiting step is the FFT, and the SE method also scales as $O(N \log N)$ as other particle mesh techniques.

The Gaussian shape in $h(\mathbf{t})$ of Eq. (41) is controlled by θ , which is parameterized as

$$\theta = \left(\frac{2\pi P\xi}{Mm} \right)^2, \quad (46)$$

on a regular grid of M^3 points with P^3 points for each shape function evaluation. The shape parameter m in Eq. (46) ensures that at the edge of $h(\mathbf{t})$ evaluation, *i.e.*, $t^2 = P^2/(2M)^2$, $h \propto e^{-m^2/2}$. Therefore, with fixed M and P , m describes the truncation of $h(\mathbf{t})$ on the discretized grid and is consistent with the original SE method of Lindbo & Tornberg [29, 30].

The computation efficiency of the SE method relies on rapidly computing the $O(NP^3)$ different Gaussian shape functions $h(\mathbf{t})$, which involves expensive exponential evaluations. To reduce these expensive operations, Lindbo & Tornberg [29, 30] introduced the fast Gaussian gridding (FGG) technique [68] to the SE method. In essence, the FGG technique evaluates the exponential function on a regular grid as

$$e^{-\alpha(\delta t + i\Delta t)^2} = e^{-\alpha(\delta t)^2} \times \left(e^{-2\alpha\delta t\Delta t} \right)^i \times \left[e^{-\alpha(\Delta t)^2} \right]^i, \quad (47)$$

where α is a constant, δt is the off-grid value, Δt is the spacing of the regular grid, and i is an integer within the range $[-P/2, P/2]$. It reduces the P exponential evaluations in each direction in the SE method to 3 exponential computations and at most $2P$ multiplications. In addition, the last term of Eq. (47) is independent of δt , and therefore only needs to be computed once.

3.2. Wave-space computation: the particle size effect

In Sec. 3.1 the terms associated with finite particle sizes in the Faxén laws and the multipole expansions are incorporated in the real-space derivatives of the shape function $h(\mathbf{t})$. For example, in a simple shear flow with lattice vectors $\mathbf{a}_1 = (L, 0, 0)$, $\mathbf{a}_2 = (\gamma L, L, 0)$, and $\mathbf{a}_3 = (0, 0, L)$, where γ is the strain, the relevant term in Eqs. (42) and (45) is

$$\begin{aligned} & \left(\frac{1}{24} a_i^2 \pi^{-2} \nabla_i^\dagger \cdot \mathbf{B}^\dagger \cdot \mathbf{B} \cdot \nabla_i \right) h(\mathbf{t}) = \\ & \frac{8}{3} \left(\frac{\pi\xi a_i}{\theta L} \right)^2 \left\{ -\theta(3 + \gamma^2) + 16\pi^2\xi^2 [(1 + \gamma^2)t_1^2 + t_2^2 + t_3^2 - 2\gamma t_1 t_2] \right\} h(\mathbf{t}). \end{aligned} \quad (48)$$

The finite particle sizes introduce additional features to the shape function, and for non-orthogonal simulation boxes, non-trivial anisotropy. As a result, compared to the case of point forces, more points P are needed to resolve the details in Eq. (48). On the other hand, the benefit of evaluating the particle size effects in the real space is that fewer FFTs are involved. To compute the mobility problem of compressible suspensions to the stresslet and the strain rate levels, only four pairs of FFTs are necessary: three are associated with \mathbf{J}_W in Eq. (23), and one associated with the \mathbf{Q} in Eq. (26).

Alternatively, the particle size effect can be completely accounted in the wave space. This requires, for each particle j , \mathbf{F}_j^H , \mathbf{T}_j^H , and \mathbf{S}_j^H , as well as $a_j^2 \mathbf{F}_j^H$ and $a_j^2 \mathbf{S}_j^H$, to be separately interpolated to the grid via $h(\mathbf{t})$ and brought to the wave space for computation. The derivatives associated with the Faxén laws and multipole expansions in Sec. 2 are carried out in the wave space as multiplication of wave vectors. The final results are then combined from different convolutions and weighted by the particle sizes. To demonstrate this, we again take the wave-space Rotne-Prager velocity, Eq. (39), as an example. In this approach, the grid assignment is split into two parts,

$$\mathbf{H}'(\mathbf{t}) = \sum_j h(\mathbf{t} - \mathbf{t}_j) \mathbf{F}_j^H \quad \text{and} \quad \mathbf{H}''(\mathbf{t}) = \sum_j h(\mathbf{t} - \mathbf{t}_j) a_j^2 \mathbf{F}_j^H. \quad (49)$$

The wave-space computation for $\mathbf{q} \neq 0$ is also split as

$$\hat{\mathbf{G}}'_q = e^{\frac{1}{4}\theta q^2 \xi^{-2}} \mathbf{g}_1(\mathbf{B} \cdot \mathbf{q}) \cdot \left[\hat{\mathbf{H}}'_q - \left(\frac{1}{6}\mathbf{q}^\dagger \cdot \mathbf{B}^\dagger \cdot \mathbf{B} \cdot \mathbf{q}\right) \hat{\mathbf{H}}''_q \right], \quad (50)$$

$$\hat{\mathbf{G}}''_q = \left(-\frac{1}{6}\mathbf{q}^\dagger \cdot \mathbf{B}^\dagger \cdot \mathbf{B} \cdot \mathbf{q}\right) e^{\frac{1}{4}\theta q^2 \xi^{-2}} \mathbf{g}_1(\mathbf{B} \cdot \mathbf{q}) \cdot \left[\hat{\mathbf{H}}'_q - \left(\frac{1}{6}\mathbf{q}^\dagger \cdot \mathbf{B}^\dagger \cdot \mathbf{B} \cdot \mathbf{q}\right) \hat{\mathbf{H}}''_q \right], \quad (51)$$

and $\hat{\mathbf{G}}'_q = \hat{\mathbf{G}}''_q = 0$ when $\mathbf{q} = 0$. The wave-space velocity disturbance is a sum of two convolutions,

$$U_i^W = \frac{1}{\eta_0 V} \int_{\mathbf{T}} d\mathbf{t} \mathbf{G}'(\mathbf{t}) h(\mathbf{t} - \mathbf{t}_i) + \frac{a_i^2}{\eta_0 V} \int_{\mathbf{T}} d\mathbf{t} \mathbf{G}''(\mathbf{t}) h(\mathbf{t} - \mathbf{t}_i). \quad (52)$$

Note that the convolution associated with $\mathbf{G}''(\mathbf{t})$ is weighted by the particle size a_i . Compared to the other approach, the wave-space computation is rather straightforward for the force interpolation and convolution. With the same P , the accuracy is expected to be higher as the derivatives are calculated in the wave space [67]. However, the computation burden is shifted to the FFTs: for the mobility problem to the \mathbf{S} and \mathbf{E} level, a total of 20 pairs of FFTs are necessary: 12 for \mathbf{F}_j^H , \mathbf{T}_j^H , and \mathbf{S}_j^H , three for $a_j^2 \mathbf{F}_j^H$, and five for the traceless part of $a_j^2 \mathbf{S}_j^H$.

A third approach, a hybridization between the wave- and the real-space approaches above, aims to reduce the errors associated with the high order derivatives of $h(\mathbf{t})$ in the real space. It retains the real-space derivatives in the force interpolation step, but when evaluating the Faxén laws, the second order derivatives are computed in the wave space for improved accuracy. The first order derivatives are computed in the real space to keep the total number of FFTs low. As a result, this hybrid approach requires 12 FFTs: four to the wave space and eight from the wave space. Taking Eq. (39) again for example, the most significant error in Sec. 3.1 is due to applying the operator $(\nabla_t^\dagger \cdot \mathbf{B}^\dagger \cdot \mathbf{B} \cdot \nabla_t)$ twice to $h(\mathbf{t})$, once during the force interpolation, and another time during the convolution. The hybrid approach retains the real-space grid assignment using $\mathbf{H}(\mathbf{t})$ in Eq. (42), but evaluates the convolution using Eq. (52) with modified $\hat{\mathbf{G}}'(\mathbf{t})$ and $\hat{\mathbf{G}}''(\mathbf{t})$: in the wave-space computations, the content in the square bracket on the right hand side of Eqs. (50) and (51) is replaced with $\hat{\mathbf{H}}_q$ in Eq. (42). We adopted this hybrid approach in this work to compute the HIs, and discuss the accuracy of various approaches in Sec. 5.1.

3.3. Real-space computation

The real-space contributions to the grand mobility tensor \mathfrak{M} are computed pairwise using the formalism in Sec. 2. Since $\mathbf{J}_R(\mathbf{r})$ [Eq. (22)] decays exponentially fast with distance, when the parameter ξ is sufficiently large, only particle pairs within a cutoff distance r_c need to be evaluated. Introducing the cutoff radius r_c for pair evaluation allows fast neighbor searching algorithms such as the linked list [69] or the chaining mesh [70] method to be used. These methods divide the simulation box into cells of size slightly larger than r_c , and sort the particles into the cells. To find the neighbors of a particle, only particles in the residing cell and its 26 neighboring cells need to be searched. This effectively improves the operation count to $\mathcal{O}(N \log N)$ for the real-space computations.

To accommodate the iterative scheme for HI computations in Sec. 4, the real-space grand mobility tensor is constructed as a sparse matrix at each time step. After the matrix construction, the action of the real-space contributions to \mathfrak{M} is simply a matrix-vector multiplication. Otherwise, neighbor searching and pair HI evaluations need to be carried out at every iteration. Note that we also include the self-contributions from the wave-space computations, *e.g.*, Eq. (37), and the self-part of the pressure Faxén law [Eq. (28)], in the real-space grand mobility tensor.

3.4. GPGPU acceleration of the mobility computation

The mobility computation with the SE method was first implemented on CPU and the performance was unsatisfactory for dynamic simulations. The bottlenecks are the force interpolation step and the convolution step. These are common speed limiting steps in particle mesh techniques due to ineffective memory caching between the particle and the grid data. For polydisperse systems in this work, the situation is aggravated as more interpolation points P are needed for satisfactory HI resolution. After a few optimization iterations on CPU, we realized that the key to the performance is the memory bandwidths. Since modern GPUs typically have significantly higher memory bandwidths compared to CPUs, in this work the entire mobility computation is carried out on GPU using CUDA C, a popular GPGPU programming model with a relatively mature environment for scientific computations.

The GPU mobility computations are carried out in Single Precision (SP) for the highest GPU performance. The cost of the performance in SP computation is the accuracy, as the SP arithmetics can be severely limited by the number

of significant digits compared to the Double Precision (DP). However, this is not a problem in this work for at least three reasons: (i) For dynamic simulations with iterative solvers, the SP accuracy is often sufficient; (ii) The SE method is able to reach the round-off error of the SP arithmetics with proper parameter selection due to its spectral accuracy; and (iii) The far-field HIs captured by the mobility computations are smooth compared to the lubrication interactions, which are evaluated in DP on CPUs. The split of the near- and far-field HIs in SD allows a natural mixed precision HI computation that captures the most significant contributions from each part.

The GPGPU computations exploit the massively parallel structure of modern GPUs by simultaneously executing a large number of similar tasks, or threads, on the data. To maintain performance, data dependencies and communications between threads should be minimized. This makes the GPU implementation of the SE method different from its CPU counterpart. Inspired by earlier GPU implementations of particle mesh techniques, this work combines the grid-based method of Ganesan *et al.* [47] for force interpolation and the particle-based approach of Harvey & De Fabritiis [48] for convolution. The grid-based force interpolation keeps a list of contributing particles for each grid point, and the list is updated when the particle configurations are changed. The grid values are computed in parallel using M^3 threads: with the particle list, each thread sums the force, torque, and stresslet contributions independently for each grid point. On the other hand, the particle-based convolution is a weighted summation on P^3 grid points for each particle. To maximize parallelization, the summation for each particle is performed by a group of P threads cooperatively. Each thread in the group first sums P^2 grid points on the transverse plane, and for the final result, the first thread in the group adds up the values from other threads using the shared memory of the GPU. Moreover, on the GPU we use the `cuffft` package for the FFTs and the `cusparse` package for the sparse matrix-vector multiplication.

4. Dynamic simulation with Stokesian Dynamics

The framework of SD [7, 63] approximates the projected grand resistance tensor \mathcal{R} in Eq. (12) as

$$\mathcal{R} = \mathfrak{M}^{-1} + \mathcal{R}^{\text{nf}}, \quad (53)$$

where \mathfrak{M} is the multipole grand mobility tensor, and \mathcal{R}^{nf} is the pairwise additive lubrication correction without the far-field contributions. Recall that the inversion of \mathfrak{M} captures the many-body aspect of HIs, and the short-range correction \mathcal{R}^{nf} captures the lubrication effects. The SD recovers the exact result for two-body problems and agrees well with the exact solutions of three-body problems [71]. It can provide significant insights to the HIs of dense suspensions [72, 73].

4.1. Iterative computation of hydrodynamic interactions

We incorporate the SE mobility computation into the framework of SD using the iterative scheme of Swan & Brady [62], and call the resulting method the Spectral Ewald Accelerated Stokesian Dynamics (SEASD). Here, a matrix-free iterative scheme is necessary as the grand mobility tensor \mathfrak{M} is not explicitly constructed. The iterative scheme splits the overall hydrodynamic force,

$$\mathcal{F}^{\text{H}} = -\mathbf{R}_{\mathcal{F}\mathcal{U}} \cdot \mathcal{U}^{\text{H}} + \mathbf{R}_{\mathcal{F}\mathbf{E}} \cdot \mathbf{E}^{\infty}, \quad (54)$$

where \mathcal{U}^{H} is the velocity disturbances due to HIs, into a near-field part and a far-field part. The near-field part satisfies

$$0 = -\mathbf{R}_{\mathcal{F}\mathcal{U}}^{\text{nf}} \cdot \mathcal{U}^{\text{H}} + \mathcal{F}^{\text{H,ff}} + \tilde{\mathcal{F}}^{\text{P}}, \quad (55)$$

where $\mathbf{R}_{\mathcal{F}\mathcal{U}}^{\text{nf}}$ is the $\mathcal{F}\mathcal{U}$ coupling in \mathcal{R}^{nf} and is stored as a sparse matrix, $\tilde{\mathcal{F}}^{\text{P}} = \mathcal{F}^{\text{P}} + \mathbf{R}_{\mathcal{F}\mathbf{E}}^{\text{nf}} \cdot \mathbf{E}^{\infty}$ contains the interparticle force \mathcal{F}^{P} and the near-field contributions from \mathbf{E}^{∞} . The far-field hydrodynamic force $\mathcal{F}^{\text{H,ff}}$ satisfies

$$\begin{bmatrix} \mathcal{U}^{\text{H}} \\ -\mathbf{E}^{\infty} \end{bmatrix} = -\mathfrak{M} \cdot \begin{bmatrix} \mathcal{F}^{\text{H,ff}} \\ \mathbf{S}^{\text{H,ff}} \end{bmatrix}, \quad (56)$$

where $\mathbf{S}^{\text{H,ff}}$ is the far-field stresslet from HIs. Solving Eqs. (55) and (56), the far-field hydrodynamic forces and stresslets are

$$\begin{bmatrix} \tilde{\mathcal{F}}^{\text{H,ff}} \\ \mathbf{S}^{\text{H,ff}} \end{bmatrix} = \tilde{\mathfrak{M}}^{-1} \cdot \left((\lambda \mathfrak{M} - \mathcal{I}) \cdot \begin{bmatrix} (\tilde{\mathbf{R}}_{\mathcal{F}\mathcal{U}}^{\text{nf}})^{-1} \cdot \tilde{\mathcal{F}}^{\text{P}} \\ \mathbf{0} \end{bmatrix} + \begin{bmatrix} \mathbf{0} \\ \mathbf{E}^{\infty} \end{bmatrix} \right), \quad (57)$$

where

$$\widetilde{\mathfrak{M}} = \left((\mathcal{I} - \lambda \mathfrak{M}) \cdot \begin{bmatrix} (\widetilde{\mathbf{R}}_{\mathcal{F}\mathcal{U}}^{\text{nf}})^{-1} & \mathbf{0} \\ \mathbf{0} & \mathfrak{M} \end{bmatrix} + \mathfrak{M} \right). \quad (58)$$

To ensure invertibility, a diagonal matrix $\lambda \mathbf{I}$, with λ a parameter, is added to $\mathbf{R}_{\mathcal{F}\mathcal{U}}^{\text{nf}}$, *i.e.*, $\widetilde{\mathbf{R}}_{\mathcal{F}\mathcal{U}}^{\text{nf}} = \mathbf{R}_{\mathcal{F}\mathcal{U}}^{\text{nf}} + \lambda \mathbf{I}$, and accordingly $\widetilde{\mathcal{F}}^{\text{H,ff}} = \mathcal{F}^{\text{H,ff}} + \lambda \mathcal{U}^{\text{H}}$. A convenient choice for λ is $6\pi\eta_0 a$, where a is the reference particle radius [62].

Solving Eq. (57) requires nested iteration as each evaluation of $\widetilde{\mathfrak{M}}$ contains the solution of the near-field problem with $\widetilde{\mathbf{R}}_{\mathcal{F}\mathcal{U}}^{\text{nf}}$. The near-field problem is efficiently solved by the Generalized Minimum Residual (GMRES) method with an Incomplete Cholesky preconditioner with zero fill-in (IC0) [74]. To reduce the IC0 breakdown, prior to applying the preconditioner particles are reordered using the reverse Cuthill-McKee algorithm. For isotropic suspensions, the near-field problem typically converges to an error of 10^{-4} within 10 iterations [36]. For suspensions with strong structural anisotropy, however, the convergence becomes more difficult and the IC0 preconditioner breaks down even with the reordering. This is resolved by increasing λ in $\widetilde{\mathbf{R}}_{\mathcal{F}\mathcal{U}}^{\text{nf}}$, or introducing a threshold value λ_{IC} in during the IC0 preconditioner computation [74]. Increasing λ in $\widetilde{\mathbf{R}}_{\mathcal{F}\mathcal{U}}^{\text{nf}}$ does not change the convergence of the near-field problem, but increases the number of expensive $\widetilde{\mathfrak{M}}$ iterations. On the other hand, increasing λ_{IC} deteriorates the quality of the IC0 preconditioner and increases the iterations required for the near-field problem, but has little effect on the far-field evaluations. In dynamic simulations, both λ and λ_{IC} are adjusted for optimal computation efficiency.

The pressure moment computation in SEASD also follows the near- and far-field splitting scheme in Eqs. (55) and (56). Due to the special coupling between the pressure moments and other force moments in compressible suspensions (Sec. 2.3), the interaction contribution to the far-field pressure moment is evaluated after $\mathbf{F}^{\text{H,ff}}$ and the traceless part of $\mathbf{S}^{\text{H,ff}}$ are solved in Eq. (57). On the other hand, the near-field part of the pressure moment is evaluated along with other parts of the stresslets using the near-field resistance functions.

The near-field pairwise lubrication corrections \mathcal{R}^{nf} are based on the exact solutions of two-body problems in series form [64, 65, 75, 76] up to s^{-300} , where $s = 2r/(a_i + a_j)$, with a_i and a_j the radii of the pair, is the scaled particle center-center distance. In the simulations, the lubrication corrections are activated when $s < 4$: for $s > 2.1$ the interpolation of tabulated data and for $s \leq 2.1$ the analytical expressions are used. Note that \mathcal{R}^{nf} constructed from two-body problems contains both the relative and the collective motions of the particle pair and, as pointed out by Cichocki *et al.* [23], the lubrication corrections corresponding to the collective motion can destroy the far-field asymptotics beyond the pair level. However, for dense suspensions, this only leads to a minor quantitative difference on the suspension static properties [11] in conventional SD. Therefore, we retain the full lubrication correction here for consistency with the existing SD framework. The SD implementations of Ando & Skolnick [13] removed the pair collective motion in the lubrication corrections.

4.2. Far-field preconditioner

Here we introduce a preconditioner for $\widetilde{\mathfrak{M}}$ to reduce the number of expensive far-field mobility evaluations when solving Eq. (57). Since \mathfrak{M} is not explicitly constructed, the preconditioner needs to be built from a suitable approximation. For mobility problems without the lubrication corrections, Saintillan *et al.* [37] and Keaveny [44] found substantial iteration improvement even with the diagonal mobility approximation. Unfortunately, the approximation of $\widetilde{\mathfrak{M}}$ is more involved due to the presence of $(\widetilde{\mathbf{R}}_{\mathcal{F}\mathcal{U}}^{\text{nf}})^{-1}$. In this work, a block diagonal approximation of $\widetilde{\mathfrak{M}}$ for the far-field preconditioner is adopted. First, the near-field resistance tensor $\widetilde{\mathbf{R}}_{\mathcal{F}\mathcal{U}}^{\text{nf}}$ is approximated by N blocks of 6×6 submatrices along its diagonal. Using the direct sum notation, this is $\bigoplus_{i=1}^N (\widetilde{\mathbf{R}}_{\mathcal{F}\mathcal{U}}^{\text{nf}})_{ii}$, where \bigoplus is the direct sum, and $(\widetilde{\mathbf{R}}_{\mathcal{F}\mathcal{U}}^{\text{nf}})_{ij}$ is the block submatrix between particles i and j in $\widetilde{\mathbf{R}}_{\mathcal{F}\mathcal{U}}^{\text{nf}}$. To approximate $\widetilde{\mathfrak{M}}$, we use

$$(\widetilde{\mathbf{R}}_{\mathcal{F}\mathcal{U}}^{\text{nf}})^{-1} \approx \bigoplus_{i=1}^N [(\widetilde{\mathbf{R}}_{\mathcal{F}\mathcal{U}}^{\text{nf}})_{ii}]^{-1}, \quad (59)$$

which only involves N inversion of 6×6 matrices. The mobility tensor \mathfrak{M} is approximated by its block-diagonal components using direct Ewald summation, *i.e.*, for each particle, the approximation only considers the interactions with its periodic images. To obtain the preconditioner, we apply the Incomplete LU decomposition with zero fill-in (ILU0) [74] on the approximated $\widetilde{\mathfrak{M}}$, which is constructed following Eq. (58) with the approximated $(\widetilde{\mathbf{R}}_{\mathcal{F}\mathcal{U}}^{\text{nf}})^{-1}$ and \mathfrak{M} .

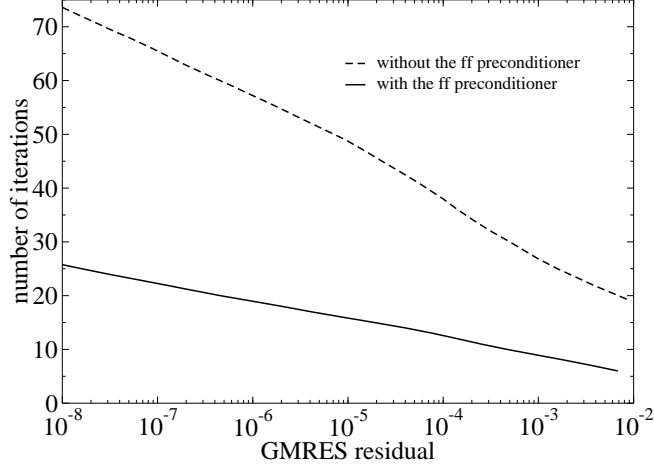


Figure 1: The number of far-field iterations, *i.e.*, the number of the grand mobility tensor $\widetilde{\mathfrak{M}}$ evaluations, as a function of the GMRES residual with (solid line) and without (dashed line) the far-field preconditioner for a bidisperse suspension of $N = 200$, $\lambda = 2$, $x_2 = 0.3$, and $\phi = 0.2$.

Unlike Saintillan *et al.* [37], including close pair interactions has an adverse effect on the preconditioner due to the diagonal approximation of $\widetilde{\mathbf{R}}_{\mathcal{F}\mathcal{U}}^{\text{nf}}$.

The effectiveness of this preconditioner on the far-field iteration is demonstrated in Fig. 1. In this case, the HIs corresponding to random forces and strain rates are solved for a random bidisperse suspension of 200 particles with $\lambda = 2$, $x_2 = 0.3$, and $\phi = 0.2$. The far-field preconditioner substantially reduces the number of GMRES iterations. Evidently, its usage is justified when the required GMRES residual is small, since constructing the approximate $\widetilde{\mathfrak{M}}$ and the ILU0 decomposition also take time. In dynamic simulations, further time saving can be achieved by updating the preconditioner every few time steps. In addition, the exact break-even time also depends on the far-field mobility computation parameters, including M , P , and r_c that indirectly affect the iterative solver. Finally, since the preconditioner construction is an $O(N)$ operation and the $\widetilde{\mathfrak{M}}$ evaluation scales as $O(N \log N)$, preconditioning is almost always justified for large systems.

4.3. Dynamics simulation of Brownian suspensions

Particle dynamics in a suspension are described by the generalized N -body Langevin equation,

$$\mathbf{m} \cdot \frac{d\mathcal{U}}{dt} = \mathcal{F}^{\text{H}} + \mathcal{F}^{\text{P}} + \mathcal{F}^{\text{B}} \quad (60)$$

where \mathbf{m} is the generalized mass/moment of inertial matrix, \mathcal{U} is the generalized particle velocity and \mathcal{F}^{H} , \mathcal{F}^{P} , and \mathcal{F}^{B} are the forces on particles. The hydrodynamic force \mathcal{F}^{H} arises from the HIs and can be computed from Eq. (54). The interparticle force \mathcal{F}^{P} originates from the interparticle potentials. The Brownian force \mathcal{F}^{B} is due to thermal fluctuations in the solvent, and from the fluctuation-dissipation theorem [77], \mathcal{F}^{B} satisfies

$$\overline{\mathcal{F}^{\text{B}}(t)} = 0 \text{ and } \overline{\mathcal{F}^{\text{B}}(0)\mathcal{F}^{\text{B}}(t)} = 2k_{\text{B}}T\delta(t)\mathbf{R}_{\mathcal{F}\mathcal{U}}, \quad (61)$$

where the overline denotes an average over the solvent fluctuations and $k_{\text{B}}T$ is the thermal energy scale.

The configuration evolution is obtained by integrating Eq. (60) twice over an appropriate time scale Δt , and the result is [38, 78]

$$\Delta X = \left[\mathcal{U}^{\infty} + \mathbf{R}_{\mathcal{F}\mathcal{U}}^{-1} \cdot (\mathbf{R}_{\mathcal{F}\text{E}} \cdot \mathbf{E}^{\infty} + \mathcal{F}^{\text{P}}) \right] \Delta t + k_{\text{B}}T \nabla \cdot \mathbf{R}_{\mathcal{F}\mathcal{U}}^{-1} \Delta t + \Delta X^{\text{B}}, \quad (62)$$

where ΔX is the suspension configuration change over time Δt , \mathcal{U}^{∞} is the generalized velocity from the imposed flow, and ΔX^{B} is the Brownian displacement which satisfies

$$\overline{\Delta X^{\text{B}}} = 0 \text{ and } \overline{\Delta X^{\text{B}}\Delta X^{\text{B}}} = 2k_{\text{B}}T\Delta t\mathbf{R}_{\mathcal{F}\mathcal{U}}^{-1}. \quad (63)$$

The second term on the right hand side of Eq. (62) is the deterministic drift due to the configuration dependent Brownian force \mathcal{F}^B , and the divergence operator is acting on the last index of $\mathbf{R}_{\mathcal{F}\mathcal{U}}^{-1}$. The divergence can be numerically evaluated following Banchio & Brady [39].

The suspension bulk stress is obtained by spatially averaging the Cauchy stress [50, 51], *i.e.*,

$$\langle \boldsymbol{\Sigma} \rangle = -\langle p \rangle_f \mathbf{I} + 2\eta_0 \langle \overline{\mathbf{E}^\infty} \rangle + (\kappa_0 - \frac{2}{3}\eta_0) E^\infty \mathbf{I} - nk_B T \mathbf{I} + n(\langle \mathbf{S}^E \rangle + \langle \mathbf{S}^P \rangle + \langle \mathbf{S}^B \rangle), \quad (64)$$

where $\langle p \rangle_f$ is the average solvent pressure, $\langle \cdot \rangle$ is the volume average over the entire suspension, κ_0 is the fluid bulk viscosity, and n is the particle number density. The particle stresslets \mathbf{S}^H are broken down as $\mathbf{S}^H = \mathbf{S}^E + \mathbf{S}^P + \mathbf{S}^B$, where \mathbf{S}^E is the contributions from the the imposed flow, \mathbf{S}^P from the interparticle potential, and \mathbf{S}^B from the Brownian motion. Their suspension averages are expressed in resistance tensors

$$\langle \mathbf{S}^E \rangle = -\langle \mathbf{R}_{S\mathcal{U}} \cdot \mathbf{R}_{\mathcal{F}\mathcal{U}}^{-1} \cdot \mathbf{R}_{\mathcal{F}E} - \mathbf{R}_{SE} \rangle, \quad (65)$$

$$\langle \mathbf{S}^P \rangle = -\langle (\mathbf{R}_{S\mathcal{U}} \cdot \mathbf{R}_{\mathcal{F}\mathcal{U}}^{-1} + \mathbf{r}\mathbf{I}) \cdot \mathbf{F}^P \rangle, \quad (66)$$

$$\langle \mathbf{S}^B \rangle = -k_B T \langle \nabla \cdot (\mathbf{R}_{S\mathcal{U}} \cdot \mathbf{R}_{\mathcal{F}\mathcal{U}}^{-1}) \rangle, \quad (67)$$

where the divergence in Eq. (67) is applied to the last index in the bracket. For hard-sphere suspensions, $\langle \mathbf{S}^P \rangle = 0$ as the HI and the interparticle force contributions exactly cancel each other [51]. The Brownian stresslet $\langle \mathbf{S}^B \rangle$ can also be computed using the modified mid-point scheme [39].

In dynamic simulations, the Brownian displacement ΔX^B is evaluated from the Brownian force \mathcal{F}^B in Eq. (61) as

$$\Delta X^B = \mathbf{R}_{\mathcal{F}\mathcal{U}}^{-1} \cdot \mathcal{F}^B \Delta t. \quad (68)$$

Following Banchio & Brady [39], the Brownian force can be split into a near-field part and a far-field part,

$$\mathcal{F}^B = \mathcal{F}^{B,\text{nf}} + \mathcal{F}^{B,\text{ff}}. \quad (69)$$

Both $\mathcal{F}^{B,\text{nf}}$ and $\mathcal{F}^{B,\text{ff}}$ have zero mean and satisfy

$$\overline{\mathcal{F}^{B,\text{nf}} \mathcal{F}^{B,\text{nf}}} = \frac{2k_B T}{\Delta t} \mathbf{R}_{\mathcal{F}\mathcal{U}}^{\text{nf}}, \quad (70)$$

$$\overline{\mathcal{F}^{B,\text{ff}} \mathcal{F}^{B,\text{ff}}} = \frac{2k_B T}{\Delta t} (\mathfrak{M}^{-1})_{\mathcal{F}\mathcal{U}}, \quad (71)$$

$$\overline{\mathcal{F}^{B,\text{ff}} \mathcal{F}^{B,\text{nf}}} = 0, \quad (72)$$

where $(\mathfrak{M}^{-1})_{\mathcal{F}\mathcal{U}}$ is the $\mathcal{F}\mathcal{U}$ block of the inverted far-field grand mobility tensor. The pairwise-additive lubrication corrections allow pairwise evaluation of the near-field Brownian force $\mathcal{F}^{B,\text{nf}}$ [39]. Since \mathfrak{M} is not explicitly constructed, to compute $\mathcal{F}^{B,\text{ff}}$, it is necessary to solve

$$\begin{bmatrix} \mathcal{F}^{B,\text{ff}} \\ \Delta \mathbf{S}^B \end{bmatrix} = \frac{2k_B T}{\Delta t} (\mathfrak{M}^{-1/2}) \cdot \boldsymbol{\Psi}, \quad (73)$$

where $\boldsymbol{\Psi}$ is a Gaussian noise of zero mean and unit variance, and $\Delta \mathbf{S}^B$ is the fluctuation part of the Brownian stress in Eq. (67). The inverse square root of the grand mobility tensor $\mathfrak{M}^{-1/2}$ in Eq. (73) can be approximated using Chebychev polynomials with eigenvalue estimations [39, 79], or solved as an Initial Value Problem (IVP) [40, 80], which was first used by Swan & Brady [40] in ASD. The solution of the following IVP [81] with matrix \mathbf{A} ,

$$\frac{d\mathbf{x}}{d\tau} = -\frac{1}{2} [\tau \mathbf{I} + (1 - \tau) \mathbf{A}]^{-1} \cdot (\mathbf{A} - \mathbf{I}) \cdot \mathbf{x}, \quad \mathbf{x}(0) = \mathbf{c}, \quad (74)$$

at $\tau = 1$ satisfies $\mathbf{x}(1) = \mathbf{A}^{-1/2} \cdot \mathbf{c}$. Swan & Brady [40] devised a numerical scheme to solve Eq. (74) in ASD: at each time step with step size $\Delta\tau$, Eq. (74) is marched first with a Euler forward half-step then a Euler backward half-step, *i.e.*,

$$\frac{\mathbf{x}_{i+\frac{1}{2}} - \mathbf{x}_i}{\Delta\tau/2} = -\frac{1}{2} [\tau_i \mathbf{I} + (1 - \tau_i) \mathbf{A}]^{-1} \cdot (\mathbf{A} - \mathbf{I}) \cdot \mathbf{x}_i, \quad (75)$$

$$\frac{\mathbf{x}_{i+1} - \mathbf{x}_{i+\frac{1}{2}}}{\Delta\tau/2} = -\frac{1}{2} [\tau_{i+1}\mathbf{I} + (1 - \tau_{i+1})\mathbf{A}]^{-1} \cdot (\mathbf{A} - \mathbf{I}) \cdot \mathbf{x}_{i+1}. \quad (76)$$

With $\mathbf{A} = \mathfrak{M}$ and $\mathbf{c} = (2k_B T / \Delta t) \Psi$, Eq. (73) is solved at $\tau = 1$. In SEASD, both Eqs. (75) and (76) are solved iteratively, usually with a smaller tolerance compared to $\Delta\tau$. The results with $\Delta\tau = 0.1$ are often satisfactory.

For dynamic simulation of Brownian suspensions under a simple shear flow with strain rate $\dot{\gamma}$, the ratio of the convective transport rate $\dot{\gamma}$ and the diffusive transport rate $k_B T / (6\pi\eta_0 a_p^3)$ defines the Péclet number,

$$\text{Pe} = \frac{6\pi\eta_0 a_p^3 \dot{\gamma}}{k_B T}. \quad (77)$$

Small Pe indicates Brownian motion dominance, and large values suggest negligible Brownian influences. For bidisperse suspensions, we define Pe based on the size of the small particles to capture the dynamics of the most rapid changes, *i.e.*, $a_p = a_1$. In dynamic simulations, the time in Eq. (62) is scaled according to the Péclet number: when $\text{Pe} \leq 1$, it is scaled with the diffusive time scale of the small particles, $6\pi\eta_0 a_1^3 / (k_B T)$, and when $\text{Pe} > 1$, the convective time scale $\dot{\gamma}^{-1}$.

4.4. The mean-field Brownian approximation

The most time-consuming step in dynamic simulations of Brownian suspensions is computing $\mathcal{F}^{\text{B,ff}}$ from Eq. (73) due to the large number of \mathfrak{M} evaluations, although the IVP approach in Sec. 4.3 is expected to be faster than the Chebychev approximation [40]. Further speed improvement is possible by introducing a mean-field approximation of the Brownian-related quantities [39]. In this approach, the far-field grand mobility tensor \mathfrak{M} is approximated as a diagonal matrix for all Brownian related computations, and the full HI computations are retained for the flow-related quantities such as \mathbf{S}^{E} . As a result, this method retains the $\mathcal{O}(N \log N)$ scaling, but with an order of magnitude smaller prefactor for monodisperse suspensions [39]. The diagonal approximation of \mathfrak{M} uses the single particle result for the ES coupling, and the far-field translational and rotational short-time self-diffusivities for the \mathcal{UF} coupling. These far-field values are from Monte-Carlo computations of equilibrium configurations at the same volume fraction *without* the lubrication corrections. Extending this approach to polydisperse suspensions is trivial: the suspension far-field diffusivities in the diagonal elements are replaced by the far-field diffusivities for each species. The mean-field Brownian approximation is especially suitable for studying dense suspension rheology, where the HIs are dominated by the near-field lubrication interactions. Following Brady & Banchio [39], we designate this approximation scheme SEASD-nf.

5. Accuracy and performance

5.1. Mobility computation accuracy

The accuracy of the mobility computation is characterized by the relative ∞ -norm of the strain rate, *i.e.*,

$$e_{\infty,r}(E) = \max_{i \in \{1, \dots, N\}} \frac{\|\mathbf{E}_i^{\text{SE}} - \mathbf{E}_i^*\|}{\|\mathbf{E}_i^*\|}, \quad (78)$$

where the \mathbf{E}_i^{SE} is the particle strain rate from the SE method and \mathbf{E}_i^* is a well-converged value from direct Ewald summation. Other error measurements can be similarly defined. For example, $e_{\infty,r}(U)$ for the linear velocity was used by Lindbo & Tornberg [29] to characterize the accuracy of the SE method for point forces. For the stresslet-strain rate level mobility computation here, we found $e_{\infty,r}(E)$ the most stringent error criteria, possibly because more derivatives are involved in Eq. (15).

To facilitate quantitative discussions, in this section we focus on a random bidisperse hard-sphere system of $N = 50$, $\phi = 0.05$, $\lambda = 2$, and $x_2 = 0.3$. The imposed force, torque, and stresslet on each particle are randomly drawn from a normal distribution, and rescaled to ensure $\|\mathbf{F}_i\| = 1$, $\|\mathbf{T}_i\| = 1$, and $\|\mathbf{S}_i\| = 1$. The simulation box lattice vectors are $\mathbf{a}_1 = (L, 0, 0)$, $\mathbf{a}_2 = (\gamma L, L, 0)$, and $\mathbf{a}_3 = (0, 0, L)$, with γ the strain. The computations are carried out in DP accuracy on CPU.




Figure 2: The wave-space accuracy measured by $e_{\infty,r}(E)$ [Eq. (78)] as a function of the interpolation point P with various shape parameter m at $M = 64$ and $\xi a_1 = 0.1$. The particle size effects are incorporated using (a): the real-space, (b): the hybrid, and (c): the wave-space approaches in Sec. 3.2. The values of m are annotated in each figure. The solid and dashed lines represent the case of $\gamma = 0$ and 0.5 , respectively. The dashed dotted lines show the exponential minimum error decay, $e_{\infty,r}(E) \sim \exp(-P\pi/2)$.

5.1.1. Wave-space accuracy

Fig. 2 presents the accuracy of wave-space computation using different SE implementations with orthogonal ($\gamma = 0$) and sheared ($\gamma = 0.5$) simulation boxes in solid and dashed lines, respectively. The error $e_{\infty,r}(E)$ is shown as a function of the interpolation point P with various shape parameter m at $M = 64$ and $\xi a_1 = 0.1$. Different particle size incorporation approaches discussed in Sec. 3.2 are presented: in Fig. 2a the real-space approach, in Fig. 2b the hybrid approach, and in Fig. 2c the wave-space approach.

There are several key observations in Fig. 2. First of all, the errors associated with orthogonal and sheared simulation boxes are almost identical. This validates the general formalism for non-orthogonal simulation boxes in Sec. 3.1. Secondly, the SE method is sensitive to P and m , which respectively correspond to the discretization and truncation of the shape function $h(t)$. At a given m , $e_{\infty,r}(E)$ first decreases exponentially, followed by a much slower reduction with increasing P . The two-stage reduction of $e_{\infty,r}(E)$ is well understood for point forces [29]: the exponential decrease is due to the improved resolution of the shape function, and the slower reduction is associated with the Gaussian truncation from the shape parameter m . Therefore, at large P and m the result is expected to be accurate; indeed, in Fig. 2 the minimum errors are all close to the machine precision. Such accuracy is inaccessible using the PME or the SPME method at this grid number ($M = 64$) due to the inherent coupling between the interpolation and the wave-space truncation errors. Moreover, for a given P , $e_{\infty,r}(E)$ first decreases to a minimum and then increases with increasing m . At the minimum, $e_{\infty,r}(E)$ is transitioning from exponential to slower decay, and the errors from the shape resolution is about the same as the errors from the Gaussian truncation. From the error estimation of Lindbo & Tornberg [29, 30], at a given P , the minimum wave-space error $e_{\infty,r}(E)$ and the corresponding shape parameter m are

$$e_{\infty,r}(E) \sim \exp(-P\pi/2) \text{ and } m \sim \sqrt{\pi P}, \quad (79)$$

respectively. The asymptotic exponential decay of the minimum $e_{\infty,r}(E)$ is also shown as dash-dotted lines in Fig. 2. The exponential decay of the minimum error with respect to P to the round-off precision at large P and m clearly demonstrate the spectral accuracy [82] of the SE method.

In Fig. 2 different particle size incorporation approaches exhibit similar qualitative behaviors with quantitative differences. For example, to achieve an accuracy of $e_{\infty,r}(E) \sim 10^{-4}$ at the optimal m , in Fig. 2a, 2b, and 2c the required P are respectively 15, 13, and 9, corresponding to the real-space, hybrid, and wave-space approaches discussed in Sec. 3.2. The latter two approaches reduce the $h(t)$ evaluations by 35% and 78% compared to the real-space approach at a cost of the number of required FFTs. Therefore, there is a subtle balance between the number of interpolation points P and the number of FFTs in the SE method implementation. The hybrid approach in Fig. 2b achieves a good balance between accuracy and computation efficiency, and therefore is adopted in SEASD.

Finally, Fig. 2 shows that, in addition to the spectral accuracy and the ease of implementation, the SE method also allows flexible error control by adjusting P and m without changing the grid points M . As a result, the errors from

the wave-space summation and the interpolation can be separated, and this permits more flexible error control when computing HIs in polydisperse systems. On the other hand, such error separation is not possible in other particle mesh techniques such as the PME and the SPME methods.

5.1.2. Overall mobility accuracy

Both the wave-space and the real-space computations affect the overall mobility accuracy, and the controlling parameters are the grid point M , the interpolation point P , the Gaussian shape parameter m , the real-space cutoff radius r_c , and the splitting parameter ξ . Out of the five parameters, only changes in ξ and m do not affect the computational cost since adjusting M affects the FFT size, changing r_c influences the neighbor search, *etc.* With fixed computation cost, *i.e.*, fixed M , P , and r_c , it is desirable to find the combination of m and ξ that minimizes the overall error.

Fig. 3 and 4 present the effects of m and ξ on the overall mobility accuracy with various P and r_c for $M = 64$ and 32, respectively. The wave-space computation uses the hybrid approach in Sec. 3.2, and the simulation box is orthogonal ($\gamma = 0$). The thick black lines in these figures indicate the theoretical optimal shape parameter $m = \sqrt{\pi P}$ [29, 30]. Note that in our implementation, the cutoff radius r_c depends on the radius a_i and a_j in a particle pair.

Fig. 3i with $M = 64$, $P = 21$, and $r_c = 6(a_i + a_j)$ best illustrates the influences of m and ξ . Here, the mobility computation can reach $e_{\infty,r}(E) < 10^{-9}$ at $(\xi a_1, m) = (0.46, 8)$. With fixed m , $e_{\infty,r}(E)$ exhibits a minimum with increasing ξa_1 , and when $m \leq 8$, the minimum degenerates to a plateau due to the wave-space Gaussian truncation, which is also illustrated in Fig. 2 at low m . At low ξ , the overall error is dominated by the real-space error, which decreases with increasing ξ . At high ξ , the overall error is mainly from the wave space, and increases with increasing ξ . With fixed ξ on the other hand, $e_{\infty,r}(E)$ also shows a minimum with increasing m . When $\xi a_1 \leq 0.46$, the $e_{\infty,r}(E)$ minimum becomes a plateau since the real-space error is independent of m . Here, the reduction of $e_{\infty,r}(E)$ with increasing m at small m comes almost entirely from the reduced Gaussian truncation. When $\xi a_1 > 0.46$, the minimum plateau disappears as in this region the wave-space error is sensitive to m , a point also illustrated in Fig. 2. Furthermore, in Fig. 3i there is a region of $e_{\infty,r}(E) > 1$ at high ξ and low m due to large wave-space errors.

Comparison across rows and columns in Fig. 3 and 4 reveals the influences of r_c and P on the overall accuracy, respectively. For both cases, reducing r_c or P increases the minimum value of $e_{\infty,r}(E)$ and changes the corresponding ξa_1 and m . Comparing Fig. 3g, 3h, and 3i shows that reducing r_c increases the real-space error and shifting the minimum of $e_{\infty,r}(E)$ towards larger ξa_1 . The decrease of $e_{\infty,r}(E)$ with respect to increasing ξ at small ξa_1 also becomes slower. In Fig. 3g, the $e_{\infty,r}(E)$ minimum is at $\xi a_1 > 1$. Comparing Fig. 3i, 3f, and 3c reveals the effects of reducing the interpolation point P . With diminishing P , the wave-space error increases due to poor Gaussian resolution, and the $e_{\infty,r}(E)$ minimum is shifted towards lower m . In addition, the overall accuracy decreases significantly for large m at small P , *e.g.*, in Fig. 3c, $e_{\infty,r}(E) > 1$ when $m > 8$.

Comparing Fig. 3 and 4 shows the effect of grid point M on the mobility accuracy. Note that the color scales in Fig. 3 and 4 are different, and the minimum $e_{\infty,r}(E)$ in Fig. 3f and 4f is approximately the same. The most apparent effect of reducing M is the shrinkage of the parameter space corresponding to $e_{\infty,r}(E) < 1$ due to the truncation of the wave-space sum. As a result, at $M = 32$, the mobility evaluation is more sensitive to ξa_1 compared to the case of $M = 64$. Otherwise, the qualitative aspects of Fig. 4 are similar to Fig. 3. Moreover, the thick black lines representing the theoretical optimal shape parameter $m = \sqrt{\pi P}$ is almost always in the vicinity of the regions of the highest accuracy in both Fig. 3 and 4. This substantially simplifies the search for the optimal ξ .

The influences of the particle number N on the overall mobility accuracy is presented in Fig. 5 for $M = 32$ and 64. The simulation box size is fixed at $L/a_1 = 23.5$ in Fig. 5a, and the suspension volume fraction is fixed at $\phi = 0.05$ in Fig. 5b. Other parameters remain unchanged from the baseline case, and the mobility computation parameters are $P = 13$, $m = 6.7$, and $r_c = 4(a_i + a_j)$. The mobility accuracy is more sensitive to changes in L than changes in ϕ . In Fig. 5a, $e_{\infty,r}(E)$ changes little, but in Fig. 5b, the $e_{\infty,r}(E)$ minimum increases drastically with different N . The almost identical decrease in $e_{\infty,r}(E)$ at small ξa_1 suggests the real-space error are not significantly changed by N in either case. The diverging $e_{\infty,r}(E)$ at higher ξa_1 in Fig. 5b suggests the wave-space computation is sensitive to the box size at fixed P and m . This is well-known for particle mesh techniques in general [29, 67]. Therefore, to retain the computational accuracy with larger systems at the same volume fraction, it is necessary to increase the grid point M or the interpolation point P . Finally, we note in passing that the same qualitative error behaviors are found in the pressure moment computations.

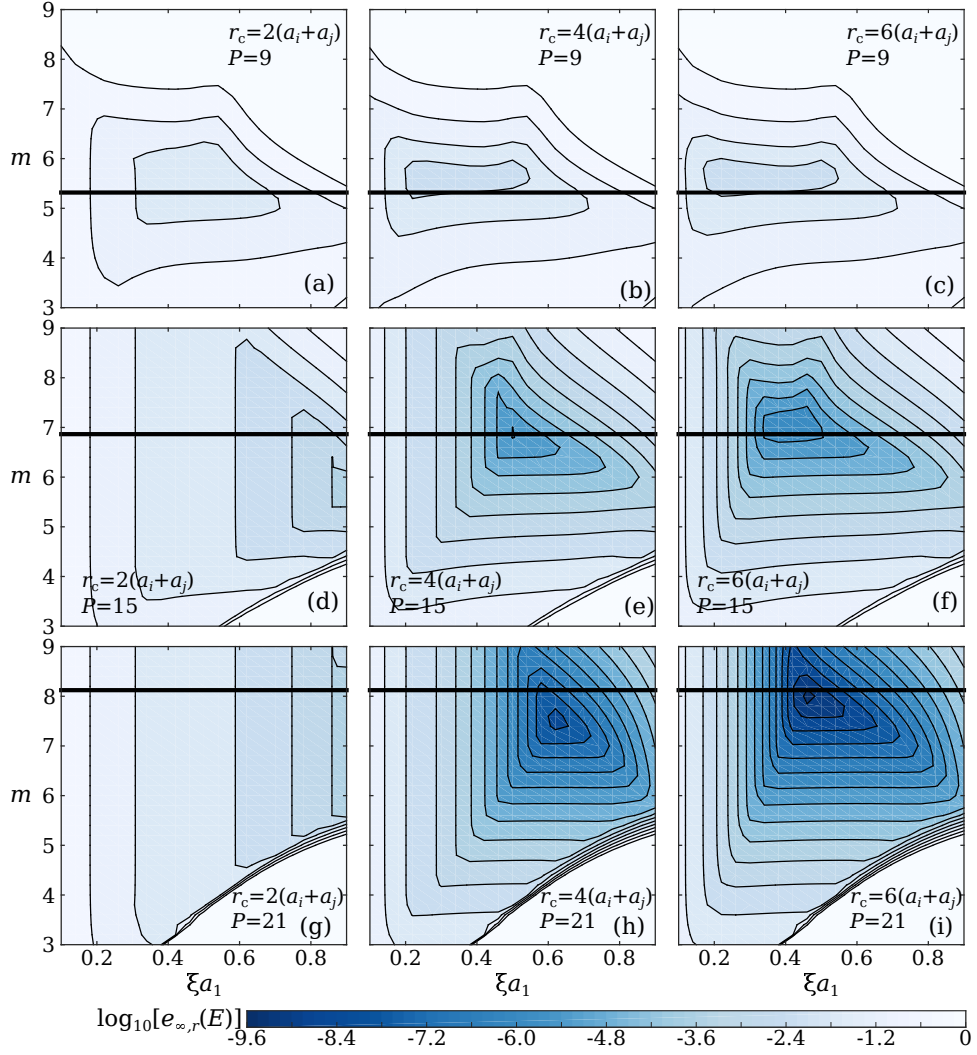


Figure 3: (Color online) The overall accuracy measured in $e_{\infty,r}(E)$ as a function of the splitting parameter ξa_1 and the shape parameter m at $M = 64$ for a real-space cutoff radius $r_c = 2(a_i + a_j)$ (left column), $4(a_i + a_j)$ (middle column), and $6(a_i + a_j)$ (right column), and the interpolation point $P = 9$ (top row), 15 (middle row), and 21 (bottom row). The thick black lines represent $m = \sqrt{\pi P}$. The simulation cell is orthogonal ($\gamma = 0$), and the particle size effects are accounted using the hybrid approach.

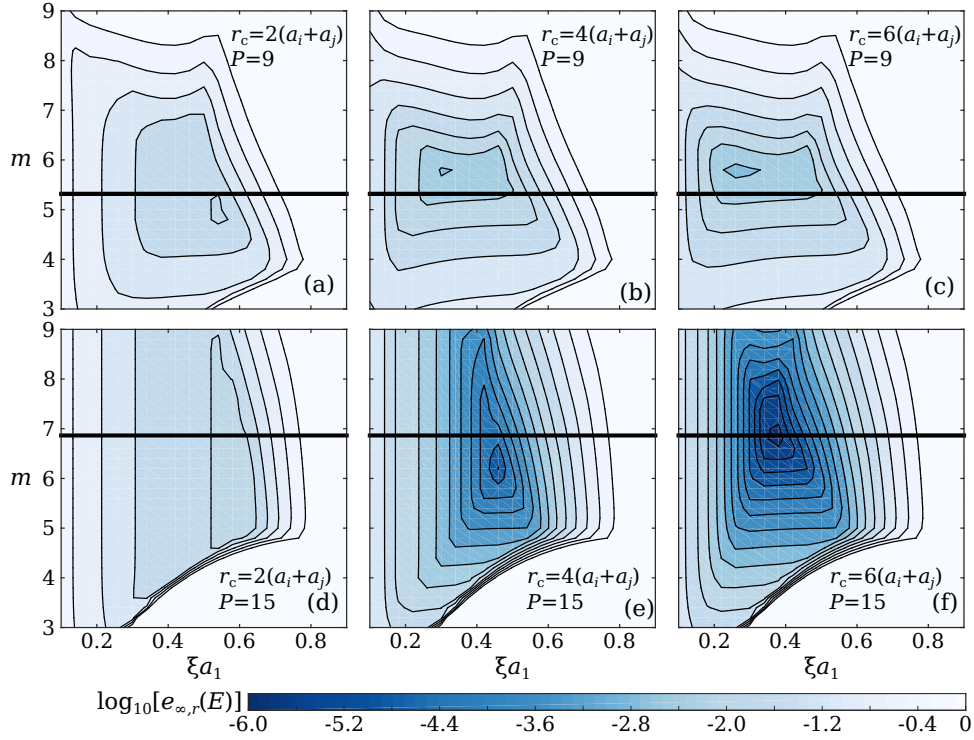


Figure 4: (Color online) The overall accuracy measured in $e_{\infty,r}(E)$ as a function of the splitting parameter ξa_1 and the shape parameter m with $M = 32$ for a real-space cutoff radius $r_c = 2(a_i + a_j)$ (left column), $4(a_i + a_j)$ (middle column), and $6(a_i + a_j)$ (right column), and the interpolation point $P = 9$ (top row) and 15 (bottom row). The thick black lines represent $m = \sqrt{\pi P}$. The simulation cell is orthogonal ($\gamma = 0$), and the particle size effects are accounted using the hybrid approach.




Figure 5: (Color online) The overall mobility accuracy measured in $e_{\infty,r}(E)$ as a function of the splitting parameter ξ with $N = 50, 100,$ and 200 , and $M = 32$ (filled symbols) and 64 (open symbols) for (a): constant box size $L/a_1 = 23.5$ and (b): constant volume fraction $\phi = 0.05$. Changes are based on the baseline case in Sec. 5.1. Other parameters are $P = 13$, $m = 6.7$, and $r_c = 4(a_i + a_j)$.




Figure 6: The accuracy of GPGPU mobility computation measured in $e_{\infty,r}(E)$. (a): the wave-space accuracy as a function of P for various m with the same parameters in Fig. 2b. The GPU results are shown in black lines, and the CPU results in Fig. 2b are reproduced in gray lines. The values of m are annotated in the figure. The solid and dashed lines represent the case of $\gamma = 0$ and 0.5, respectively. (b): The overall mobility accuracy from the GPU (solid lines) and the CPU (dashed lines) computations as a function ξa_1 with $r_c = 4(a_i + a_j)$ and $m = \sqrt{\pi P}$. The corresponding M and P are annotated in the figure.

5.2. Accuracy of the GPGPU implementation

The accuracy of mobility computation using GPGPU programming discussed in Sec. 3.4 is presented in Fig. 6. Clearly, the GPU computations provide sufficient accuracy for dynamic simulations. Fig. 6a shows the GPU wave-space accuracy as a function of the interpolation point P for various shape parameters m for orthogonal ($\gamma = 0$) and sheared ($\gamma = 0.5$) simulation boxes. Here, the particle size effects are incorporated using the hybrid approach in Sec. 3.2, and the SE method parameters are identical to those of Fig. 2b. Moreover, for comparison the data in Fig. 2b are reproduced in gray. In Fig. 6a, the GPU results in black lines are indistinguishable from the CPU results in gray lines when $e_{\infty,r}(E) > 10^{-5}$ for all m and γ , indicating that the GPU computations are only limited by the SP arithmetics. When the error $e_{\infty,r}(E)$ reaches 10^{-5} , increasing the interpolation point P does not improve the computation accuracy on GPUs, while the error in the CPU computations using DP arithmetics continue to decrease until $e_{\infty,r}(E) \sim 10^{-14}$. In addition, the wave-space error remain $e_{\infty,r}(E) \sim 10^{-5}$ after reaching the SP limit even with further increase in P , *i.e.*, increasing P does not adversely affect the wave-space accuracy.

The overall GPU mobility accuracy as a function of ξa_1 is presented in Fig. 6b for two M and P combinations with $m = \sqrt{\pi P}$ and $r_c = 4(a_i + a_j)$ in orthogonal simulation boxes. The errors $e_{\infty,r}(E)$ are computed using the baseline case of Sec. 5.1. The GPU results are shown in solid lines and the CPU results in dashed lines. When the overall error $e_{\infty,r}(E) > 10^{-5}$, *i.e.*, the case of $(M, P) = (32, 13)$ in Fig. 6b, the GPU and the CPU results are indistinguishable from each other. However, the differences are evident for the case of $(M, P) = (64, 21)$. When $0.5 < \xi a_1 < 0.85$, the GPU computations deviate from the CPU results with larger errors due to the SP arithmetics. Beyond this range, the CPU and the GPU results overlap again. In both cases, the accuracy achieved by the GPU mobility computation is sufficient for dynamic simulations, where the error tolerance is typically set at 10^{-3} . The results in Fig. 6 dispel any concerns over the SP accuracy in the GPU mobility computations for dynamic simulations.

5.3. Overall performance

Fig. 7 presents the overall performance of various implementations of the SEASD and the conventional SD as a function of the system size N . The program performance is characterized by the wall time, *i.e.*, the actual time of program execution, to march 100 steps in a dynamic simulation of Brownian suspensions at $Pe = 1$ starting from an equilibrium configuration. The suspension composition is $\lambda = 2$, $y_2 = 0.5$, and $\phi = 0.45$. The SEASD mobility computation parameters are fixed at $M = 32$, $P = 11$, $r_c = 4(a_i + a_j)$ with appropriate ξ and m as they provide sufficient accuracy. The tolerance of the iterative solvers is set at 10^{-3} . For SEASD the far-field Brownian forces are calculated using Eqs. (75) and (76) with $\Delta\tau = 0.2$, and for SEASD-nf the far-field diffusivities are from Table 1.

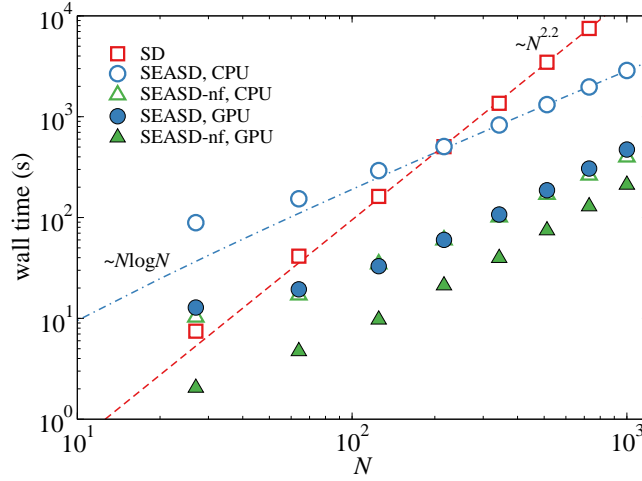


Figure 7: (Color online) The wall time (in second) of 100 time steps in dynamic simulations at $Pe = 1$ as a function of the particle number N using the conventional SD, SEASD, and SEASD-nf. The open symbols represent the CPU mobility computation and the filled symbols the GPU mobility computation. The dashed line show the $O(N^{2.2})$ scaling, and the dash-dotted line show the $O(N \log N)$ scaling. The suspension is bidisperse with $\lambda = 2$, $\gamma_2 = 0.5$, and $\phi = 0.45$ starting from equilibrium configurations.

The conventional SD result is from an efficient polydisperse implementation [11, 12, 83]. All the timing results are collected from a workstation with Intel i7-3770K CPU and NVIDIA GeForce GTX 680 GPU.

Fig. 7 demonstrates the expected $O(N \log N)$ asymptotic scaling of various SEASD implementations, highlighted by the dash-dotted line. The implementations with the CPU mobility computation are shown in open symbols and the GPU mobility computation in filled symbols. The GPU SEASD has almost the same time scaling as the CPU SEASD-nf at all the system size N . Both are almost an order of magnitude faster than the CPU SEASD at a typical system size $N \approx 200$. This clearly demonstrates the power and promise of GPGPU programming in the dynamic simulation of colloidal suspensions. More significant speedup is achieved by combining the mean-field Brownian approximation and the GPU mobility computation. In this case, the speedup of GPU SEASD-nf computation relative to the CPU SEASD ranges between 40 times for small systems and 15 times for large systems. We believe further speedup is still possible by optimizing the GPU implementation. With the speedup shown in Fig. 7, we are able to study dynamics of larger systems at longer times. In addition, compared to the conventional SD, all the SEASD implementations are faster at large enough N due to their favorable scaling. Here, the conventional SD scales as $O(N^{2.2})$, highlighted by the dashed line in Fig. 7. This peculiar scaling is a combined effect of the pairwise grand mobility tensor construction and explicit matrix inversion. At $N \gg 1000$, the scaling should recover $O(N^3)$. In Fig. 7, the break-even between the CPU SEASD and SD is $N = 216$, and for GPU SEASD at $N \approx 40$. At all the system sizes studied here, the GPU SEASD-nf is always faster than the conventional SD.

6. Static and dynamic simulation results

6.1. Short-time transport properties

In this section we present static SEASD simulation results on the short-time transport properties of monodisperse and bidisperse hard-sphere suspensions. With the iterative computation scheme in Sec. 4.1, the short-time translational and rotational self-diffusivities, instantaneous sedimentation velocities, and high-frequency dynamic shear and bulk viscosities can be straightforwardly evaluated. Other transport properties can also be calculated with an appropriate computation scheme.

The suspension short-time limit refers to a time scale t satisfying $\tau_I \ll t \ll \tau_D$, where τ_I is the inertial time and τ_D is the diffusion time. The inertia time $\tau_I = \frac{2}{9} \rho_p a_p^2 / \eta_0$, where ρ_p and a_p are the characteristic particle density and radius, describes the time required for the *particle* momentum to dissipate by interacting with the solvent. When $\tau_I \ll t$, the particle momentum dissipates almost instantaneously and the particle dynamics are completely overdamped. The

Table 1: The polynomial coefficient fitted from the far-field diffusivities in Fig. 9. The data is for polydisperse suspensions with $\lambda = 2$ and $y_2 = 0.5$. The far-field self-diffusivity d_s^{ff} can be expressed as $d_s^{\text{ff}}/d_0 = 1 + c_1\phi + c_2\phi^2 + c_3\phi^3$, where d_0 is the single particle diffusivity.

	$d_{s,1}^{t,\text{ff}}$	$d_{s,2}^{t,\text{ff}}$	$d_{s,1}^{r,\text{ff}}$	$d_{s,2}^{r,\text{ff}}$
c_1	-1.27	-1.70	-0.207	-0.538
c_2	0.536	1.005	-0.131	-0.312
c_3	-0.018	-0.12	-0.091	0.19

diffusion time $\tau_D = 6\pi\eta_0 a_p^3/k_B T$ characterizes the time scale of suspension configuration change and $t \ll \tau_D$ ensures that the transport properties entirely arise from the (instantaneous) HIs. Therefore, they are only determined by the configuration X , and can be calculated by sampling independent but equivalent configurations. In this work we use the Monte-Carlo procedure of Wang & Brady [11]: the hard-sphere configurations are first generated by an event-driven Lubachesky-Stillingr algorithm [84, 85], followed by a short equilibration. The transport properties are then computed statically. Here we compare the results from the SEASD with CPU mobility computation with our recent conventional SD results [11]. Although SEASD and SD are based on the same formalism, the grand mobility tensor \mathfrak{M} constructed from SD includes an additional mean-field quadrupole term [63], which can have quantitative consequences. For bidisperse hard-sphere suspensions, we focus on the composition with $\lambda = 2$ and $y_2 = 0.5$. In the SEASD computations, the system size is $N = 800$, and the results are averaged over 500 independent configurations. Note that for simple cubic array of monodisperse particles, SEASD produces identical results as those of Sierou & Brady [36].

6.1.1. Short-time translational and rotational self-diffusivities

The microscopic definition of the short-time translational and rotational self-diffusivities, $d_{s,\alpha}^t$ and $d_{s,\alpha}^r$ respectively, for homogeneous suspensions are,

$$d_{s,\alpha}^t = \frac{k_B T}{N_\alpha} \left\langle \sum_{i \in \alpha} \hat{q} \cdot \boldsymbol{\mu}_{ii}^t \cdot \hat{q} \right\rangle, \text{ and } d_{s,\alpha}^r = \frac{k_B T}{N_\alpha} \left\langle \sum_{i \in \alpha} \hat{q} \cdot \boldsymbol{\mu}_{ii}^{rr} \cdot \hat{q} \right\rangle, \quad (80)$$

where \hat{q} is a vector of unit length for the averaging process and $\boldsymbol{\mu}_{ii}^t$ and $\boldsymbol{\mu}_{ii}^{rr}$ are respectively the diagonal blocks of the force-linear velocity and torque-angular velocity couplings in $\mathbf{R}_{\mathcal{F}\mathcal{U}}^{-1}$. Note that $i \in \alpha$ in Eq. (80) suggests the summation is restricted to particles of species α . The diffusivities are computed using the matrix-free approach of Sierou & Brady [36]: the velocity disturbance \mathcal{U}^R corresponding to a stochastic external force \mathcal{F}^R satisfying $\langle \mathcal{F}^R \rangle = 0$ and $\langle \mathcal{F}^R \mathcal{F}^R \rangle = \mathcal{I}$ is evaluated. It is straightforward to show that the ensemble average $\langle \mathcal{U}^R \mathcal{F}^R \rangle = \text{diag}(\mathbf{R}_{\mathcal{F}\mathcal{U}}^{-1})$, allowing extraction of the diffusivities in Eq. (80).

The computed short-time translational self-diffusivities $d_{s,\alpha}^t$ exhibit a strong $N^{-1/3}$ size dependence due to the periodic boundary conditions. The size dependence from an N -particle system can be eliminated by adding the following quantity to the results,

$$\Delta_N d_{s,\alpha}^t = \frac{1.76 d_{0,1}^t}{(x_1 + x_2 \lambda^3)^{1/3}} \frac{\eta_0}{\eta_s} \left(\frac{\phi}{N} \right)^{1/3}, \quad (81)$$

where $d_{0,1}^t = k_B T / (6\pi\eta_0 a_1)$ is Stokes-Einstein-Sutherland diffusivity for species 1, and η_s is the high-frequency dynamic shear viscosity from the same configurations. The shear viscosity exhibits little size dependence, and can be directly used. The effectiveness of Eq. (81) has been demonstrated by Wang & Brady [11] in the wave-number-dependent hydrodynamic functions. The results here always contain this finite size N correction.

Fig. 8a and Fig. 8b respectively present $d_{s,\alpha}^t/d_{0,\alpha}^t$ and $d_{s,\alpha}^r/d_{0,\alpha}^r$ of monodisperse and bidisperse suspensions, where the single particle translational and rotational self-diffusivities are $d_{0,\alpha}^t = k_B T / (6\pi\eta_0 a_\alpha)$ and $d_{0,\alpha}^r = k_B T / (8\pi\eta_0 a_\alpha^3)$. The SEASD results, shown in symbols, agree well with the conventional SD results shown in lines. As expected, both $d_{s,\alpha}^t$ and $d_{s,\alpha}^r$ decrease with increasing volume fraction ϕ , and for bidisperse suspensions, the small particles show diffusivity enhancement while the large particles exhibit diffusivity suppression. Compared to $d_{s,\alpha}^t$, $d_{s,\alpha}^r$ are less sensitive to the volume fractions ϕ , but more sensitive to the particle sizes λ . The SEASD results for large particles show larger error bars compared to the SD results [11], most likely due to the stochastic computation procedure.




Figure 8: (Color online) The species short-time (a): translational and (b): rotational self-diffusivities, $d_{s,\alpha}^t$ and $d_{s,\alpha}^r$ respectively, as a function of the total volume fraction ϕ for monodisperse and bidisperse hard-sphere suspensions with $\lambda = 2$, $\gamma_2 = 0.5$. The results are scaled with the single particle translation and rotational diffusivity, $d_{0,\alpha}^t$ and $d_{0,\alpha}^r$, respectively. The SEASD results are shown in symbols and the conventional SD results from Wang & Brady [11] are shown as lines.




Figure 9: (Color online) The species far-field short-time translational and rotational self-diffusivities, $d_{s,\alpha}^{t,ff}$ and $d_{s,\alpha}^{r,ff}$ respectively, as a function of the total volume fraction ϕ for bidisperse hard-sphere suspensions with $\lambda = 2$ and $\gamma_2 = 0.5$. The results are scaled with the single particle translation and rotational diffusivity, $d_{0,\alpha}^t$ and $d_{0,\alpha}^r$, respectively. The symbols are the computation results, and the dashed and the dash-dotted lines are polynomial fittings for the small and the large particles, respectively.




Figure 10: (Color online) The scaled species instantaneous sedimentation velocity, $U_{s,\alpha}/U_{0,\alpha}$, as a function of the total volume fraction ϕ for monodisperse and bidisperse hard-sphere suspensions with $\lambda = 2$ and $\gamma_2 = 0.5$. The single particle sedimentation velocity is $U_{0,\alpha}$. The SEASD results are shown in symbols and the conventional SD results from Wang & Brady [11] are shown as lines.

We also calculated the far-field short-time translational and rotational self-diffusivities $d_{s,\alpha}^{t,\text{ff}}$ and $d_{s,\alpha}^{r,\text{ff}}$, where “ff” suggests only the far-field HIs without the lubrication corrections are considered. They are the input for subsequent SEASD-nf computations in Sec. 6.2 and 6.3. The $N^{-1/3}$ size dependency in the far-field translational diffusivity $d_{s,\alpha}^{t,\text{ff}}$ is corrected using Eq. (81) with the corresponding far-field viscosity. Fig. 9 shows $d_{s,\alpha}^{t,\text{ff}}$ and $d_{s,\alpha}^{r,\text{ff}}$ for bidisperse suspensions up to $\phi = 0.62$. Compared to Fig. 8, the far-field diffusivities exhibit weaker volume fraction dependence, and they do not have sharp reductions at high volume fractions. Consistent with Fig. 8, $d_{s,\alpha}^{r,\text{ff}}$ also exhibits stronger particle size dependence compared to its translational counterpart. In general, the ϕ dependence of any scaled far-field diffusivity d_s^{ff}/d_0 , with d_0 the corresponding single-particle data, can be adequately captured by a cubic polynomial $d_s^{\text{ff}}/d_0 = 1 + c_1\phi + c_2\phi^2 + c_3\phi^3$, where the coefficients c_i , $i \in \{1, 2, 3\}$, only depend on the suspension composition. The fitting coefficients for bidisperse suspensions with $\lambda = 2$ and $\gamma_2 = 0.5$ are presented Table 1. The polynomial fittings, also shown in Fig. 9 in dashed and dash-dotted lines for the small and the large particles, respectively, indeed describe the computation data. Not shown in Fig. 9 are the SEASD far-field diffusivities for monodisperse suspensions, which are identical to those of Banchio & Brady [39].

6.1.2. Instantaneous sedimentation velocity

The species instantaneous sedimentation velocities $U_{s,\alpha}$ are computed by applying a uniform external force F_α to each species. For bidisperse suspensions, the sedimentation velocity $U_{s,\alpha}$ also depends on the species density ratio [8], $\gamma = \Delta\rho_2/\Delta\rho_1$, with $\Delta\rho_\alpha = \rho_\alpha - \rho_0$ the density difference of species α . The species force ratio satisfies $F_2/F_1 = \gamma\lambda^3$, and here we set $\gamma = 1$ to facilitate comparison with earlier results. To eliminate the $N^{-1/3}$ size dependence, the following corrections are added to the results:

$$\Delta_N U_{s,1} = \frac{1.76U_{0,1}}{(x_1 + x_2\lambda^3)^{\frac{1}{3}}} \frac{\eta_0}{\eta_s} \left(\frac{\phi}{N}\right)^{\frac{1}{3}} \left[S_{11}(0) + \lambda^3\gamma \sqrt{\frac{x_2}{x_1}} S_{12}(0) \right], \quad (82)$$

$$\Delta_N U_{s,2} = \frac{1.76U_{0,1}}{(x_1 + x_2\lambda^3)^{\frac{1}{3}}} \frac{\eta_0}{\eta_s} \left(\frac{\phi}{N}\right)^{\frac{1}{3}} \left[\sqrt{\frac{x_1}{x_2}} S_{21}(0) + \lambda^3\gamma S_{22}(0) \right], \quad (83)$$

where $U_{0,\alpha} = F_\alpha/(6\pi\eta_0 a_\alpha)$ is the single particle sedimentation velocity and $S_{\alpha\beta}(0)$ is the partial static structural factors in the zero wave number limit. Eqs. (82) and (83) are based on the finite-size correction for partial hydrodynamic functions [11]. Here, the partial static structural factors are computed from the polydisperse Percus-Yevic integral equations [86–89].




Figure 11: (Color online) The high-frequency dynamic (a): shear viscosity η_s and (b): bulk viscosity κ_s as functions of the total volume fraction ϕ for monodisperse and bidisperse hard-sphere suspensions with $\lambda = 2$ and $\gamma_2 = 0.5$. The results are scaled with the solvent viscosity η_0 , and only the particle contributions, $\eta_s/\eta_0 - 1$ and $(\kappa_s - \kappa_0)/\eta_0$ are presented. The SEASD results are shown as symbols and the conventional SD results [11] are shown as lines.

Fig. 10 presents the SEASD $U_{s,\alpha}/U_{0,\alpha}$ in symbols, which are not the identical to the conventional SD results shown in lines. The difference is especially pronounced at high volume fractions. For monodisperse suspensions, the SEASD and the conventional SD agree with each other satisfactorily up to $\phi \approx 0.3$, and at higher ϕ , the SEASD results become significantly higher. This difference is from the mean-field quadrupole term, which is absent in SEASD. Despite the quantitative differences, the SEASD monodisperse sedimentation velocity remain positive and physical. A similar overestimation of the sedimentation velocity is also found when comparing ASD results [36] and the conventional SD results [63] for simple cubic arrays.

The differences between the SEASD and the conventional SD results are more significant for bidisperse suspensions. For $U_{s,1}$ of the small particles, the differences are not evident until $\phi = 0.3$, and for $U_{s,2}$ of the large particles, the differences are obvious even at $\phi \approx 0.2$. Moreover, $U_{s,2}$ exhibits a minimum and increases with ϕ at higher volume fraction, leading to a crossing of $U_{s,1}$ and $U_{s,2}$ at $\phi = 0.45$. These unphysical behaviors are caused by inaccurate HI computations at the stresslet-strain rate level. Apparently, the HIs of the large particles, which are surrounded by many small particles, are more complicated than those of the small particles and more difficult to capture accurately. Note that for sedimentation the lubrication interactions are not important and one must rely on the far-field mobility for all HIs.

Fig. 10 also illustrates that sedimentation problems in dense bidisperse suspensions, even at $\lambda = 2$, is challenging for SEASD. Incorporating the mean-field quadrupole term [63], $(1 - \frac{1}{5}\phi)$, in the grand mobility tensor can significantly improve the results [11]. However, such incorporation is not carried out in this work.

6.1.3. High-frequency dynamic shear and bulk viscosities

The high-frequency dynamic shear and bulk viscosities, η_s and κ_s , are respectively defined as,

$$\eta_s = \eta_0 + n\langle \mathbf{S}^E \rangle_{xy} / \dot{\gamma}, \text{ and } \kappa_s = \kappa_0 + \frac{1}{3}n\langle \mathbf{S}^E \rangle : \mathbf{I} / \dot{\epsilon}, \quad (84)$$

where $\dot{\gamma}$ is the imposed strain rate, $\dot{\epsilon}$ is the imposed uniform expansion rate, \mathbf{S}^E is the hydrodynamic stresslet in Eq. (65), and the subscript xy denotes the velocity-velocity gradient component. They are directly computed from SEASD and exhibit little size dependencies. Experimentally, η_s and κ_s are measured by imposing high-frequency, low-amplitude deformations, such that the suspension microstructures are only slightly perturbed, and the Brownian stress contributions are out of phase with the applied deformations [90].

Fig. 11a and 11b present the volume fraction ϕ dependency of the particle contributions to the high-frequency dynamic shear and bulk viscosities, $\eta_s/\eta_0 - 1$ and $(\kappa_s - \kappa_0)/\eta_0$, respectively. The SEASD calculations are shown in symbols, and the corresponding conventional SD results are shown in lines. For η_s , the SEASD and the conventional SD results agree well over the entire ϕ range. The results for monodisperse and bidisperse suspensions with $\lambda = 2$ are




Figure 12: (Color online) The equilibrium osmotic pressure $\Pi/(nk_B T)$ of monodisperse and bidisperse Brownian suspensions with $\lambda = 2$ and $\gamma_2 = 0.5$, as a function of volume fraction ϕ . The dashed line represents the CS equation of state, Eq. (86), and the dash-dotted line represents the BMCSL equation of state, Eq. (87).

almost identical when $\phi < 0.55$. At higher volume fractions, the monodisperse η_s are more sensitive to ϕ compared to the bidisperse results, as introducing particles of difference sizes significantly alters the suspension hydrodynamic environment in this limit. Unlike sedimentation, for the shear viscosity lubrication interactions are important and dominate the behavior at high ϕ .

For the high-frequency dynamics bulk viscosity κ_s in Fig. 11b, the SEASD and conventional SD results show qualitative agreement with noticeable quantitative differences at moderate ϕ : the SEASD results are higher and less sensitive to the particle size ratio λ . The differences are caused by different pressure moment computation procedures. Recall that the far-field grand mobility tensor \mathfrak{M} is not symmetric by construction, and the symmetry of \mathfrak{M}^{-1} must be restored for subsequent calculations. This is done in conventional SD by explicit copy of matrix elements after the matrix inversion [91]. This is not applicable for the matrix-free computation of \mathfrak{M} in SEASD. Here, the pressure moment is computed from the far-field forces and stresses. Fig. 11b shows that the two conceptually equivalent approaches do lead to small quantitative differences. Moreover, for dense suspensions, such differences are masked by the dominance of lubrication interactions. Therefore, the SEASD and the conventional SD results agree well at low and high ϕ . Near the close packing limit, κ_s for bidisperse suspensions is significantly lower than that of the monodisperse case, since the particle size polydispersity improves the particle packing.

6.2. Equilibrium suspensions

Here we present the dynamic simulation results with SEASD and SEASD-nf for monodisperse and bidisperse Brownian suspensions at zero Péclet number. In particular, we are interested in the following equilibrium properties: the osmotic pressure Π , the high-frequency dynamic bulk modulus K'_∞ , and high-frequency dynamic shear modulus, G'_∞ . The dynamic simulations are carried out with 100 particles over 200 diffusive time units with a time step $\Delta t d_{0,1}^t / a_1^2 = 10^{-3}$. The mobility computation in SEASD is performed on GPUs with $M = 32$, $P = 11$, and $r_c = 4(a_i + a_j)$, and the far-field Brownian force is calculated using the IVP method in Sec. 4.3 with $\Delta\tau = 0.1$. The tolerance for the iterative solver is 10^{-3} and the tolerance for matrix inversion in Eqs. (75) and (76) is 0.02. The composition of bidisperse suspensions are $\lambda = 2$ and $\gamma_2 = 0.5$. Therefore, for the SEASD-nf computations the coefficients in Table 1 are used. Note that with $Pe = 0$, SEASD-nf computations do not contain far-field mobility evaluations.

6.2.1. Osmotic pressure

The osmotic pressure of an equilibrium suspension is defined as

$$\Pi = nk_B T - \frac{1}{3} n \langle \mathbf{S}^B \rangle : \mathbf{I}, \quad (85)$$




Figure 13: (Color online) The high-frequency dynamic moduli: (a) the bulk modulus $K'_\infty a_1^3 / (k_B T)$, and (b) the shear modulus $G'_\infty a_1^3 / (k_B T)$, as functions of volume fraction ϕ for equilibrium monodisperse and bidisperse Brownian suspensions with $\lambda = 2$ and $y_2 = 0.5$. The results are computed from SEASD (filled symbols) and SEASD-nf (open symbols).

where $\langle \mathbf{S}^B \rangle$ is the Brownian stresslet in Eq. (67). For rigid particles with no-slip boundary conditions, Brady [51] showed that the osmotic pressure is purely hydrodynamic in origin, and is identical to that of a hard-sphere fluid. The osmotic pressure of monodisperse suspensions is well described by the Carnahan-Starling (CS) equation up to the fluid-solid transition,

$$\frac{\Pi}{nk_B T} = \frac{1 + \phi + \phi^2 - \phi^3}{(1 - \phi)^3}. \quad (86)$$

The CS equation of state is extended to polydisperse suspensions as the Boublik-Mansoori-Carnahan-Starling-Leland (BMCSL) equation [92]:

$$\frac{\Pi}{nk_B T} = \frac{1 + \phi + \phi^2 - 3\phi(z_1 + z_2\phi) - z_3\phi^3}{(1 - \phi)^3}, \quad (87)$$

where $z_1 = \Delta_{12}(1 + \lambda) / \sqrt{\lambda}$, $z_2 = \Delta_{12}(y_1\lambda + y_2) / \sqrt{\lambda}$, and $z_3 = [(y_1^2 x_1)^{1/3} + (y_2^2 x_2)^{1/3}]^3$ with $\Delta_{12} = \sqrt{y_1 y_2} \sqrt{x_1 x_2} (\lambda - 1)^2 / \lambda$.

Fig. 12 presents the equilibrium osmotic pressure of monodisperse and bidisperse suspensions with $\lambda = 2$ and $y_2 = 0.5$ as functions of ϕ using SEASD and SEASD-nf computations. The CS [Eq. (86)] and the BMCSL [Eq. (87)] equations of state at the corresponding bidisperse compositions are respectively shown in dashed and dash-dotted lines. Also shown in Fig. 12 are the static computation results with $N = 200$, denoted “static”. The static computations do not consider particle dynamics, and calculate the osmotic pressure by taking a full Brownian step from independent particle configurations in a Monte-Carlo fashion. In Fig. 12, at each volume fraction 500 independent configurations are used in the static computations.

The osmotic pressures from the SEASD, the SEASD-nf, and the static computations agree with the CS and BMCSL predictions in Fig. 12. The static computations show the best agreement over the entire ϕ range, and this directly validates the Brownian stress computation method in Sec. 4.3. The dynamic SEASD results are slightly higher than the theoretical predictions because the configuration evolution is affected by the finite $\Delta\tau$ in the far-field Brownian force computation. The slight difference does not invalidate this approach as it is well within the discretization errors of Eqs. (75) and (76). Note that, as long as the tolerances for the iterative solution of Eqs. (75) and (76) are smaller than the discretization step size $\Delta\tau$, the principal source of error is the time discretization. We have verified that reducing the iterative solver tolerance with fixed $\Delta\tau$ does not improve the results. Finally, the agreement in the bidisperse osmotic pressures from SEASD-nf and the BMCSL equation validates the extension of the mean-field Brownian approximation to polydisperse systems. The SEASD-nf results are only slightly lower than the theoretical predictions, which is acceptable considering the substantial speedup offered by this approach.

6.2.2. High-frequency dynamic moduli

The suspension high-frequency dynamic bulk and shear moduli, K'_∞ and G'_∞ respectively, can be computed from the short-time limit of the pressure-pressure and stress-stress autocorrelation functions [91, 93, 94], *i.e.*,

$$K'_\infty = \lim_{t \rightarrow 0} \frac{V}{k_B T} \langle \delta \Pi(t) \delta \Pi(0) \rangle, \text{ and } G'_\infty = \lim_{t \rightarrow 0} \frac{V}{k_B T} \langle \sigma(t) \sigma(0) \rangle, \quad (88)$$

where $\delta \Pi$ is the osmotic pressure fluctuations and σ is the off-diagonal components of the bulk stress $\langle \mathbf{\Sigma} \rangle$ in Eq. (64). Note that the viscoelasticity of colloidal suspensions is entirely of hydrodynamic origin, and without HIs, *e.g.*, in hard-sphere fluids, these moduli are infinite.

Fig. 13a and 13b respectively present K'_∞ and G'_∞ of monodisperse and bidisperse suspensions as functions of ϕ from the same SEASD and SEASD-nf dynamic simulations of Fig. 12. Both K'_∞ and G'_∞ grow rapidly with ϕ , and at the same volume fraction, the monodisperse moduli are always higher. In Fig. 13a, the bulk modulus K'_∞ computed from SEASD and SEASD-nf share the same qualitative behavior. However, the SEASD results are almost always higher than the SEASD-nf results except at small ϕ , and their differences grow with increasing ϕ . This is consistent with the growing differences in Π with increasing ϕ in Fig. 12. On the other hand, in Fig. 13b the differences in the shear modulus G'_∞ between the SEASD and the SEASD-nf results decrease with increasing ϕ , with the SEASD-nf data higher at low volume fractions. Note that the bidisperse SEASD results show large fluctuations when $\phi = 0.2 \sim 0.25$, most likely due to the small number of large particles at $N = 100$ and the particular particle spacing at this volume fraction. Finally, small differences in fluctuation quantities such as K'_∞ and G'_∞ are expected for SEASD and SEASD-nf because the Brownian stresses are computed differently. However, more importantly, the same qualitative behaviors are followed in both methods.

6.3. Rheology of bidisperse suspensions

The final validation of SEASD and SEASD-nf is the steady shear rheology of Brownian suspensions at constant strain rate. Both monodisperse and bidisperse hard-sphere suspensions are considered: the volume fractions are fixed at $\phi = 0.45$ in both cases, and the bidisperse composition is $\lambda = 2$ and $y_2 = 0.5$. The results are extracted from SEASD and SEASD-nf simulations with GPU mobility computation over a wide range of Péclet number $Pe = 6\pi\eta_0 a_1^3 \dot{\gamma} / (k_B T)$. Moreover, we introduce a small excluded volume on each particle to emulate the effects of surface asperities or polymer coating and to prevent particle overlap. It is characterized by,

$$\delta = 1 - a_i / b_i, \quad (89)$$

where b_i is the excluded volume radius for each particle. The SEASD and SEASD-nf simulations are carried out at $\delta = 5 \times 10^{-4}$ with $N = 200$ over 150 dimensionless time units with a step size 10^{-3} . Other simulation parameters are similar to those in Sec. 6.2. The data are averaged in segments after the steady state is reached, usually after 20 dimensionless time units. As is customary, the x -direction is the velocity direction, the y -direction is the velocity gradient direction, and the z -direction is the vorticity direction.

6.3.1. Shear viscosity

Fig. 14a and 14b respectively present the Brownian viscosity η^B and the flow viscosity η^E as functions of the Péclet number. These viscosities are defined as

$$\eta^B = n \langle \mathbf{S}^B \rangle_{xy} / \dot{\gamma} \text{ and } \eta^E = n \langle \mathbf{S}^E \rangle_{xy} / \dot{\gamma}, \quad (90)$$

with $\langle \mathbf{S}^B \rangle$ in Eq. (67) and $\langle \mathbf{S}^E \rangle$ in Eq. (65). In this figure, the monodisperse data are shown in squares and the bidisperse data in triangles, with the SEASD results in filled symbols and the SEASD-nf results in open symbols. For comparison, the SD results of Foss & Brady [72] for monodisperse suspensions are presented in open circles. To clarify the effects of the excluded volume parameter δ on viscosities, another set of monodisperse SD simulations with $N = 30$ are performed at $\delta = 5 \times 10^{-4}$ and 10^{-5} , and the results are shown as crosses and pluses respectively. In all cases, the stress contributions from inter-particle forces are negligible, and therefore are not presented.

In Fig. 14 both the Brownian viscosity η^B and the flow viscosity η^E exhibit the expected behaviors: with increasing Pe , η^B decreases (shear-thinning) and η^E grows (shear-thickening). In addition, there are several important

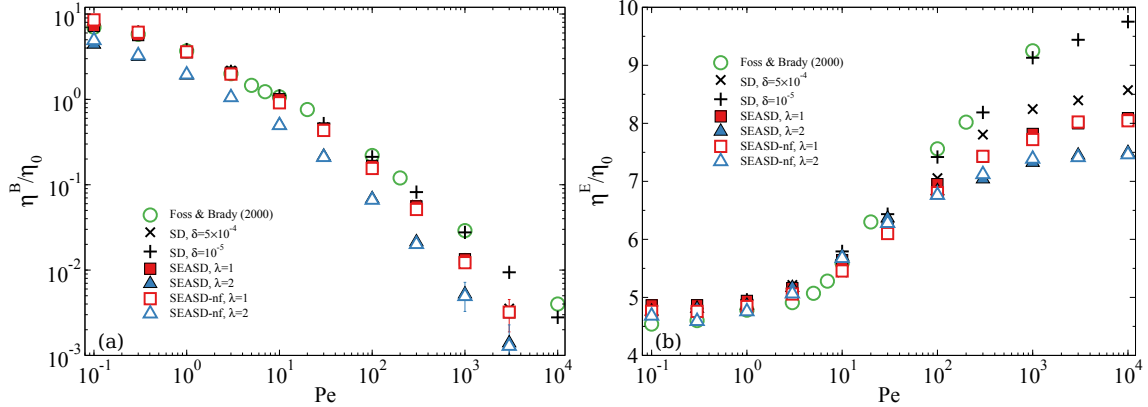


Figure 14: (Color online) Different viscosity contributions to the rheology of monodisperse and bidisperse hard-sphere suspensions: (a) the Brownian viscosity η^B/η_0 and (b) the flow viscosity η^E/η_0 , as functions of Pe. The volume fraction $\phi = 0.45$ in both cases, and the bidisperse composition is $\lambda = 2$ and $\gamma_2 = 0.5$.

observations. First of all, the excluded volume parameter δ introduces quantitative effects on the suspension rheology, especially at high Pe. Comparing the SD results with $\delta = 5 \times 10^{-4}$ and 10^{-5} , increasing δ enhances the shear-thinning of η^B and weakens the shear-thickening of η^E , especially at high Pe. At low Pe, the effect of δ is almost unnoticeable. The SD results at $\delta = 10^{-5}$ agree well with those of Foss & Brady [72], and the results at $\delta = 5 \times 10^{-4}$ are consistent with the monodisperse SEASD and SEASD-nf results, with larger differences shown in η^E . This difference is most likely due to the number of particles in the computations. Next, the bidisperse Brownian viscosity η^B is always lower than the monodisperse value at all Pe, and for the flow viscosity η^E , their difference is most apparent at high Pe. The large difference in η^E at high Pe suggests distinct HIs and structures between the monodisperse and the bidisperse suspensions, since Fig. 11a suggests η^E is insensitive to equilibrium suspension structures at $\phi = 0.45$. Finally, the SEASD and SEASD-nf results in Fig. 14 almost always overlap each other, showing that the mean-field Brownian approximation is valid over the entire Péclet number range. At high Pe, the Brownian viscosity η^B from SEASD shows larger fluctuations compared to the SEASD-nf results as the Brownian stresses are difficult to compute with highly anisotropic structures. However, these fluctuations do not affect the overall viscosity since the Brownian contribution at high Pe is insignificant.

6.3.2. Non-equilibrium osmotic pressures

Fig. 15a and 15b present the Brownian and the flow contributions to the suspension osmotic pressure,

$$\Pi^B = nk_B T - \frac{1}{3}n\langle \mathbf{S}^B \rangle : \mathbf{I} \text{ and } \Pi^E = -\frac{1}{3}n\langle \mathbf{S}^E \rangle : \mathbf{I}, \quad (91)$$

respectively, as functions of Péclet number Pe. In these figures, the scaling for the Brownian contribution is $nk_B T$ and the scaling for the flow contribution Π^E is $\eta_0 \dot{\gamma}$. Similar to Fig. 15, the monodisperse data are presented in squares and the bidisperse data in triangles, with the SEASD results in filled symbols and SEASD-nf results in open symbols. Fig. 15 also presents the $N = 30$ monodisperse SD results with $\delta = 5 \times 10^{-4}$ and 10^{-5} in crosses and pluses, respectively. Similarly to the shear stresses, the inter-particle contribution to the osmotic pressures is also negligible compared to the contributions from HIs.

In Fig. 15, both $\Pi^B/(nk_B T)$ and $\Pi^E/(\dot{\gamma}\eta_0)$ grow with increasing Pe when $Pe < 100$. The Brownian contribution $\Pi^B/(nk_B T)$ asymptotes the equilibrium value as $Pe \rightarrow 0$. At higher Pe, the influence of the excluded volume parameter δ becomes apparent. For the Brownian osmotic pressure contribution $\Pi^B/(nk_B T)$, the SD results at $\delta = 10^{-5}$ continuously grow with Pe up to $Pe = 10^4$, the highest value in our study, while with $\delta = 5 \times 10^{-4}$, a maximum in $\Pi^B/(nk_B T)$ around $Pe = 10^3$ is apparent. After the maximum, $\Pi^B/(nk_B T)$ decreases slowly with growing Pe. In this case, the parameter δ not only brings quantitative, but also qualitative differences. On the other hand, the flow osmotic pressure contribution $\Pi^E/(\dot{\gamma}\eta_0)$ increases and reaches a plateau at high Pe. Comparing the SD results with $\delta = 5 \times 10^{-4}$ and 10^{-5} , increasing δ reduces the final plateau value of $\Pi^E/(\dot{\gamma}\eta_0)$ at a smaller Pe. Apparently, the high




Figure 15: (Color online) Different contributions to the osmotic pressures of monodisperse and bidisperse hard-sphere suspensions: (a) the Brownian contribution scaled with $nk_B T$, $\Pi^B / (nk_B T)$, and (b) the flow contribution scaled with $\eta_0 \dot{\gamma}$, $\Pi^E / (\dot{\gamma} \eta_0)$, as functions of Pe. The volume fraction is $\phi = 0.45$ in both cases, and the bidisperse composition is $\lambda = 2$ and $y_2 = 0.5$.

Pe osmotic pressure is very sensitive to the excluded volume parameter δ . In terms of the normal viscosity, *i.e.*, $\Pi / \dot{\gamma}$ with $\Pi = \Pi^B + \Pi^E$, increasing δ weakens the shear thickening of the normal viscosity. Furthermore, the SD results at $\delta = 10^{-5}$ agree qualitatively with the results of Yurkovetsky & Morris [53], with slight quantitative difference due to different osmotic pressure computations. At $\delta = 5 \times 10^{-4}$, the Brownian osmotic pressures Π^B from SD and SEASD almost overlap each other in Fig. 15a, and Π^E from SEASD is lower than the SD results in Fig. 15b. Similarly to Fig. 14b, the difference is most likely due to the small system sizes in the SD computations. Moreover, the SEASD Π^B also exhibits larger error bars at high Pe due to the Brownian stress computation, but such errors are of little consequences on the suspension total osmotic pressures.

For the bidisperse results shown in triangles in Fig. 15, the Brownian osmotic pressure Π^B is always lower than its monodisperse counterpart, and the bidisperse Π^E is first slightly higher than the monodisperse results at low Pe and then lower at high Pe. The crossing of the monodisperse and bidisperse Π^E demonstrates the complex interplay between HIs and structures in polydisperse systems.

The SEASD-nf results in Fig. 15 agree qualitatively with the SEASD computations. However, for Π^B , there are quantitative differences at both $\lambda = 1$ and $\lambda = 2$, with the SEASD-nf results systematically lower. This difference is inherently associated with the far-field Brownian force computations in Sec. 4.3 and the mean-field Brownian approximations, and is also encountered in Fig. 12. However, the quantitative discrepancies in Π^B are still within the discretization errors of $\Delta \tau$ in Eqs. (75) and (76). On the other hand, for Π^E , the SEASD-nf and SEASD results almost always overlap each other over the entire Pe range for both bidisperse and monodisperse suspensions. SEASD-nf satisfactorily captures both contributions of the suspension osmotic pressures, Π^B and Π^E .

6.3.3. Normal stress differences

The first normal stress difference N_1 and the second normal stress difference N_2 , defined as

$$N_1 = \langle \boldsymbol{\Sigma} \rangle_{xx} - \langle \boldsymbol{\Sigma} \rangle_{yy} \text{ and } N_2 = \langle \boldsymbol{\Sigma} \rangle_{yy} - \langle \boldsymbol{\Sigma} \rangle_{zz}, \quad (92)$$

describe the stress anisotropy in sheared suspensions, and are important for understanding phenomena such as the shear-induced particle migrations [52]. The normal stress differences N_1 and N_2 are respectively shown in Fig. 16a and Fig. 16b. The monodisperse data are shown in squares and the bidisperse data in triangles, with SEASD results in filled symbols and SEASD-nf results in open symbols. In addition, in Fig. 16, the SD results of Foss & Brady [72] are presented in circles, and the SD computations at $N = 30$ with $\delta = 5 \times 10^{-4}$ and 10^{-5} are respectively shown in crosses and pluses.

In general, the first normal stress difference N_1 in Fig. 16a changes sign from positive to negative with increasing Pe, and the second normal stress N_2 in Fig. 16b remains negative for all Pe studied and exhibits weak Pe dependence. The data with small systems are strongly scattered, particularly at small Pe. For monodisperse suspensions, the

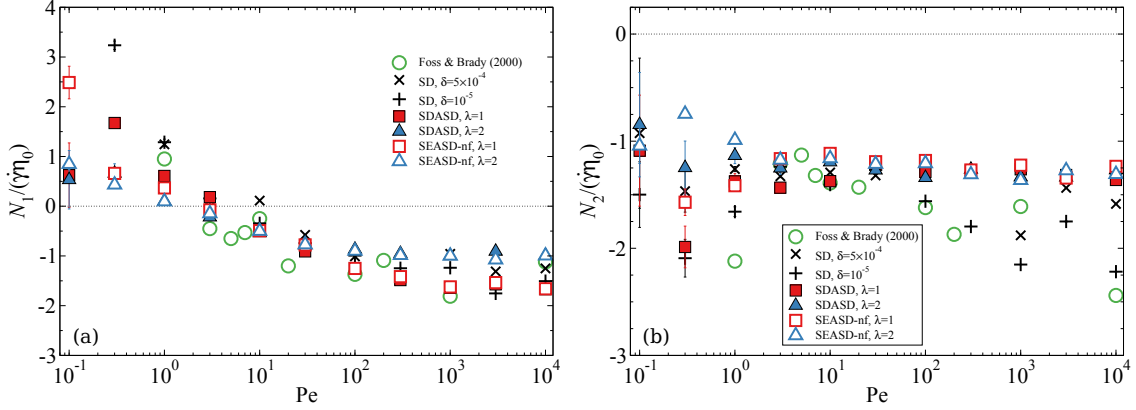


Figure 16: (Color online) The normal stress differences: (a) the first normal stress difference N_1 and (b) the second normal stress difference N_2 as functions of Péclet number Pe . The volume fraction is $\phi = 0.45$ in both cases and the bidisperse composition is $\lambda = 2$ and $y_2 = 0.5$.

excluded volume parameter δ has little effect on N_1 or N_2 , as there lacks a qualitative difference for the SD results at $\delta = 5 \times 10^{-4}$ and 10^{-5} in Fig. 16. These SD results in general agree with the data of Foss & Brady [72] when $Pe > 1$. At smaller Pe , the data exhibit large errors due to fluctuations in Brownian stresses, making quantitative comparisons difficult.

In Fig. 16 the SEASD results at $\lambda = 1$ follow the SD data with the same qualitative behaviors. The differences at low Pe is likely associated with the difficulties in measuring the fluctuating Brownian normal stresses. In addition, the SEASD results show clearer trends at high Pe thanks to larger system sizes: both N_1 and N_2 asymptote toward constant values with increasing Pe . Particle size polydispersity weakens the influences of Pe on the first normal stress difference N_1 . In Fig. 16a, the bidisperse N_1 are less sensitive to Pe compared to the monodisperse case, and as $Pe \rightarrow \infty$, the bidisperse N_1 asymptotes towards a negative value with a smaller magnitude. On the other hand, the size polydispersity has little effect on the second normal stress N_2 , as the bidisperse N_2 almost overlaps the monodisperse N_2 , especially at large Pe .

The SEASD-nf and the SEASD results agree satisfactorily when $Pe \geq 10$ for both the monodisperse and bidisperse suspensions. As expected, larger differences are found at low Pe , as the mean-field Brownian approximation in SEASD-nf explicitly removes the anisotropy in the far-field mobility tensor. However, the SEASD-nf results still capture the qualitative aspect of N_1 and N_2 even in the low Pe limit.

6.3.4. Species stress distribution

Stress distributions across different species are key to understand the phenomena of particle migration and segregation in polydisperse suspensions [95], and are presently only accessible from simulations. Fig. 17 presents the stress distribution, expressed as the stress fraction taken up by the small particles (species 1), as functions of Pe for bidisperse suspensions with $\phi = 0.45$, $\lambda = 2$, and $y_2 = 0.5$. Fig. 17a shows various shear stress fractions. In terms of the definitions in Eqs. (64)–(67), σ_1/σ (circles), σ_1^B/σ^B (squares), and σ_1^E/σ^E (triangles) in Fig. 17a are

$$\sigma_1/\sigma = x_1 \langle \boldsymbol{\Sigma} \rangle_{1,xy} / \langle \boldsymbol{\Sigma} \rangle_{xy}, \quad \sigma_1^B/\sigma^B = x_1 \langle \boldsymbol{S}^B \rangle_{1,xy} / \langle \boldsymbol{S}^B \rangle_{xy}, \quad \text{and} \quad \sigma_1^E/\sigma^E = x_1 \langle \boldsymbol{S}^E \rangle_{1,xy} / \langle \boldsymbol{S}^E \rangle_{xy}, \quad (93)$$

where $\langle \cdot \rangle_\alpha$ indicates averaging with respect to species α . Fig. 17b presents various normal stress fractions. The normal stress fractions S_1/S (circles), S_1^B/S^B (squares), and S_1^E/S^E (triangles) in Fig. 17b are similarly defined as

$$S_1/S = x_1 (\boldsymbol{I} : \langle \boldsymbol{\Sigma} \rangle_1) / (\boldsymbol{I} : \langle \boldsymbol{\Sigma} \rangle), \quad S_1^B/S^B = x_1 (\boldsymbol{I} : \langle \boldsymbol{S}^B \rangle_1) / (\boldsymbol{I} : \langle \boldsymbol{S}^B \rangle), \quad \text{and} \quad S_1^E/S^E = x_1 (\boldsymbol{I} : \langle \boldsymbol{S}^E \rangle_1) / (\boldsymbol{I} : \langle \boldsymbol{S}^E \rangle). \quad (94)$$

In both figures, the SEASD results are shown in filled symbols and the SEASD-nf results are shown in open symbols.

Fig. 17a illustrates that the total shear stress is roughly equally partitioned between the two species, and the fraction σ_1/σ is almost constant with respect to Pe . This is largely because the flow shear stress fraction σ_1^E/σ^E is insensitive to Pe . The Brownian shear stress fraction σ_1^B/σ^B , on the other hand, exhibits weak Pe dependence: the ratio σ_1^B/σ^B




Figure 17: (Color online) The fraction of stresses taken up by the small particles (species 1) in a bidisperse suspension: (a) the fraction of the shear stress and (b) the fraction of the normal stress. The stress fractions are shown as functions of Pe. The composition of the bidisperse hard-sphere suspension is $\phi = 0.45$, $\lambda = 2$, and $y_2 = 0.5$.

increases with Pe from less than 0.45 at Pe = 0.1 to close to 0.6 at Pe = 100. At higher Pe, the Brownian stress fraction shows large fluctuations, also due to the difficulties associated with the anisotropic structures. However, in this limit, the Brownian contribution to the total stress is small, and the large fluctuations in Fig. 17a is inconsequential. On the other hand, the total normal stress fraction S_1/S in Fig. 17b shows stronger Pe dependency, and it decreases from 0.6 at Pe = 0.1 to 0.45 at Pe = 10⁴. Contrary to shear stress distributions in Fig. 17a, the Brownian normal stress distribution S_1^B/S^B is almost constant at 0.6, but S_1^E/S^E increases from 0.3 at Pe = 0.1 and asymptotes towards 0.45 as Pe $\rightarrow \infty$. Since the Brownian stresslet dominates at low Pe and the flow stresslet dominates at high Pe, the normal stress distributions in Fig. 17b are distinctively affected by both the flow and the Brownian contributions. Fig. 17 demonstrates that both the shear and the normal stresses in bidisperse suspensions are distributed based on the species volume and the distribution weakly depends on Pe. This is a useful insight for modelling polydisperse systems.

The stress distributions from SEASD-nf accurately capture the SEASD results except the Brownian shear stress distribution σ_1^B/σ^B at high Pe in Fig. 17a, where the SEASD-nf results is slightly lower. This difference, however, is expected since the mean-field Brownian approximation ignores the structural anisotropy in the suspension. Moreover, the discrepancies are only evident at Péclet numbers where the Brownian stress does not affect the overall suspension rheology. From this perspective, the overall quality of the SEASD-nf approximation is deemed satisfactory.

6.3.5. Long-time diffusion

An important characterization of the overall suspension dynamics is the translational long-time self-diffusivities. The long-time limit refers to a time scale $t \gg \tau_D$, where, recall that, $\tau_D = 6\pi\eta_0 a_p^3/k_B T$ is the single particle diffusive time scale. In this limit, the particle movement is diffusive due to extensive interactions with their neighbors. The corresponding diffusivities are obtained from the particle mean-square displacement. In the velocity gradient and the vorticity directions, these self-diffusivities are respectively defined as

$$d_{\infty,\alpha}^{t,yy} = \lim_{t \rightarrow \infty} \frac{1}{2} d\langle (\Delta y)^2 \rangle_\alpha / dt \quad \text{and} \quad d_{\infty,\alpha}^{t,zz} = \lim_{t \rightarrow \infty} \frac{1}{2} d\langle (\Delta z)^2 \rangle_\alpha / dt, \quad (95)$$

where Δy and Δz are the particle trajectory fluctuations in y- and z-directions. Fig. 18a and 18b respectively present the long-time diffusivities $d_{\infty,\alpha}^{t,yy}$ and $d_{\infty,\alpha}^{t,zz}$ as functions of Péclet number. The monodisperse results are shown in squares. For bidisperse suspensions, the small and the large particle long-time self-diffusivities are presented in triangles and circles, respectively. For comparison, Fig. 18 also shows the results from Foss & Brady [72] in crosses. Moreover, the SEASD and the SEASD-nf results are shown in filled and open symbols, respectively.

For monodisperse suspensions in Fig. 18, both $d_{\infty}^{t,yy}$ and $d_{\infty}^{t,zz}$ grow with Pe due to the imposed shear flow, with the velocity direction diffusivity $d_{\infty}^{t,yy}$ slightly higher. At low Pe, $d_{\infty}^{t,yy}$ and $d_{\infty}^{t,zz}$ grow weakly with Pe, and at large Pe, both diffusivities are proportional to Pe. The SEASD results is consistent with the SD results of Foss & Brady [72] at




Figure 18: (Color online) The species long-time self-diffusivities: (a) the velocity gradient direction diffusivity $d_{\infty, \alpha}^{t, yy} / d_{0, \alpha}^t$ and (b) the vorticity direction diffusivity $d_{\infty, \alpha}^{t, zz} / d_{0, \alpha}^t$ of monodisperse and bidisperse hard-sphere suspensions as functions of Pe. The volume fraction is $\phi = 0.45$ for both cases, and the bidisperse composition is $\lambda = 2$ and $y_2 = 0.5$.

intermediate Pe. The differences at large and small Pe are most likely due to the system size, as in this work $N = 200$ while in Foss & Brady [72] $N = 27$. For bidisperse suspensions, the long-time self-diffusivities $d_{\infty, \alpha}^{t, yy}$ and $d_{\infty, \alpha}^{t, zz}$ for both species exhibit similar Pe dependencies as the monodisperse case. However, introducing a second species to the suspension apparently enhances the long-time self-diffusivities of both species, particularly at high Pe. This mutual diffusivity enhancement is in contrast to the short-time diffusivities in Fig. 8a, where at $\phi = 0.45$, the small particle diffusivity enhancement is always accompanied by the large particle diffusivity suppression. Moreover, the diffusivity enhancement in y -direction is stronger than those in z -direction.

In Fig. 18 the diffusivities from SEASD-nf in general agree with the SEASD results for both monodisperse and bidisperse suspensions. At low Pe, the SEASD-nf diffusivity is lower, particularly for the large particles. The agreement between SEASD and SEASD-nf improves with increasing Pe due to the reduced influences of Brownian motion.

6.3.6. Suspension structures

Finally, we examine the structures of sheared bidisperse suspensions via the projections of the partial pair-distribution functions $g_{\alpha\beta}(\mathbf{r})$, which are defined as the conditional probability of finding another particle in species β given a particle of species α , *i.e.*,

$$g_{\alpha\beta}(\mathbf{r}) = \frac{1}{n_{\alpha}n_{\beta}} \left\langle \sum'_{\substack{i \in \alpha, \\ j \in \beta}} \frac{1}{V} \delta(\mathbf{r} - \mathbf{r}_i + \mathbf{r}_j) \right\rangle. \quad (96)$$

They are related to the pair-distribution function $g(\mathbf{r})$ through

$$g(\mathbf{r}) = \sum_{\alpha, \beta} x_{\alpha} x_{\beta} g_{\alpha\beta}(\mathbf{r}). \quad (97)$$

Fig. 19, 20, and 21 present projections of $g(\mathbf{r})$ and $g_{\alpha\beta}(\mathbf{r})$ on the velocity-velocity gradient (xy -) plane, the velocity-vorticity (xz -) plane, and the velocity gradient-vorticity (yz -) plane, respectively, at selected Péclet numbers. These figures are based on particle trajectories from SEASD simulation, and are indistinguishable from the SEASD-nf results.

Fig. 19 clearly displays the structural anisotropy caused by the shear flow in the xy -plane, characterized by the distortion of the otherwise isotropic pair-distribution rings. With increasing Pe, the overall pair-distribution function $g(\mathbf{r})$ shows an accumulation of neighboring particles in the compressional quadrant. This is indicated by the brightening and thinning of the rings at $2a_1$, $a_1 + a_2$, and $2a_2$, corresponding to the particle pairs of two small particles, a large and a small particle, and two large particles, respectively. Meanwhile, the particle pairs are depleted in the extensional quadrant.




Figure 19: (Color online) The velocity-velocity gradient (xy -) plane projection of the pair-distribution function $g(\mathbf{r})$ and the partial pair-distribution functions $g_{\alpha\beta}(\mathbf{r})$ at various Pe for bidisperse suspensions with $\phi = 0.45$, $\lambda = 2$, and $y_2 = 0.5$.




Figure 20: (Color online) The velocity-vorticity (xz -) plane projection of the pair-distribution function $g(\mathbf{r})$ and the partial pair-distribution functions $g_{\alpha\beta}(\mathbf{r})$ at various Pe for bidisperse suspensions with $\phi = 0.45$, $\lambda = 2$, and $y_2 = 0.5$.




Figure 21: (Color online) The velocity gradient-vorticity (yz -) plane projection of the pair-distribution function $g(\mathbf{r})$ and the partial pair-distribution functions $g_{\alpha\beta}(\mathbf{r})$ at various Pe for bidisperse suspensions with $\phi = 0.45$, $\lambda = 2$, and $y_2 = 0.5$.

Specific changes in different types of particle pairs are revealed by examining the corresponding partial pair-distribution function $g_{\alpha\beta}(\mathbf{r})$ in Fig. 19. The distribution of the small-small particle pairs is presented in $g_{11}(\mathbf{r})$. Similarly to $g(\mathbf{r})$, $g_{11}(\mathbf{r})$ is increasingly distorted and compressed in the compressional quadrant with increasing Pe , forming a boundary layer. At higher Pe , the pair structure remain approximately unchanged. In the extensional quadrant, the pair breakup point shifts from the extensional axis towards the velocity (x -) direction due to the lubrication interactions, with a clear tail of high probability outlining the trajectory of small-small pair disengagement. The distribution of the small-large particle pairs in $g_{12}(\mathbf{r})$ shows a similar structural distortion in the compressional quadrant with increasing Pe . Moreover, in the extensional quadrant, the trajectory of particle disengagement is more diffusive compared to $g_{11}(\mathbf{r})$ at the same Pe . This suggests that particle movement in bidisperse suspensions are facilitated by the breakup of small-large particle pairs, and partially explains the mutual enhancement of long-time self-diffusivity in Fig. 18. For the distribution of large-large particle pairs, $g_{22}(\mathbf{r})$ also exhibits anisotropy with increasing Pe in Fig. 19. However, due to the limited particle number, information beyond the first coordinate shell is difficult to analyze.

Fig. 20 displays the total and partial pair-distribution function projections in the xz -plane. Unlike the xy -plane projections in Fig. 19 which exhibits strong anisotropy, the suspension structures here are less sensitive to Pe . With increasing Pe , the particles are compressed towards each other, which is evidenced by the thinning and brightening of the first coordinate shells. More interestingly, at higher $Pe \geq 100$, $g_{12}(\mathbf{r})$ shows a belt of particle enrichment along the flow direction, while $g_{11}(\mathbf{r})$ and $g_{22}(\mathbf{r})$ exhibit a corresponding particle depletion. This indicates that the small-large pairs are preferred in the xz -plane, and that the shear flow promotes species mixing in the flow direction.

Fig. 21 shows the projection of $g(\mathbf{r})$ and $g_{\alpha\beta}(\mathbf{r})$ in the yz -plane. With increasing Pe , the shear flow also compresses the particle pairs in this plane without apparent anisotropy. Note that even at $Pe = 10^4$, the suspension does not exhibit string ordering [96] due to the HIs. The lack of structural formation is also confirmed by the continuous increase of the long-time self-diffusivities $d_{\infty,\alpha}^{l,yy}$ and $d_{\infty,\alpha}^{l,yy}$ with Pe in Fig. 18.

7. Conclusions

In this work we presented the Spectral Ewald Accelerated Stokesian Dynamics (SEASD) for dynamic simulations of polydisperse colloidal suspensions. Using the framework of Stokesian Dynamics (SD), SEASD can accurately and rapidly compute HIs in dense polydisperse suspensions. Other features of SEASD include (i) direct inclusion of the

solvent compressibility and pressure evaluations; *(ii)* the use of the Spectral Ewald (SE) method for accurate mobility computation with flexible error control; *(iii)* a far-field preconditioner to accelerate the convergence of the nested iterative scheme; *(iv)* GPGPU accelerated mobility evaluation for almost an order of magnitude speed improvement; and *(v)* the incorporation of SEASD-nf, an extension of the mean-field Brownian approximation of Banchio & Brady [39] to polydisperse suspensions.

We extensively discussed the accuracy of mobility computation using the SE method, established the baseline for parameter selection, and demonstrated the adequate accuracy in the GPU single precision (SP) mobility computation. We found that compared to the full SEASD computations, SEASD-nf can achieve significant speedup without substantially sacrificing accuracy. Indeed, for all the dynamic simulations in this work, the SEASD and SEASD-nf results agree satisfactorily. In addition, we verified the $O(N \log N)$ computational scaling of SEASD and SEASD-nf in dynamic simulations.

We rigorously validated SEASD and SEASD-nf for monodisperse and bidisperse colloidal suspensions via: *(i)* the short-time transport properties, *(ii)* the equilibrium osmotic pressure and viscoelastic moduli, and *(iii)* the steady Brownian shear rheology at $\phi = 0.45$. For *(i)*, the SEASD diffusivities and shear viscosity agree with the conventional SD calculations. The SEASD sedimentation velocity differ qualitatively from the SD results due to the absence of a mean-field quadrupole term in the mobility computation. For the bulk viscosity computation, different procedures to eliminate the spurious HIs lead to slight differences in the SEASD and the SD results. In *(ii)*, SEASD and SEASD-nf reproduced the equilibrium suspension osmotic pressure for monodisperse and bidisperse suspensions within the error tolerance, with the SEASD data higher. For the steady shear rheology in *(iii)*, the agreement between SEASD-nf and SEASD is satisfactory in the suspension mechanics, dynamics, and structures. Moreover, we found that the particle size polydispersity reduces the suspension viscosity and osmotic pressure, and enhances the long-time translational self-diffusivities of both species. Our rheological simulations also improve our understanding on the structure, dynamics, and rheology of polydisperse suspensions.

The SEASD and SEASD-nf developed in this work are important tools for studying dynamics of dense, polydisperse colloidal suspensions, and have significantly extended the parameter space accessible to computational studies. For example, they can provide otherwise inaccessible details on a wide range of experimental observations including the yielding phenomena in glass rheology and the continuous and discontinuous shear-thickening.

Finally, through SEASD and SEASD-nf we demonstrated the generality and versatility of the SD framework, particularly the splitting of the far- and near-field interactions: with a suitable far-field computation, the lubrication interactions can be added pairwise for free. We believe that many far-field HI computational methods can and should be used with the SD framework to expand their accessible parameter range, particularly for dense systems.

Acknowledgements

We thank Wen Yan for helpful discussions on GPGPU programming and particle mesh techniques. M.W. gratefully acknowledges supports from the Natural Sciences and Engineering Research Council of Canada (NSERC) by a Postgraduate Scholarship (PGS), and the National Science Foundation (NSF) grant CBET-1337097.

References

- [1] W. B. Russel, D. A. Saville, W. R. Schowalter, Colloidal Dispersions, Cambridge University Press, 1991.
- [2] J. S. Chong, E. B. Christiansen, A. D. Baer, Rheology of concentrated suspensions, J. Appl. Polym. Sci. 15 (1971) 2007.
- [3] B. E. Rodriguez, E. W. Kaler, M. S. Wolfe, Binary mixtures of monodisperse latex dispersions. 2. viscosity, Langmuir 8 (1992) 2382.
- [4] T. Shikata, H. Niwa, Y. Morishima, Viscoelastic behavior of bimodal suspensions, J. Rheo. 42 (1998) 765.
- [5] T. Sentjabrskaja, E. Babaliari, J. Hendricks, M. Laurati, G. Petekidis, S. U. Egelhaaf, Yielding of binary colloidal glasses, Soft Matter 9 (2013) 4524.
- [6] D. Semwogerere, E. R. Weeks, Shear-induced particle migration in binary colloidal suspensions, Phys. Fluids 20 (2008) 043306.
- [7] L. Durlofsky, J. F. Brady, G. Bossis, Dynamic simulation of hydrodynamically interacting particles, J. Fluid Mech. 180 (1987) 21.
- [8] G. K. Batchelor, Sedimentation in a dilute polydisperse system of interacting spheres. Part 1. General theory, J. Fluid Mech. 119 (1982) 379.
- [9] N. J. Wagner, A. T. J. M. Woutersen, The viscosity of bimodal and polydisperse suspensions of hard spheres in the dilute limit, J. Fluid Mech. 278 (1994) 267.
- [10] H. Zhang, G. Nägele, Tracer-diffusion in binary colloidal hard-sphere suspensions, J. Chem. Phys. 117 (2002) 5908.
- [11] M. Wang, J. F. Brady, Short-time transport properties of bidisperse suspensions and porous media: a Stokesian Dynamics study, J. Chem. Phys. 142 (2015) 094901.

- [12] M. Wang, M. Heinen, J. F. Brady, Short-time diffusion in concentrated bidisperse hard-sphere suspensions, *J. Chem. Phys.* 142 (2015) 064905.
- [13] T. Ando, J. Skolnick, Crowding and hydrodynamic interactions likely dominate in vivo macromolecular motion, *Proc. Natl. Acad. Sci. U. S. A.* 107 (2010) 18457.
- [14] A. J. C. Ladd, Hydrodynamic transport coefficients of random dispersions of hard spheres, *J. Chem. Phys.* 93 (1990) 3484.
- [15] B. Cichocki, B. U. Felderhof, K. Hinsen, E. Wajnryb, J. Blawdziewicz, Friction and mobility of many spheres in Stokes flow, *J. Chem. Phys.* 100 (1994) 3780.
- [16] G. K. Youngren, A. Acrivos, Stokes flow past a particle of arbitrary shape: a numerical method of solution, *J. Fluid Mech.* 69 (1975) 377.
- [17] C. Pozrikidis, *Boundary Integral and Singularity Methods for Linearized Viscous Flow*, Cambridge University Press, 1992.
- [18] L. B. Lucy, A numerical approach to the fission hypothesis, *Astronomical Journal* 82 (1977) 1013.
- [19] G. R. Liu, M. B. Liu, *Smoothed Particle Hydrodynamics: A Meshfree Particle Method*, World Scientific, 2003.
- [20] P. J. Hoogerbrugge, J. M. V. A. Koelman, Simulating microscopic hydrodynamic phenomena with Dissipative Particle Dynamics, *Europhys. Lett.* 19 (1992) 155.
- [21] B. Dünweg, A. J. C. Ladd, Lattice Boltzmann simulations of soft matter systems, *Adv. Polym. Sci.* 221 (2009) 89.
- [22] M. Abbas, E. Climent, O. Simonin, M. R. Maxey, Dynamics of bidisperse suspensions under Stokes flows: Linear shear flow and sedimentation, *Phys. Fluids* 18 (2006) 121504.
- [23] B. Cichocki, M. L. Ekiel-Jeżewska, E. Wajnryb, Lubrication corrections for three-particle contribution to short-time self-diffusion coefficients in colloidal dispersions, *J. Chem. Phys.* 111 (1999) 3265.
- [24] B. Cichocki, R. B. Jones, R. Kutteh, E. Wajnryb, Friction and mobility for colloidal spheres in Stokes flow near a boundary: The multipole method and applications, *J. Chem. Phys.* 112 (2000) 2548.
- [25] K. Yeo, M. R. Maxey, Simulation of concentrated suspensions using the force-coupling method, *J. Comput. Phys.* 229 (2010) 2401.
- [26] N.-Q. Nguyen, A. J. C. Ladd, Lubrication corrections for lattice-Boltzmann simulations of particle suspensions, *Phys. Rev. E* 66 (2002) 046708.
- [27] S. Gallier, E. Lemaire, L. Lobry, F. Peters, A fictitious domain approach for the simulation of dense suspensions, *J. Comput. Phys.* 256 (2014) 367.
- [28] A. Lefebvre-Lepot, B. Merlet, T. N. Nguyen, An accurate method to include lubrication forces in numerical simulations of dense Stokesian suspensions, *J. Fluid Mech.* 769 (2015) 369.
- [29] D. Lindbo, A.-K. Tornberg, Spectrally accurate fast summation for periodic Stokes potentials, *J. Comput. Phys.* 229 (2010) 8994.
- [30] D. Lindbo, A.-K. Tornberg, Spectral accuracy in fast Ewald-based methods for particle simulations, *J. Comput. Phys.* 230 (2011) 8744.
- [31] L. af Klingteberg, A.-K. Tornberg, Fast Ewald summation for Stokesian particle suspensions, *Int. J. Numer. Methods Fluids* 76 (2014) 669.
- [32] T. Darden, D. York, L. Pedersen, Particle mesh Ewald: An Nlog(N) method for Ewald sums in large systems, *J. Chem. Phys.* 98 (1993) 10089.
- [33] U. Essmann, L. Perera, M. L. Berkowitz, T. Darden, H. Lee, L. G. Pedersen, A smooth particle mesh Ewald method, *J. Chem. Phys.* 103 (1995) 8577.
- [34] L. Greengard, V. Rokhlin, A fast algorithm for particle simulations, *J. Comput. Phys.* 73 (1987) 325.
- [35] K. Ichiki, Improvement of the Stokesian Dynamics method for systems with a finite number of particles, *J. Fluid Mech.* 452 (2002) 231.
- [36] A. Sierou, J. F. Brady, Accelerated Stokesian Dynamics simulations, *J. Fluid Mech.* 448 (2001) 115.
- [37] D. Saintillan, E. Darve, E. S. G. Shaqfeh, A smooth particle-mesh Ewald algorithm for Stokes suspension simulations: The sedimentation of fibers, *Phys. Fluids* 17 (2005) 033301.
- [38] D. L. Ermak, J. A. McCammon, Brownian dynamics with hydrodynamic interactions, *J. Chem. Phys.* 69 (1978) 1352.
- [39] A. J. Banchio, J. F. Brady, Accelerated Stokesian Dynamics: Brownian motion, *J. Chem. Phys.* 118 (2003) 10323.
- [40] J. W. Swan, J. F. Brady, Colloidal dispersions deformed by a steady shear stress, Submitted.
- [41] A. Kumar, J. J. L. Higdon, Origins of the anomalous stress behavior in charged colloidal suspensions under shear, *Phys. Rev. E* 82 (2010) 051401.
- [42] L. D. Landau, E. M. Lifshitz, *Fluid Mechanics*, Butterworth-Heinemann, 1987.
- [43] A. J. C. Ladd, H. Gang, J. X. Zhu, D. A. Weitz, Temporal and spatial dependence of hydrodynamic correlations: Simulation and experiment, *Phys. Rev. E* 52 (1995) 6550.
- [44] E. E. Keaveny, Fluctuating force-coupling method for simulations of colloidal suspensions, *J. Comput. Phys.* 269 (0) (2014) 61.
- [45] S. Delong, F. B. Usabiaga, R. Delgado-Buscalioni, B. E. Griffith, A. Donev, Brownian dynamics without Green's functions, *J. Chem. Phys.* 140 (13) (2014) –.
- [46] M. Kopp, F. Höfling, GPU-accelerated simulation of colloidal suspensions with direct hydrodynamic interactions, *Eur. Phys. J.-Spec. Top.* 210 (2012) 101.
- [47] N. Ganesan, M. Taufer, B. Bauer, S. Patel, FENZI: Gpu-enabled Molecular Dynamics simulations of large membrane regions based on the CHARMM force field and PME, in: *IEEE IPDPSW*, 2011, p. 472.
- [48] M. J. Harvey, G. De Fabritiis, An implementation of the smooth particle mesh Ewald method on GPU hardware, *J. Chem. Theory Comput.* 5 (2009) 2371.
- [49] F. Boyer, E. Guazzelli, O. Pouliquen, Unifying suspension and granular rheology, *Phys. Rev. Lett.* 107 (2011) 188301.
- [50] J. F. Brady, A. S. Khair, M. Swaroop, On the bulk viscosity of suspensions, *J. Fluid Mech.* 554 (2006) 109.
- [51] J. F. Brady, Brownian motion, hydrodynamics, and the osmotic pressure, *J. Chem. Phys.* 98 (1993) 3335.
- [52] J. F. Morris, F. Boulay, Curvilinear flows of noncolloidal suspensions: The role of normal stresses, *J. Rheo.* 43 (1999) 1213.
- [53] Y. Yurkovetsky, J. F. Morris, Particle pressure in sheared Brownian suspensions, *J. Rheo.* 52 (2008) 141.
- [54] S. Kim, S. J. Karrila, *Microhydrodynamics*, Dover, 2005.
- [55] P. Szymczak, B. Cichocki, A diagrammatic approach to response problems in composite systems, *J. Stat. Mech. Theor. Exp.* (2008) P01025.
- [56] W. van Saarloos, P. Mazur, Many-sphere hydrodynamic interactions and mobilities in a suspension, *Physica* 115A (1982) 21.
- [57] A. J. C. Ladd, Hydrodynamic interactions in a suspension of spherical particles, *J. Chem. Phys.* 88 (1988) 5051.
- [58] R. Schmitz, B. U. Felderhof, Creeping flow about a spherical particle, *Physica* 113A (1982) 90.

- [59] H. Hasimoto, On the periodic fundamental solutions of the Stokes equations and their application to viscous flow past a cubic array of spheres, *J. Fluid Mech.* 5 (1959) 317.
- [60] C. W. J. Beenakker, Ewald sum of the Rotne-Prager tensor, *J. Chem. Phys.* 85 (1986) 1581.
- [61] J. W. Swan, J. F. Brady, Particle motion between parallel walls: Hydrodynamics and simulation, *Phys. Fluids* 22 (2010) 103301.
- [62] J. W. Swan, J. F. Brady, The hydrodynamics of confined dispersions, *J. Fluid Mech.* 687 (2011) 254.
- [63] J. F. Brady, R. J. Phillips, J. C. Lester, G. Bossis, Dynamic simulation of hydrodynamically interacting suspensions, *J. Fluid Mech.* 195 (1988) 257.
- [64] D. J. Jeffrey, J. F. Morris, J. F. Brady, The pressure moments for two rigid spheres in low-Reynolds-number flow, *Phys. Fluids A* 5 (1993) 10.
- [65] A. S. Khair, M. Swaroop, J. F. Brady, A new resistance function for two rigid spheres in a uniform compressible low-Reynolds-number flow, *Phys. Fluids* 18 (2006) 043102.
- [66] H. G. Petersen, Accuracy and efficiency of the particle mesh Ewald method, *J. Chem. Phys.* 103 (1995) 3668.
- [67] M. Deserno, C. Holm, How to mesh up ewald sums. I. A theoretical and numerical comparison of various particle mesh routines, *J. Chem. Phys.* 109 (1998) 7678.
- [68] L. Greengard, J.-Y. Lee, Accelerating the nonuniform Fast Fourier Transform, *SIAM Rev.* 46 (2004) 443.
- [69] M. P. Allen, D. J. Tildesley, *Computer Simulation of Liquids*, Clarendon Press, 1991.
- [70] R. W. Hockney, J. W. Eastwood, *Computer Simulation Using Particles*, Adam Hilger, 1989.
- [71] H. J. Wilson, Stokes flow past three spheres, *J. Comput. Phys.* 245 (2013) 302.
- [72] D. R. Foss, J. F. Brady, Structure, diffusion and rheology of Brownian suspensions by Stokesian Dynamics simulation, *J. Fluid Mech.* 407 (2000) 167.
- [73] A. Sierou, J. F. Brady, Rheology and microstructure in concentrated noncolloidal suspensions, *J. Rheo.* 46 (2002) 1031.
- [74] Y. Saad, *Iterative Methods for Sparse Linear Systems*, SIAM, 2003.
- [75] D. J. Jeffrey, Y. Onishi, Calculation of the resistance and mobility functions for two unequal rigid spheres in low-Reynolds-number flow, *J. Fluid Mech.* 139 (1984) 261.
- [76] D. J. Jeffrey, The calculation of the low Reynolds number resistance functions for two unequal spheres, *Phys. Fluids A* 4 (1992) 16.
- [77] D. J. Evans, G. Morriss, *Statistical Mechanics of Nonequilibrium Liquids*, Cambridge University Press, 2008.
- [78] G. Bossis, J. F. Brady, Self-diffusion of Brownian particles in concentrated suspensions under shear, *J. Chem. Phys.* 87 (1987) 5437.
- [79] R. M. Jendrejack, M. D. Graham, J. J. de Pablo, Hydrodynamic interactions in long chain polymers: Application of the Chebyshev polynomial approximation in stochastic simulations, *J. Chem. Phys.* 113 (2000) 2894.
- [80] N. J. Higham, *Functions of Matrices: Theory and Computation*, SIAM, 2008.
- [81] S. K. Boyd, J. Baglama, E. J. Allen, Numerical approximation of the product of the square root of a matrix with a vector, *Linear Algebra Appl.* 310 (2000) 167.
- [82] L. N. Trefethen, *Spectral Methods in MATLAB*, SIAM, 2000.
- [83] M. Wang, J. F. Brady, The suspension balance model revisited: Revisited, In preparation.
- [84] B. D. Lubachevsky, F. H. Stillinger, Geometric properties of random disk packings, *J. Stat. Phys.* 60 (1990) 561.
- [85] M. Skoge, A. Donev, F. H. Stillinger, S. Torquato, Packing hyperspheres in high-dimensional Euclidean spaces, *Phys. Rev. E* 74 (2006) 041127.
- [86] J. K. Percus, G. J. Yevick, Analysis of classical statistical mechanics by means of collective coordinates, *Phys. Rev.* 110 (1958) 1.
- [87] J. L. Lebowitz, Exact solution of generalized Percus-Yevick equation for a mixture of hard spheres, *Phys. Rev.* 133 (1964) A895.
- [88] N. W. Ashcroft, D. C. Langreth, Structure of binary liquid mixtures. I, *Phys. Rev.* 156 (1967) 685.
- [89] N. W. Ashcroft, D. C. Langreth, Errata, *Phys. Rev.* 166 (1968) 934.
- [90] R. J. Phillips, J. F. Brady, G. Bossis, Hydrodynamic transport properties of hard-sphere dispersions. I. Suspensions of freely mobile particles, *Phys. Fluids* 31 (1988) 3462.
- [91] M. Swaroop, The bulk viscosity of suspensions, Ph.D. thesis, California Institute of Technology (2010).
- [92] G. A. Mansoori, N. S. Carnahan, K. E. Starling, T. W. Leland, Equilibrium thermodynamic properties of the mixture of hard spheres, *J. Chem. Phys.* 54 (1971) 1523.
- [93] J. F. Brady, The rheological behavior of concentrated colloidal dispersions, *J. Chem. Phys.* 99 (1993) 567.
- [94] G. Nägele, J. Bergenholtz, Linear viscoelasticity of colloidal mixtures, *J. Chem. Phys.* 108 (1998) 9893.
- [95] H. M. Vollebregt, R. G. M. van der Sman, R. M. Boom, Model for particle migration in bidisperse suspensions by use of effective temperature, *Faraday Discuss.* 158 (2012) 89.
- [96] D. R. Foss, J. F. Brady, Brownian Dynamics simulation of hard-sphere colloidal dispersions, *J. Rheo.* 44 (2000) 629.