

# Python-JAX-based fast Stokesian dynamics

Kim William Torre<sup>1\*</sup>, Raoul D. Schram<sup>2†</sup> and Joost de Graaf<sup>1‡</sup>

<sup>1</sup> Institute for Theoretical Physics, Center for Extreme Matter and Emergent Phenomena,  
Utrecht University, Princetonplein 5, 3584 CC Utrecht, The Netherlands

<sup>2</sup> Research and Data Management Services, Information and Technology Services,  
Utrecht University, Heidelberglaan 8, 3584 CS Utrecht, The Netherlands

\* [k.w.torre@uu.nl](mailto:k.w.torre@uu.nl), † [r.d.schram@uu.nl](mailto:r.d.schram@uu.nl), ‡ [j.degraaf@uu.nl](mailto:j.degraaf@uu.nl),

## Abstract

Stokesian Dynamics (SD) is a powerful computational framework for simulating the motion of particles in a viscous Newtonian fluid under Stokes-flow conditions. Traditional SD implementations can be computationally expensive as they rely on the inversion of large mobility matrices to determine hydrodynamic interactions. Recently, however, the simulation of thermalized systems with large numbers of particles has become feasible [Fiore and Swan, *J. Fluid Mech.* 878, 544 (2019)]. Their “fast Stokesian dynamics” (FSD) method leverages a saddle-point formulation to ensure overall scaling of the algorithm that is linear in the number of particles  $\mathcal{O}(N)$ ; performance relies on dedicated graphics-processing-unit computing. Here, we present a different route toward implementing FSD, which instead leverages the Just-in-Time (JIT) compilation capabilities of Google JAX. We refer to this implementation as JFSD and perform benchmarks on it to verify that it has the right scaling and is sufficiently fast by the standards of modern computational physics. In addition, we provide a series of physical test cases that help ensure accuracy and robustness, as the code undergoes further development. Thus, JFSD is ready to facilitate the study of hydrodynamic effects in particle suspensions across the domains of soft, active, and granular matter.



Copyright K. W. Torre *et al.*

This work is licensed under the Creative Commons

[Attribution 4.0 International License](https://creativecommons.org/licenses/by/4.0/).

Published by the SciPost Foundation.

Received 2025-03-12

Accepted 2025-06-05

Published 2025-06-30

doi:[10.21468/SciPostPhysCodeb.56](https://doi.org/10.21468/SciPostPhysCodeb.56)



Check for updates

**This publication is part of a bundle:** Please cite both the article and the release you used.

DOI	Type
<a href="https://doi.org/10.21468/SciPostPhysCodeb.56">doi:10.21468/SciPostPhysCodeb.56</a>	Article
<a href="https://doi.org/10.21468/SciPostPhysCodeb.56-r0.2">doi:10.21468/SciPostPhysCodeb.56-r0.2</a>	Codebase release

## Contents

<b>1</b>	<b>Introduction</b>	<b>2</b>
<b>2</b>	<b>Method</b>	<b>3</b>

<b>3</b>	<b>Method validations</b>	<b>5</b>
3.1	Sedimenting particles	5
3.2	Hydrodynamic effects under shear flow	6
3.3	Thermal motion	8
3.4	Algorithmic performance and scaling	9
<b>4</b>	<b>Discussion</b>	<b>11</b>
<b>5</b>	<b>Conclusion and outlook</b>	<b>11</b>
	<b>References</b>	<b>12</b>

---

## 1 Introduction

The study of hydrodynamic interactions in particle suspensions has attracted considerable interest in colloid science and fluid mechanics over the past decades [2–4]. The development of Stokesian Dynamics (SD) by Brady and Bossis [5] in the late 1980s provided an accurate framework for simulating the many-body hydrodynamic interactions between spherical particles at zero (low) Reynolds number. That is, in the regime where inertia is completely dominated by friction. The Stokes equations that describe the fluid dynamics of such systems are linear. This is leveraged in SD to divide the mobility problem into a far- and near-field component. The split enables the precise modeling of long-range, many-body hydrodynamic interactions and short-range, lubrication forces [6–8], both of which can strongly influence the behavior of colloidal suspensions [9–19].

However, despite its accuracy, SD did not become as widely adopted as other hydrodynamics methods for particle suspensions [20]. Among these, lattice-based methods for computational fluid dynamics gained considerable traction, including lattice-Boltzmann (LB) [21–25], multi-particle collision dynamics (MPCD) [26, 27] — or stochastic rotation dynamics (SRD) [28] — and Fluid Particle Dynamics (FPD) [29].

Contrasting the features of these methods, we see the following advantages for SD. The algorithm relies on multipole expansion and pairwise approximations to compute hydrodynamic interactions without directly solving for the fluid flow. This means that at low dilution, there is an intrinsic advantage to using SD. However, LB and MPCD are particularly effective for determining fluid-structure interactions involving complex geometries [21–25]. Yet, achieving the same level of accuracy in both many-body and lubrication interactions as SD requires the use of prohibitively fine resolution in LB, leading to higher computational demands [24, 30]. MPCD and SRD operate on similar principles [26–28]; however, they are intrinsically stochastic, with control parameters such as viscosity emerging as collective properties of the system. SD instead provides explicit control over thermal and athermal motion, like LB. Lastly, FPD, as introduced by Tanaka and Araki [29], directly solves the Navier-Stokes equation on a lattice. It treats colloidal particles as regions of higher viscosity fluid with smooth interfacial profiles. This eliminates the need to explicitly resolve solid-fluid boundaries, but it introduces approximations. These depend on the viscosity ratio between the particles and surrounding fluid, as well as on the lattice resolution.

The major bottlenecks for the use of SD were: (i) the computational cost involved in propagating the dynamics and (ii) that it is challenging to implement the algorithm correctly. Advancements following the original implementation [5] have therefore focused on improving

SD’s computational performance [31–33]. These include the incorporation of iterative solvers and preconditioning strategies. Fiore and Swan overcame some of the last remaining hurdles in making SD truly efficient, through their Fast Stokesian Dynamics (FSD) approach [1]. They combined matrix-free techniques with spectral Ewald methods, which allowed them to further reduce computational costs, improve the scaling to  $\mathcal{O}(N)$  overall, and maintain accuracy. They also made their algorithm widely available through integration in the HOOMD [34] simulation package. This means that, provided that complex geometries are not critical to the analysis of a system, FSD provides a competitive alternative to more widely used lattice-based approaches. However, the efficiency obtained in FSD was limited to NVIDIA<sup>®</sup> graphics processing units (GPUs), through an internal reliance on CUDA<sup>®</sup> programming.

In this paper, we present an alternative route toward making FSD more widely available and cross-platform compatible. We have ported the FSD algorithm to a Python-based framework that relies on the Google JAX library [35]. We name this version of SD, Python-JAX Fast Stokesian Dynamics, or Python-JFSD for short. Our implementation ensures compatibility with contemporary CUDA<sup>®</sup> and CuDNN<sup>®</sup> libraries. It also addresses limitations in the original code caused by updates to the HOOMD application programming interface (API). Python-JFSD further offers a user-friendly interface and provides flexibility in choosing between periodic and open boundary conditions for hydrodynamic interactions. By leveraging Google JAX’s Just-In-Time (JIT) compilation [35], the framework supports the efficient addition of new features, ensuring modularity and extensibility.

The remainder of this paper is organized as follows. In Section 2, we first provide a brief overview of the FSD method. Section 3 contains the results of a set of physical unit test that probe the implementation accuracy and benchmark its performances. In Section 4, we discuss the applicability range of JFSD, as well as its strong and weak points. We close with a summary and outlook in Section 5.

## 2 Method

Consider a system comprising  $N$  spherical particles of diameter  $\sigma$ , suspended in a Newtonian solvent with viscosity  $\eta$ . We have a prescribed background flow  $\mathcal{U}^\infty$ , which is a  $6N$ -dimensional vector containing the linear and angular fluid velocities at the position of each particle. We assume viscous forces dominate inertial ones, such that the Reynolds number  $\text{Re} \ll 1$ , and the particles undergo over-damped motion. Under these conditions, the particle dynamics are prescribed by a linear system of equations that couple the dynamical degrees of freedom (forces and torques  $\mathcal{F}$ , and stresses  $\mathcal{S}$ ) to kinematic degrees of freedom (linear and angular velocities  $\mathcal{U}$  and strain rates  $\mathcal{E}$ ) via the grand-resistance matrix  $\mathcal{R}$ . The relation is given by

$$\begin{pmatrix} \mathcal{F}^H \\ \mathcal{S}^H \end{pmatrix} = -\mathcal{R} \cdot \begin{pmatrix} \mathcal{U} - \mathcal{U}^\infty \\ -\mathcal{E}^\infty \end{pmatrix}, \quad (1)$$

where  $\mathcal{U}$  is a  $6N$ -dimensional vector containing the linear and angular velocities of all particles,  $\mathcal{E}^\infty$  is the strain rate of the background flow,  $\mathcal{F}^H$  denotes the hydrodynamic drag forces and torques, and  $\mathcal{S}^H$  is a  $5N$ -dimensional vector of the independent components of the hydrodynamic stresslets.

The generalized velocities of all  $N$  particles can then be expressed as

$$\mathcal{U} - \mathcal{U}^\infty = \mathbf{R}_{\mathcal{FU}}^{-1} \cdot (\mathcal{F}^P + \mathbf{R}_{\mathcal{FE}} \cdot \mathcal{E}^\infty) + \sqrt{\frac{2k_B T}{\Delta t}} \mathbf{R}_{\mathcal{FU}}^{-1/2} \cdot \psi + k_B T \nabla \cdot \mathbf{R}_{\mathcal{FU}}^{-1}. \quad (2)$$

Here,  $\mathcal{F}^{\mathcal{P}}$  is a  $6N$ -dimensional vector whose first  $3N$  components are the forces applied to the particles, and the last  $3N$  components are the applied torques. The vector  $\psi$  is a  $6N$ -dimensional vector containing random variables normally distributed with zero mean and unit variance, and  $\Delta t$  is the simulation timestep. The term  $k_{\text{B}}T$  represents the thermal energy;  $k_{\text{B}}$  is the Boltzmann constant and  $T$  the temperature. The tensor  $\mathbf{R}_{\mathcal{F}\mathcal{U}}$  is the upper-left subblock of  $\mathcal{R}$ , which maps velocities to forces, while  $\mathbf{R}_{\mathcal{F}\mathcal{E}}$  is the upper-right subblock that couples strain rates to forces in  $\mathcal{R}$ , respectively. In Eq. (2), the first two terms represent deterministic contributions to the particle dynamics coming from the applied forces and torques, as well as the background flow, while the last two terms account for thermal fluctuations.

In SD [5], hydrodynamic interactions are incorporated through a combination of the mobility and resistance frameworks. The construction of the grand resistance matrix  $\mathcal{R}$  begins with the far-field mobility matrix  $\mathcal{M}^{\text{ff}}$ . Short-range interactions are resolved in a pairwise manner by adding near-field resistance contributions  $\mathcal{R}^{\text{nf}}$  [6, 8]. To avoid double-counting, the far-field pairwise contributions  $\mathcal{R}_{2\text{B}}^{\text{ff}}$  are subtracted

$$\mathcal{R} = (\mathcal{M}^{\text{ff}})^{-1} + \mathcal{R}^{\text{nf}} - \mathcal{R}_{2\text{B}}^{\text{ff}}. \quad (3)$$

The FSD method [1] extends the SD framework by formulating the particle dynamics as a linear system of  $17N$  equations using a saddle-point matrix that integrates near-field (short-ranged, pairwise additive) and far-field (long-ranged, many-body) interactions. This formulation enables iterative solvers with specialized preconditioners to efficiently handle the resulting linear systems. The approach avoids the explicit inversion of ill-conditioned hydrodynamic operators and drastically reducing computational costs. Brownian forces are computed by combining the positively-split Ewald method [36] and an iterative Krylov subspace method [37], both seamlessly integrated into the saddle-point formulation. We refer the interested reader to Ref. [1] for the full details.

Our implementation of the method, JFSD, is structured as a modular Python package optimized for scalability and high-performance hydrodynamic simulations. In particular, the configuration management module leverages a user-friendly TOML file that specifies simulation parameters — such as the number of steps, particle count, time step, initialization settings, and output options — providing a transparent and easily modifiable interface for setting up simulations. The modular design allows for easy extension of pairwise interaction models, enabling users to implement custom hydrodynamic kernels or introduce *ad hoc* force laws with minimal modifications. The framework also provides built-in support for steady and oscillatory shear flows, with a straightforward pathway toward incorporating more complex shear protocols by modifying existing velocity gradients. Boundary conditions can be adjusted efficiently by altering the scalar mobility functions, as demonstrated in the already implemented transition between open and periodic boundary conditions. Lastly, our neighbor list management was derived from the one used in JAX-MD [38], a widely used molecular dynamics framework. Our implementation ensures FSD-compatible handling of particle interactions in both dense and dilute systems. Full details of our implementation can be found in the documentation accompanying the JFSD release [39].

A key advantage of JFSD is its integration with Google JAX [35], which enables Just-In-Time (JIT) compilation and GPU acceleration for significant performance improvements. JAX compiles entire computation graphs into highly optimized routines that eliminate Python overhead by fusing multiple operations into a single efficient kernel. Moreover, its vectorized operations and automatic backend dispatch ensure that computations run optimally on both GPU and CPU, making the framework highly scalable across diverse hardware configurations.

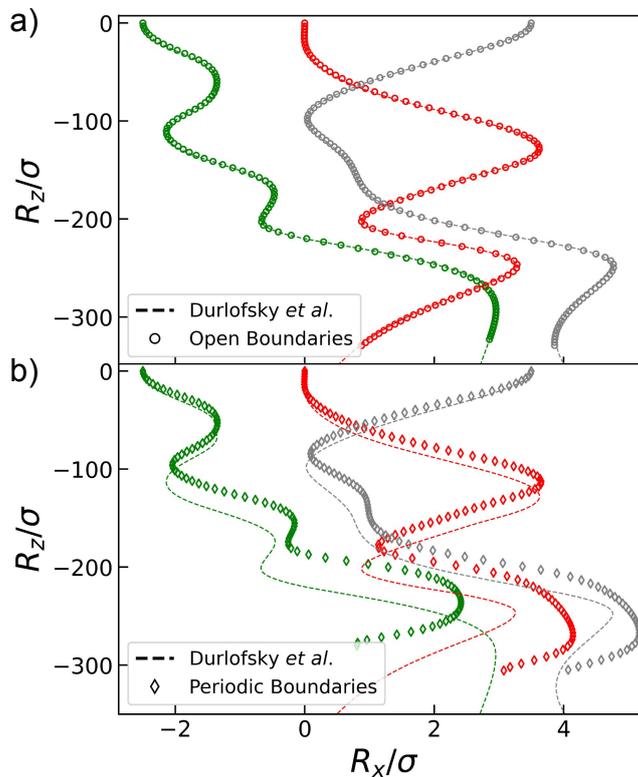


Figure 1: Trajectories of three equal spheres sedimenting vertically, starting from initial positions at  $\mathbf{r}_1 = (-2.5\sigma, 0, 0)$  (green),  $\mathbf{r}_2 = (0, 0, 0)$  (red), and  $\mathbf{r}_3 = (3.5\sigma, 0, 0)$  (gray). Results obtained using the JFSD method under (a) open boundary conditions (circles) and (b) periodic boundary conditions (rhombi), in a box of size  $L = 25\sigma$ . Both panels include reference data from Ref. [40] (dashed curves), where the original data points have been smoothed using spline interpolation.

### 3 Method validations

Following a general discussion of (F)SD, we now show results obtained using JFSD. These are contrasted against two-body analytical expressions and/or independently obtained numerical data. We also showcase the efficiency of the JFSD software by benchmarking the code for the case of a hard-sphere system across a range of volume fractions  $\phi$ .

#### 3.1 Sedimenting particles

The deterministic mapping from forces to velocities was validated by simulating the sedimentation of three particles aligned along the  $x$ -axis. Our results are compared to the original Stokesian Dynamics (SD) trajectories by Durlofsky *et al.* [40] for an unbounded fluid solvent in Fig. 1. This shows the sedimentation trajectories of the three particles. Here, JFSD demonstrates excellent agreement with the SD reference result. The trajectories align closely over the full time interval considered. With JFSD we also considered periodic boundaries, for which discrepancies become pronounced as time progresses. These deviations are expected and arise due to interaction with periodic images. Despite these differences, JFSD faithfully reproduces the initial dynamics. The combined result is sufficient to conclude that JFSD is consistent with the original SD framework.

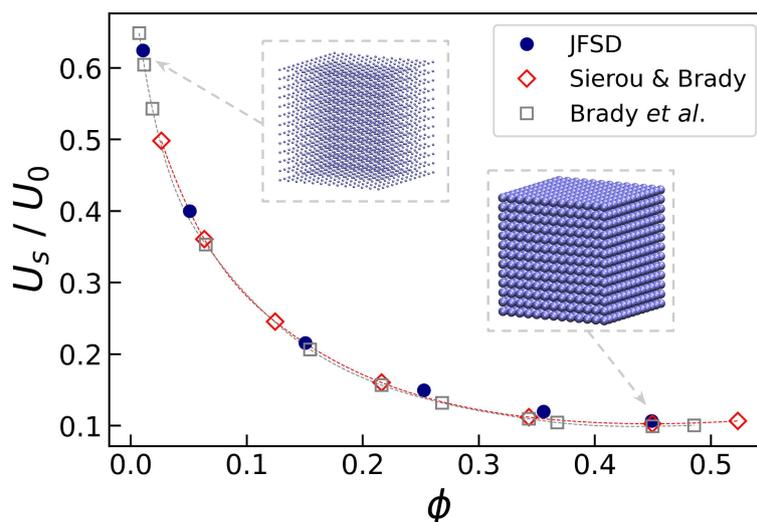


Figure 2: Comparison of sedimenting velocity  $U_s$  (blue dots) for a simple cubic array of  $N = 2,197$  particles (insets show examples) as a function of the volume fraction  $\phi$ . The sedimentation velocity is normalized by the infinite dilution limit  $U_0$ . Results from the JFSD method under periodic boundary conditions are contrasted with those obtained using Sierou’s accelerated Stokesian dynamics [31] (red rhombi) and Brady *et al.*’s original Stokesian dynamics [41] (gray squares). The curves connecting the two reference data sets, obtained through spline interpolation, serve to guide the eye.

Next, we computed the sedimentation velocity  $U_s$  of a periodic cubic array of  $N = 2,197$  hard spheres across a range of  $\phi$ , see Fig. 2. These results were compared against the benchmark data from Sierou and Brady [31], which revealed that JFSD accurately captures the non-linear decay of  $U_s$  with increasing  $\phi$ . Our implementation gives consistent results even in dense regimes, which implies that it can reliably propagate suspension dynamics across a wide range of concentrations.

### 3.2 Hydrodynamic effects under shear flow

Now that we have examined the quality of the force-torque and (angular) velocity coupling, we turn to shear. We investigated the behavior of particle pairs in simple shear flow to benchmark the accuracy of JFSD under these conditions, see Fig. 3, which shows three example trajectories. The results align closely with analytical predictions [4, 42], demonstrating JFSD’s performance.

In addition to pair interactions, we computed the shear viscosity  $\eta_{\text{eff}}$  of a periodic simple cubic array of (again)  $N = 2,197$  hard spheres as a function of  $\phi$ . Figure 4 compares our result with data from the original FSD implementation by Fiore and Swan [1]. The accurate agreement shows that both algorithms produce the same results and are capable of computing rheological properties even at high  $\phi$ .

Our combined shear results allow us to conclude that this part of the algorithm is correctly implemented. Of course, additional testing was performed on the behavior of the  $\mathcal{R}$  matrix, but we chose only to report on physical observables here.

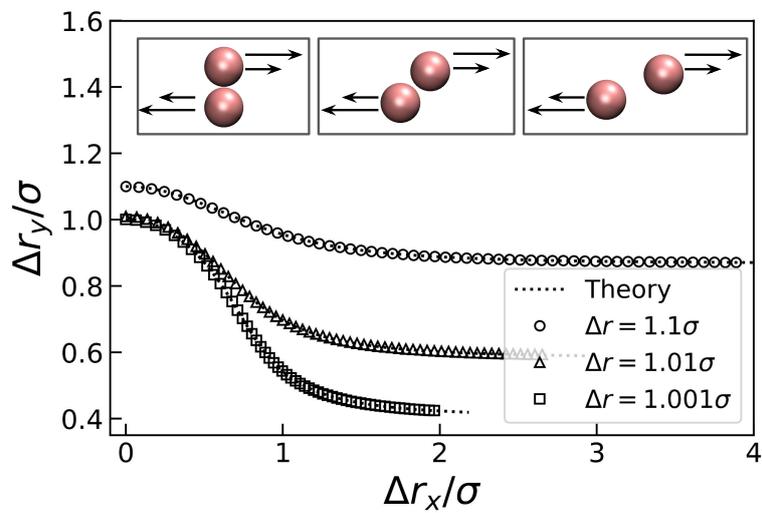


Figure 3: Trajectories of a particle pair in simple shear flow under open boundary conditions. Results from JFSD (symbols) are compared against analytical predictions [4, 42] (dotted lines) for three initial center-to-center separations  $\Delta r$ . These trajectories are represented in the distance along the  $x$ - and  $y$ -axes, respectively, normalized by the particle diameter  $\sigma$ . The insets show three snapshots of a trajectory for spheres spaced  $\Delta r = 1.01\sigma$  apart orthogonal to the shear plane.

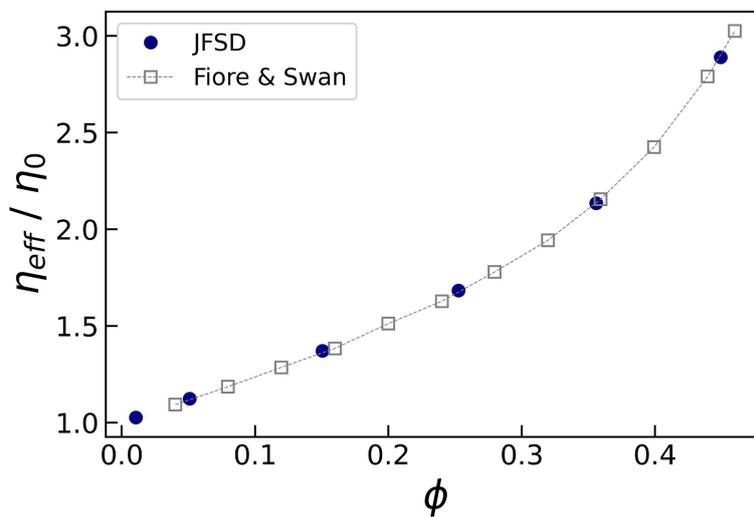


Figure 4: Shear viscosity  $\eta_{\text{eff}}$  of a simple cubic array of  $N = 2,197$  particles as a function of volume fraction  $\phi$ . The results are normalized by the bare viscosity of the suspending medium  $\eta_0$ . Python-JFSD results (blue circles) are compared with data from Fiore and Swan [1] (open squares). The dotted connecting curve serves to guide the eye.

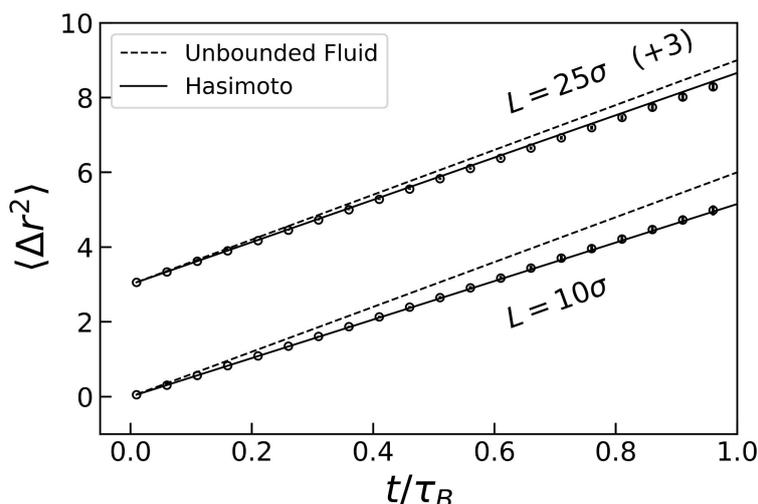


Figure 5: Mean squared displacement  $\langle r^2 \rangle$  of a single particle as a function of time  $t$  divided by the Brownian time  $t_B$ . The JFSD data is given by open circles with error bars representing the standard error of the mean for two box sizes  $L$  as labelled. The data for the  $L = 25\sigma$  ( $\sigma$  is the sphere diameter) box is shifted by the constant offset  $+3$  in the vertical direction to help aid the presentation. The dotted line shows the analytic  $\langle r^2 \rangle = 6Dt$  result for an unbounded fluid, while the solid line applies Hasimoto’s correction to obtain the appropriate expression for a periodic domain [43].

### 3.3 Thermal motion

Having verified the deterministic part of the algorithm, we turned our attention to thermal fluctuations. The mean squared displacement (MSD)  $\langle r^2 \rangle$  of a single particle in 3D was computed under periodic boundary conditions and benchmarked against Hasimoto’s corrected expression for periodic boundaries [43]. To generate the MSD curves in Fig. 5, we simulated for 100 Brownian times and repeated the process for 10 independent samples for each box size  $L$ , averaging the results across these runs. The figure reveals that the simulated MSD aligns closely with theoretical predictions, confirming the accurate implementation of periodic images and the validity of thermal diffusion in the JFSD framework.

However, MSD analysis alone is not sufficient to validate the thermal part of the algorithm. This is because in SD, there is a ‘Brownian drift’ term that is required to ensure stationarity under the Gibbs-Boltzmann distribution and generate particle configurations with the correct statistics at equilibrium [1, 44]. To demonstrate that we have correctly implemented this aspect, we draw inspiration from a test originally considered by Fiore and Swan [1, 44]. That is, we simulated a pair of particles, in an unbounded three-dimensional (3D) fluid, interacting via a linear potential

$$V(r) = k|r - r_0|. \tag{4}$$

Here,  $k$  is the potential strength,  $r$  is the radial distance between the sphere centers, and  $r_0$  is the equilibrium distance. Figure 6 compares the average distribution of radial distances  $p(r)$  derived from 5 independent simulations — with duration of 5 Brownian times each — against the theoretical Boltzmann probability density function [1]. The good agreement demonstrates Python-JFSD’s ability to reproduce expected equilibrium distributions. Small deviations in the tail of the distribution arise primarily from the accumulation of numerical errors in long trajectories, due to the use of a first-order Euler integration scheme [45].

We also considered large particle number systems ( $N = 2,500$ ) with periodic boundary conditions. For these, we computed the short-time self-diffusion coefficient  $D_s$  as a function of

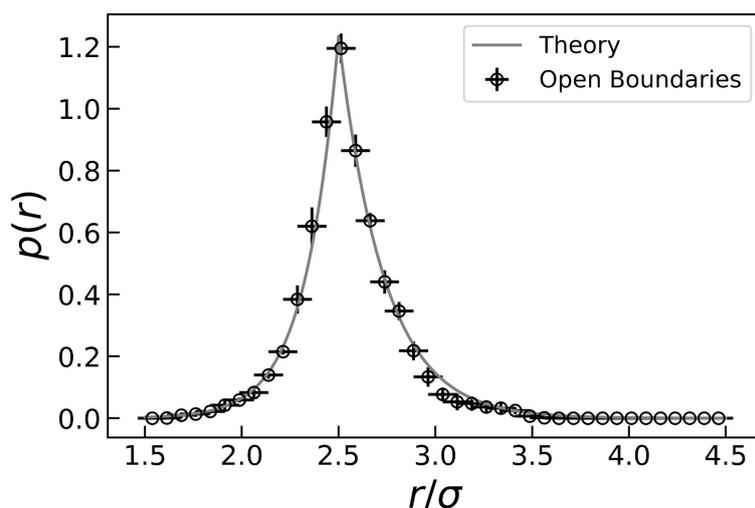


Figure 6: Radial distribution  $p(r) \propto r^2 \exp(-V(r)/k_B T)$  of a pair of particles interacting via a linear potential (see Eq. (4)) under open boundary conditions, with  $r$  the center-to-center distance between the two spheres. The JFSD data is shown using open circles, with the vertical error bars showing the standard error of the mean (computed from 5 independent samples). The horizontal error bars follow from the fact that the data was obtained by a binning procedure.

$\phi$  for hard-sphere suspensions, see Fig. 7. Our results are compared therein with two independently derived datasets. The first is by Ladd [46] and was obtained from a multipole-moment expansion, combined with pairwise lubrication interactions, that contains a higher number of moments than conventional SD. The second follows from Sierou’s and Brady’s work [31] and represents the Accelerated Stokesian Dynamics method. Both datasets were obtained for open systems, hence we have corrected for the periodicity effects using Hasimoto’s scaling factor [31, 43, 46]. The close agreement between our results and the (scaled) literature values confirms JFSD’s ability to resolve thermal fluctuations in systems with many particles in periodic boundary conditions.

### 3.4 Algorithmic performance and scaling

We assessed the computational performance of Python-JFSD using an NVIDIA GeForce RTX 4060 Ti with 8GB of VRAM. As a standard performance metric, we define the particle time steps per second (PTPS) as  $\text{PTPS} = N/T_{\text{step}}$ , where  $T_{\text{step}}$  is the clock time in seconds required to perform a single simulation step as a function of the number of particles  $N$ . The setup consists of  $N$  hard spheres arranged in a cubic array, interacting only via excluded-volume interactions to prevent overlap. The particles undergo thermal motion, and a simple shear flow is applied in the  $xy$ -plane (with vorticity in the  $z$ -direction, flow in the  $x$ -direction, and flow gradient in the  $y$ -direction). Periodic boundary conditions are imposed. To ensure robust benchmarking, the system is simulated for 50 steps, and the average  $T_{\text{step}}$  is computed after excluding the first step, which involves JIT compilation and can take up to 1-2 minutes.

Since computational performance depends on the volume fraction  $\phi$ , we considered four representative values,  $\phi \in \{0.05, 0.15, 0.4, 0.5\}$ , to assess scaling behavior. This dependence arises from the iterative solver used in the FSD algorithm [1], where the number of iterations required for convergence increases due to lubrication interactions, which become more significant at higher  $\phi$  as particles are closer together. The results, shown in Fig. 8, present PTPS as a function of  $N$ .

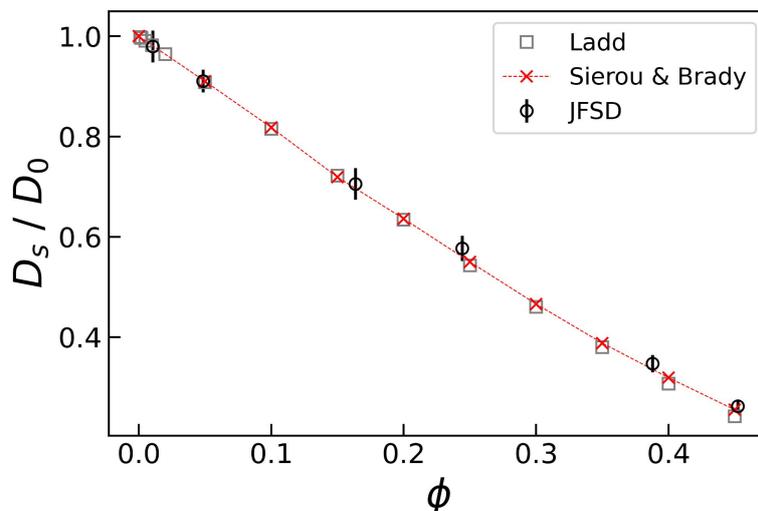


Figure 7: Short-time self-diffusion coefficient  $D_s$  for hard-sphere suspensions under periodic boundary conditions as a function of volume fraction  $\phi$ . The diffusion coefficient is normalized by the bulk value  $D_0$ . Results from Python-JFSD (open circles) have error bars that show the standard error of the mean. These are compared to the work by Ladd [46] (open circle) and Sierou [31] (red crosses). Both datasets are corrected for finite-size effects *via* the Hasimoto factor [43] and the dashed curves help guide the eye.

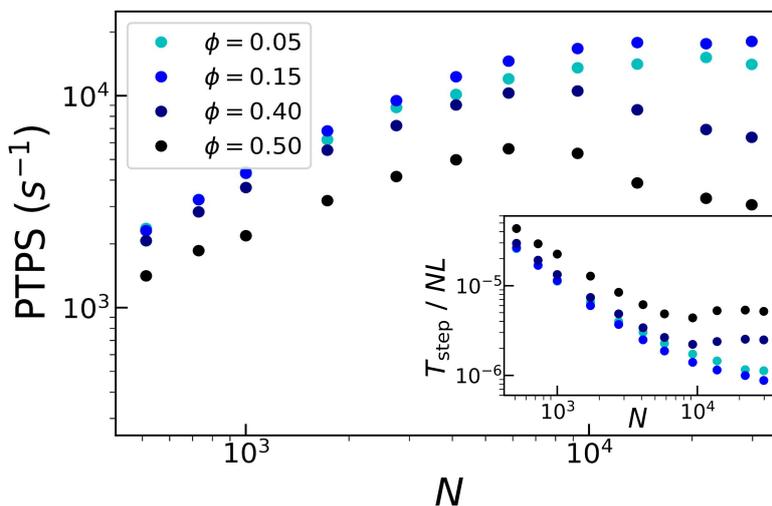


Figure 8: Particle time steps per second (PTPS) as a function of the number of particles  $N$  in a periodic simulation volume. Four volume fractions were considered:  $\phi \in \{0.05, 0.15, 0.4, 0.5\}$ , as labeled. At low  $\phi$ , PTPS plateaus for large  $N$ , while for high  $\phi$ , PTPS decreases with  $N$  due to system-size dependence. The inset shows the normalized step time per particle per unit box size,  $T_{\text{step}}/(NL)$ , which remains constant for large  $\phi$ , confirming that the algorithm maintains linear scaling despite the decrease in PTPS in the main plot.

The PTPS trends indicate that at low  $\phi$ , performance reaches a steady plateau, as expected for sufficiently large  $N$ . Conversely, at higher volume fractions, PTPS continues to decrease with increasing  $N$ . This behavior suggests that system size influences the cost of a single solver iteration by increasing the number of Fourier grid points required for far-field hydrodynamics. At high volume fractions, where lubrication slows convergence, this added cost becomes more pronounced. The inset in Fig. 8 clarifies this effect: for large  $\phi$ , the normalized time step per particle per unit box size remains unchanged, verifying that the scaling of the FSD algorithm is preserved. This shows that although PTPS decreases at high  $\phi$ , the core computational efficiency of the method remains intact, consistent with theoretical predictions [1].

## 4 Discussion

SD is a well-established framework for simulating particle suspensions [1, 5, 31–33, 40], capable of accurately resolving both many-body hydrodynamic interactions and near-field lubrication forces. It is particularly effective for systems requiring precise particle-level hydrodynamic modeling without complicated boundary conditions. SD relies on direct calculations of forces and torques, which allows for computational efficiency without the need for fine spatial discretization. Recent applications include sedimentation in colloidal suspensions [13], colloidal gelation [18], and the dynamics of active matter [47] in confining geometries [48, 49]. These examples showcase SD’s versatility across a broad range of (colloidal) particulate systems.

Implementing SD methods, however, presents significant challenges due to the algorithm’s reliance on complex mathematical algorithms [50] and the precision required to accurately model lubrication forces and many-body interactions. Ensuring correctness and robustness needs meticulous validation. To address this, our JFSD implementation incorporates continuous integration workflows, automatically checking each pull request – originating from Git-based version control – against an extensive suite of benchmarks [39]. This rigorous testing, combined with a modular design, allows users to choose between periodic or open boundary conditions and various hydrodynamic models, including FSD, which has been the focus of this paper. In addition to these features, our software suite is also capable of modeling hydrodynamic interactions at the Rotne-Prager-Yamakawa level [36, 51, 52], which is suitable for capturing pair-wise, far-field coupling. We have further implemented a simple Brownian Dynamics code, which focuses on the hydrodynamic self-interaction only. However, this last algorithm has not been optimized to the same degree as is available in dedicated software packages, as our primary focus has been on realizing a new version of FSD.

JFSD was created to provide greater accessibility to SD by leveraging Python’s simplicity and interoperability, while maintaining the competitive computational performance of dedicated implementations. Thanks to Google JAX’s JIT compilation and reliance on GPU acceleration, our software achieves speeds comparable to previous CUDA<sup>®</sup> implementations [1]. Presently, JFSD is capable of simulating approximately 30,000-50,000 particles within the memory constraints of mid-range GPUs. While this imposes a practical upper limit, ongoing developments in Google JAX, particularly the support for multi-core CPU and distributed computing, could further expand the range of the software in the future.

## 5 Conclusion and outlook

In summary, we have presented a Python-based implementation of fast Stokesian Dynamics that relies on the Google JAX, which we refer to as JFSD. Our work addresses the use issues posed by the currently outdated CUDA<sup>®</sup> code base [1] originally developed by Fiore and Swan. By utilizing Just-In-Time compilation [35], JFSD achieves identical scaling, while maintaining performance and offering a user-friendly and modular interface. Our software is capable of

simulations involving tens of thousands of particles on a modern desktop GPU.

We built upon the extensive literature for SD and hydrodynamic interactions in suspensions to validate the code. Key benchmarks performed in this paper include: sedimentation, shear rheology, and thermal motion. We found strong agreement with the literature, demonstrating the quality of our implementation. The inclusion of both open and periodic boundary conditions, along with alternative hydrodynamic models, ensures flexibility and broad applicability for diverse physical systems. JFSD is thus ready to enable new research into the effect of hydrodynamic interactions in particle suspensions.

## Acknowledgments

We developed JFSD with the aim of giving back to the fluid-dynamics community, whose efforts and developments [1, 34, 36, 40, 44, 52, 53] helped support our own research [18, 54]. We hope that our implementation of FSD will allow researchers to continue to make use of the excellent algorithms developed by Fiore and Swan [1] for many years to come. Thank you Jim for the kindness that you showed us. We are also grateful to Andrew M. Fiore, Zhouyang Ge, and Gwynn Elfring; Luca Leone and Henri Menke; and Athanasios Machas and Dimos Aslanis for useful discussions, which helped us significantly improve our understanding of FSD, Google JAX, and the user/developer perspective, respectively.

**Data availability statement** An open data package containing the means to reproduce the results of the simulations is available at: doi:[10.24416/UU01-QO36GQ](https://doi.org/10.24416/UU01-QO36GQ). The specific version of JFSD used to generate the results in this paper is tagged as v0.2.0, which will aid with reproducibility.

**Author contributions** Conceptualization, J.d.G. & K.W.T.; Methodology, K.W.T. & R.S.; Numerical calculations, K.W.T., Validation, K.W.T. & R.S.; Investigation, K.W.T.; Writing — Original Draft, K.W.T.; Writing — Review & Editing, J.d.G.; Funding Acquisition, J.d.G.; Resources, J.d.G.; Supervision, J.d.G.

**Funding information** The authors acknowledge the Dutch Research Council NWO for funding through OCENW.KLEIN.354, as well as the International Fine Particle Research Institute for funding through collaboration grant CRR-118-01.

## References

- [1] A. M. Fiore and J. W. Swan, *Fast Stokesian dynamics*, *J. Fluid Mech.* **878**, 544 (2019), doi:[10.1017/jfm.2019.640](https://doi.org/10.1017/jfm.2019.640).
- [2] G. K. Batchelor, *The effect of Brownian motion on the bulk stress in a suspension of spherical particles*, *J. Fluid Mech.* **83**, 97 (1977), doi:[10.1017/S0022112077001062](https://doi.org/10.1017/S0022112077001062).
- [3] J. Happel and H. Brenner, *Low reynolds number hydrodynamics: With special applications to particulate media*, Springer, Dordrecht, Netherlands, ISBN 9789400983526 (1983), doi:[10.1007/978-94-009-8352-6](https://doi.org/10.1007/978-94-009-8352-6).
- [4] S. Kim and S. J. Karrila, *Microhydrodynamics: Principles and selected applications*, Butterworth-Heinemann, Oxford, UK, ISBN 9780750691734 (1991), doi:[10.1016/C2013-0-04644-0](https://doi.org/10.1016/C2013-0-04644-0).

- [5] J. F. Brady and G. Bossis, *Stokesian dynamics*, *Annu. Rev. Fluid Mech.* **20**, 111 (1988), doi:[10.1146/annurev.fl.20.010188.000551](https://doi.org/10.1146/annurev.fl.20.010188.000551).
- [6] D. J. Jeffrey and Y. Onishi, *Calculation of the resistance and mobility functions for two unequal rigid spheres in low-Reynolds-number flow*, *J. Fluid Mech.* **139**, 261 (1984), doi:[10.1017/S0022112084000355](https://doi.org/10.1017/S0022112084000355).
- [7] S. Kim and R. T. Mifflin, *The resistance and mobility functions of two equal spheres in low-Reynolds-number flow*, *Phys. Fluids* **28**, 2033 (1985), doi:[10.1063/1.865384](https://doi.org/10.1063/1.865384).
- [8] D. J. Jeffrey, *The calculation of the low Reynolds number resistance functions for two unequal spheres*, *Phys. Fluids A: Fluid Dyn.* **4**, 16 (1992), doi:[10.1063/1.858494](https://doi.org/10.1063/1.858494).
- [9] E. C. Eckstein, D. G. Bailey and A. H. Shapiro, *Self-diffusion of particles in shear flow of a suspension*, *J. Fluid Mech.* **79**, 191 (1977), doi:[10.1017/S0022112077000111](https://doi.org/10.1017/S0022112077000111).
- [10] A. Furukawa and H. Tanaka, *Key role of hydrodynamic interactions in colloidal gelation*, *Phys. Rev. Lett.* **104**, 245702 (2010), doi:[10.1103/PhysRevLett.104.245702](https://doi.org/10.1103/PhysRevLett.104.245702).
- [11] J. K. Whitmer and E. Luijten, *Influence of hydrodynamics on cluster formation in colloid-polymer mixtures*, *J. Phys. Chem. B* **115**, 7294 (2011), doi:[10.1021/jp111388m](https://doi.org/10.1021/jp111388m).
- [12] C. P. Royall, J. Eggers, A. Furukawa and H. Tanaka, *Probing colloidal gels at multiple length scales: The role of hydrodynamics*, *Phys. Rev. Lett.* **114**, 258302 (2015), doi:[10.1103/PhysRevLett.114.258302](https://doi.org/10.1103/PhysRevLett.114.258302).
- [13] Z. Varga, J. L. Hofmann and J. W. Swan, *Modelling a hydrodynamic instability in freely settling colloidal gels*, *J. Fluid Mech.* **856**, 1014 (2018), doi:[10.1017/jfm.2018.725](https://doi.org/10.1017/jfm.2018.725).
- [14] K. A. Whitaker, Z. Varga, L. C. Hsiao, M. J. Solomon, J. W. Swan and E. M. Furst, *Colloidal gel elasticity arises from the packing of locally glassy clusters*, *Nat. Commun.* **10**, 2237 (2019), doi:[10.1038/s41467-019-10039-w](https://doi.org/10.1038/s41467-019-10039-w).
- [15] L. Turetta and M. Lattuada, *The role of hydrodynamic interactions on the aggregation kinetics of sedimenting colloidal particles*, *Soft Matter* **18**, 1715 (2022), doi:[10.1039/D1SM01637G](https://doi.org/10.1039/D1SM01637G).
- [16] K. W. Torre and J. de Graaf, *Structuring colloidal gels via micro-bubble oscillations*, *Soft Matter* **19**, 2771 (2023), doi:[10.1039/D2SM01450E](https://doi.org/10.1039/D2SM01450E).
- [17] J. de Graaf, K. W. Torre, W. C. K. Poon and M. Hermes, *Hydrodynamic stability criterion for colloidal gelation under gravity*, *Phys. Rev. E* **107**, 034608 (2023), doi:[10.1103/PhysRevE.107.034608](https://doi.org/10.1103/PhysRevE.107.034608).
- [18] K. W. Torre and J. de Graaf, *Hydrodynamic lubrication in colloidal gels*, *Soft Matter* **19**, 7388 (2023), doi:[10.1039/D3SM00784G](https://doi.org/10.1039/D3SM00784G).
- [19] K. W. Torre and J. de Graaf, *Delayed gravitational collapse of attractive colloidal suspensions*, *J. Fluid Mech.* **1000**, A73 (2024), doi:[10.1017/jfm.2024.1076](https://doi.org/10.1017/jfm.2024.1076).
- [20] D. S. Bolintineanu, G. S. Grest, J. B. Lechman, F. Pierce, S. J. Plimpton and P. R. Schunk, *Particle dynamics modeling methods for colloid suspensions*, *Comput. Part. Mech.* **1**, 321 (2014), doi:[10.1007/s40571-014-0007-6](https://doi.org/10.1007/s40571-014-0007-6).
- [21] S. Chen and G. D. Doolen, *Lattice Boltzmann method for fluid flows*, *Annu. Rev. Fluid Mech.* **30**, 329 (1998), doi:[10.1146/annurev.fluid.30.1.329](https://doi.org/10.1146/annurev.fluid.30.1.329).

- [22] A. J. C. Ladd and R. Verberg, *Lattice-Boltzmann simulations of particle-fluid suspensions*, J. Stat. Phys. **104**, 1191 (2001), doi:[10.1023/A:1010414013942](https://doi.org/10.1023/A:1010414013942).
- [23] B. Dünweg and A. J. C. Ladd, *Lattice Boltzmann simulations of soft matter systems*, in *Advanced computer simulation approaches for soft matter sciences III*, Springer, Berlin, Heidelberg, Germany, ISBN 9783540877059 (2009), doi:[10.1007/978-3-540-87706-6\\_2](https://doi.org/10.1007/978-3-540-87706-6_2).
- [24] C. K. Aidun and J. R. Clausen, *Lattice-Boltzmann method for complex flows*, Annu. Rev. Fluid Mech. **42**, 439 (2010), doi:[10.1146/annurev-fluid-121108-145519](https://doi.org/10.1146/annurev-fluid-121108-145519).
- [25] D. Roehm and A. Arnold, *Lattice Boltzmann simulations on GPUs with ESPResSo*, Eur. Phys. J. Spec. Top. **210**, 89 (2012), doi:[10.1140/epjst/e2012-01639-6](https://doi.org/10.1140/epjst/e2012-01639-6).
- [26] G. Gompper, T. Ihle, D. M. Kroll and R. G. Winkler, *Multi-particle collision dynamics: A particle-based mesoscale simulation approach to the hydrodynamics of complex fluids*, in *Advanced computer simulation approaches for soft matter sciences III*, Springer, Berlin, Heidelberg, Germany, ISBN 9783540877059 (2009), doi:[10.1007/978-3-540-87706-6\\_1](https://doi.org/10.1007/978-3-540-87706-6_1).
- [27] M. P. Howard, A. Z. Panagiotopoulos and A. Nikoubashman, *Efficient mesoscale hydrodynamics: Multiparticle collision dynamics with massively parallel GPU acceleration*, Comput. Phys. Commun. **230**, 10 (2018), doi:[10.1016/j.cpc.2018.04.009](https://doi.org/10.1016/j.cpc.2018.04.009).
- [28] A. Malevanets and R. Kapral, *Mesosopic model for solvent dynamics*, J. Chem. Phys. **110**, 8605 (1999), doi:[10.1063/1.478857](https://doi.org/10.1063/1.478857).
- [29] H. Tanaka and T. Araki, *Simulation method of colloidal suspensions with hydrodynamic interactions: Fluid particle dynamics*, Phys. Rev. Lett. **85**, 1338 (2000), doi:[10.1103/PhysRevLett.85.1338](https://doi.org/10.1103/PhysRevLett.85.1338).
- [30] N.-Q. Nguyen and A. J. C. Ladd, *Lubrication corrections for lattice-Boltzmann simulations of particle suspensions*, Phys. Rev. E **66**, 046708 (2002), doi:[10.1103/PhysRevE.66.046708](https://doi.org/10.1103/PhysRevE.66.046708).
- [31] A. Sierou and J. F. Brady, *Accelerated Stokesian dynamics simulations*, J. Fluid Mech. **448**, 115 (2001), doi:[10.1017/S0022112001005912](https://doi.org/10.1017/S0022112001005912).
- [32] A. J. Banchio and J. F. Brady, *Accelerated Stokesian dynamics: Brownian motion*, J. Chem. Phys. **118**, 10323 (2003), doi:[10.1063/1.1571819](https://doi.org/10.1063/1.1571819).
- [33] M. Wang and J. F. Brady, *Spectral Ewald acceleration of Stokesian dynamics for polydisperse suspensions*, J. Comput. Phys. **306**, 443 (2016), doi:[10.1016/j.jcp.2015.11.042](https://doi.org/10.1016/j.jcp.2015.11.042).
- [34] J. A. Anderson, J. Glaser and S. C. Glotzer, *HOOMD-blue: A Python package for high-performance molecular dynamics and hard particle Monte Carlo simulations*, Comput. Mater. Sci. **173**, 109363 (2020), doi:[10.1016/j.commatsci.2019.109363](https://doi.org/10.1016/j.commatsci.2019.109363).
- [35] J. Bradbury, R. Frostig, P. Hawkins, M. J. Johnson, C. Leary, D. Maclaurin, G. Necula, A. Paszke, J. VanderPlas, S. Wanderman-Milne and Q. Zhang, *JAX: Composable transformations of Python+NumPy programs* (2018), <http://github.com/google/jax>.
- [36] A. M. Fiore, F. Balboa Usabiaga, A. Donev and J. W. Swan, *Rapid sampling of stochastic displacements in Brownian dynamics simulations*, J. Chem. Phys. **146**, 124116 (2017), doi:[10.1063/1.4978242](https://doi.org/10.1063/1.4978242).
- [37] E. Chow and Y. Saad, *Preconditioned Krylov subspace methods for sampling multivariate Gaussian distributions*, SIAM J. Sci. Comput. **36**, A588 (2014), doi:[10.1137/130920587](https://doi.org/10.1137/130920587).

- [38] S. S. Schoenholz and E. D. Cubuk, *JAX, M.D. A framework for differentiable physics*, J. Stat. Mech.: Theory Exp. 124016 (2021), doi:[10.1088/1742-5468/ac3ae9](https://doi.org/10.1088/1742-5468/ac3ae9).
- [39] K. W. Torre, J. De Graaf and R. Schram, *Python-JAX-fast-Stokesian-dynamics*, <https://github.com/torrewk/Python-Jax-Fast-Stokesian-Dynamics>.
- [40] L. Durlofsky, J. F. Brady and G. Bossis, *Dynamic simulation of hydrodynamically interacting particles*, J. Fluid Mech. **180**, 21 (1987), doi:[10.1017/S002211208700171X](https://doi.org/10.1017/S002211208700171X).
- [41] J. F. Brady, R. J. Phillips, J. C. Lester and G. Bossis, *Dynamic simulation of hydrodynamically interacting suspensions*, J. Fluid Mech. **195**, 257 (1988), doi:[10.1017/S0022112088002411](https://doi.org/10.1017/S0022112088002411).
- [42] F. R. Da Cunha and E. J. Hinch, *Shear-induced dispersion in a dilute suspension of rough spheres*, J. Fluid Mech. **309**, 211 (1996), doi:[10.1017/S0022112096001619](https://doi.org/10.1017/S0022112096001619).
- [43] 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**, 317 (1959), doi:[10.1017/S0022112059000222](https://doi.org/10.1017/S0022112059000222).
- [44] A. M. Fiore and J. W. Swan, *Rapid sampling of stochastic displacements in Brownian dynamics simulations with stresslet constraints*, J. Chem. Phys. **148**, 044114 (2018), doi:[10.1063/1.5005887](https://doi.org/10.1063/1.5005887).
- [45] D. L. Ermak and J. A. McCammon, *Brownian dynamics with hydrodynamic interactions*, J. Chem. Phys. **69**, 1352 (1978), doi:[10.1063/1.436761](https://doi.org/10.1063/1.436761).
- [46] A. J. C. Ladd, *Hydrodynamic transport coefficients of random dispersions of hard spheres*, J. Chem. Phys. **93**, 3484 (1990), doi:[10.1063/1.458830](https://doi.org/10.1063/1.458830).
- [47] Z. Ge and G. J. Elfring, *Hydrodynamic diffusion in apolar active suspensions of squirmers*, J. Fluid Mech. **1003**, A17 (2025), doi:[10.1017/jfm.2024.1071](https://doi.org/10.1017/jfm.2024.1071).
- [48] G. J. Elfring and J. F. Brady, *Active Stokesian dynamics*, J. Fluid Mech. **952**, A19 (2022), doi:[10.1017/jfm.2022.909](https://doi.org/10.1017/jfm.2022.909).
- [49] V. A. Shaik, Z. Peng, J. F. Brady and G. J. Elfring, *Confined active matter in external fields*, Soft Matter **19**, 1384 (2023), doi:[10.1039/D2SM01135B](https://doi.org/10.1039/D2SM01135B).
- [50] A. K. Townsend, *Generating, from scratch, the near-field asymptotic forms of scalar resistance functions for two unequal rigid spheres in low Reynolds number flow*, Phys. Fluids **35**, 127126 (2023), doi:[10.1063/5.0175697](https://doi.org/10.1063/5.0175697).
- [51] J. Rotne and S. Prager, *Variational treatment of hydrodynamic interaction in polymers*, J. Chem. Phys. **50**, 4831 (1969), doi:[10.1063/1.1670977](https://doi.org/10.1063/1.1670977).
- [52] R. P. Peláez, *Complex fluids in the GPU era: Algorithms and simulations*, PhD thesis, Universidad Autónoma de Madrid, Madrid, Spain (2021).
- [53] A. K. Townsend, *Stokesian dynamics in Python*, J. Open Source Softw. **9**, 6011 (2024), doi:[10.21105/joss.06011](https://doi.org/10.21105/joss.06011).
- [54] J. de Graaf, W. C. K. Poon, M. J. Haughey and M. Hermes, *Hydrodynamics strongly affect the dynamics of colloidal gelation but not gel structure*, Soft Matter **15**, 10 (2019), doi:[10.1039/C8SM01611A](https://doi.org/10.1039/C8SM01611A).