

A smooth particle-mesh Ewald algorithm for Stokes suspension simulations: The sedimentation of fibers

David Saintillan and Eric Darve^{a)}

Department of Mechanical Engineering, Stanford University, Stanford, California 94305

Eric S. G. Shaqfeh

Department of Chemical Engineering and Department of Mechanical Engineering, Stanford University, Stanford, California 94305

(Received 17 May 2004; accepted 21 December 2004; published online 1 March 2005)

Large-scale simulations of non-Brownian rigid fibers sedimenting under gravity at zero Reynolds number have been performed using a fast algorithm. The mathematical formulation follows the previous simulations by Butler and Shaqfeh [“Dynamic simulations of the inhomogeneous sedimentation of rigid fibres,” *J. Fluid Mech.* **468**, 205 (2002)]. The motion of the fibers is described using slender-body theory, and the line distribution of point forces along their lengths is approximated by a Legendre polynomial in which only the total force, torque, and particle stresslet are retained. Periodic boundary conditions are used to simulate an infinite suspension, and both far-field hydrodynamic interactions and short-range lubrication forces are considered in all simulations. The calculation of the hydrodynamic interactions, which is typically the bottleneck for large systems with periodic boundary conditions, is accelerated using a smooth particle-mesh Ewald (SPME) algorithm previously used in molecular dynamics simulations. In SPME the slowly decaying Green’s function is split into two fast-converging sums: the first involves the distribution of point forces and accounts for the singular short-range part of the interactions, while the second is expressed in terms of the Fourier transform of the force distribution and accounts for the smooth and long-range part. Because of its smoothness, the second sum can be computed efficiently on an underlying grid using the fast Fourier transform algorithm, resulting in a significant speed-up of the calculations. Systems of up to 512 fibers were simulated on a single-processor workstation, providing a different insight into the formation, structure, and dynamics of the inhomogeneities that occur in sedimenting fiber suspensions. © 2005 American Institute of Physics.

[DOI: 10.1063/1.1862262]

I. INTRODUCTION

Sedimentation of small particles in a fluid occurs in a very wide range of both natural phenomena and industrial processes: sedimentation in a river, pollution in the atmosphere, waste processing, or the production of composite materials are all instances in which particle settling plays a central role. While being one of the simplest hydrodynamic particulate flows, it also remains one of the least well understood, and has been the subject of numerous experimental and theoretical investigations. Batchelor,¹ working on the premise of homogeneity and diluteness, was the first to calculate the correction to the sedimentation speed of non-Brownian spheres in an infinite fluid. His result was confirmed by others,^{2–4} although it appears that the functional dependence of the hindered settling velocity depends on the statistical structure of the suspension.⁵ Later, Caffish and Luke⁶ demonstrated that the average velocity in an infinite suspension has little meaning since the variance is unbounded, growing like the linear dimension of the sedimenting system. This prediction is contradicted by experiments, which suggest that the velocity fluctuations and hydrody-

namic diffusivities do not depend on the size of the container.⁷ This raises the question of whether an initially random suspension remains so during sedimentation. Several interpretations have been proposed, including a screening mechanism for the velocity disturbance of the particles.⁸ More recent investigations^{9–13} suggest that the presence of container walls and of a small stratification in the suspension create a decay of the velocity fluctuations.

The sedimentation of nonspherical particles such as fibers is at least as complicated. Experiments demonstrate a concentration instability, by which an initially homogeneous distribution of fibers can become highly inhomogeneous: high-concentration streamers made of distinct clusters tend to form, surrounded by clarified regions. This clustering has a strong effect on the sedimentation speed, which is enhanced in the dilute regime and can become larger than the maximum possible value for an isolated fiber.^{14,15} This concentration instability was predicted theoretically by the linear stability analysis of Koch and Shaqfeh,¹⁶ which demonstrated that such a clustering for suspensions of nonspherical axisymmetric particles is a consequence of a coupling between the mean flow fluctuations and the anisotropic mobility of the particles. Their linear analysis, however, predicted that the density perturbations with the maximum growth rate

^{a)}Author to whom correspondence should be addressed. Electronic mail: darve@stanford.edu

are those of arbitrarily large wavelength, and hence was not able to provide any information on the wavenumber selection observed in the experiments. Thus, this process is yet to be elucidated.

Several attempts have been made to numerically simulate this instability. Mackaplow and Shaqfeh¹⁷ performed both Monte Carlo simulations and point-particle dynamic simulations, and managed to capture most of the qualitative features of the instability in the dilute regime: formation of streamers with increased sedimentation velocity and alignment of the fibers with gravity with occasional flipping. Their simulations predicted the correct form of the steady-state orientation distribution, but generally overestimated the mean sedimentation speed. This discrepancy was attributed to the influence of the dimensions of their periodic unit cell, and to the absence of lubrication or contact forces in their simulation method, where the latter would have the effect of limiting cluster densification. Their study did not provide any quantitative results on the velocity distributions, pair distributions, or on the influence of the effective particle volume fraction $n\ell^3$.

Butler and Shaqfeh¹⁸ performed more sophisticated simulations in which they accounted for both far-field hydrodynamic interactions and short-range lubrication forces. The fibers were no longer approximated as point particles, but modeled using slender-body theory as line distributions of point forces. A spectral approximation of the line distributions was used, in which they retained the first two moments, namely, the total force and dipole term, whose antisymmetric and symmetric parts are the total torque and particle stresslet. Short-range pairwise interactions were also included using the lubrication approximation. The results they obtained were dependent on the aspect ratio of their periodic simulation box, but good orientation and velocity statistics were obtained using boxes that were highly elongated in the direction of gravity. The size of the systems they were able to simulate was, however, greatly limited by the cost of evaluating the interparticle long-range hydrodynamic interactions, so that only very few clusters were typically observed in their simulations. Moreover the size of the clusters was influenced by the boundary conditions. In particular, the limitations on the box aspect ratio and number of fibers did not allow them to observe more than one streamer in the horizontal direction. To reduce the influence of the boundary conditions, and study in more details the structure of the suspension (e.g., wavenumber selection, cluster size, and dynamics), larger systems need to be simulated: this can only be achieved using efficient algorithms to evaluate the long-range interactions.

In the Stokes flow regime, the disturbance of a point force or Stokeslet on the flow decays as the inverse of the distance from the force location, so that even distant particle pairs can have significant interactions. Evaluating the disturbance velocity at each particle location due to the presence and motion of the other particles is typically an $O(N^2)$ operation, where N is the number of particles, and can become very prohibitive when the size of the system increases. Several alternate methods for the summation of the interactions have been suggested in the last few years, most of which

were inspired by existing algorithms used in molecular dynamics simulations, where the electrostatic potentials have the same slow decay as the disturbance velocity in creeping flow. Sangani and Mo¹⁹ developed an $O(N)$ fast-multipole method for hydrodynamic interactions, based on the famous algorithm by Greengard and Rokhlin.²⁰ In the fast-multipole method the simulation domain is decomposed into a tree of cells, and in each cell a compressed representation of the flow disturbances is constructed using a truncated multipole expansion. Far-field interactions with the particles located in distant cells are then calculated by means of this compressed representation. The fast-multipole method was also applied more recently to suspensions of many deformable drops²¹ and to Stokesian dynamics simulations with finite numbers of particles.²² Another efficient method was also developed by Ladd²³ and is based on a discretized Boltzmann equation for the fluid phase; this technique, however, only applies to finite Reynolds numbers.

The most notable method for both electrostatic and hydrodynamic interactions is the so-called particle-particle-particle-mesh (PPPM) algorithm.²⁴ This algorithm and its many variants have been used for over a decade in plasma simulations, astrophysics, and molecular dynamics, and are based on a decomposition of the interactions into two fast-converging contributions, one of which is efficiently evaluated using an underlying grid and the fast Fourier transform algorithm. A good choice of parameters allows one to reduce the cost of evaluating the interactions to $O(N \ln N)$, which is a significant improvement for large systems. A thorough review of different variants of the method for electrostatic interactions can be found in Deserno and Holm.²⁵

A version of the PPPM algorithm, called particle-mesh Ewald (PME), was first applied to hydrodynamic interactions between suspended spheres by Guckel,²⁶ and was subsequently systematized by Sierou and Brady²⁷ in their accelerated Stokesian dynamics (ASD). ASD is a new version of the Stokesian dynamics method,²⁸ in which the direct calculation of the far-field grand mobility matrices is replaced by the use of the PME algorithm along with an iterative solver for the matrix inversion. ASD was proven to be very efficient in evaluating the rheological properties of suspensions of spheres,²⁹ allowing one to simulate systems of up to 1000 spheres.

In this paper we are implementing a variant of the PME algorithm, called smooth particle-mesh Ewald (SPME), and based on the work by Essman *et al.*³⁰ for electrostatic interactions. The main advantage of SPME over the original PME method is the improved accuracy resulting from the new force assignment and interpolation schemes used for the Fourier sum, which are based on an approximation of structure factors using high-order B -splines. The common points and differences with accelerated Stokesian dynamics will become apparent in the subsequent discussion. Section II gives an overview of the mathematical formulation of the sedimentation problem: the equations for the motion of the fibers, lubrication forces, and hydrodynamic interactions are all presented in detail. We then proceed to explain the smooth particle-mesh Ewald algorithm in Sec. III, and assess its performance by comparing it to the standard Ewald summation

technique. In Sec. IV we discuss the iterative method and preconditioning techniques used to solve for the stresslets and lubrication forces. Finally the results of the simulations are given in Sec. V, where the structure of the suspension, orientation, and velocity statistics are discussed in turn, and we conclude in Sec. VI.

II. MATHEMATICAL FORMULATION

The simulation method follows the previous work by Butler and Shaqfeh.¹⁸ Slender-body theory is used to model the motion of the fibers, along with a spectral approximation of the force distribution in which the total force, torque, and particle stresslet are retained. The fundamental periodic solution of the Stokes equation is used for the hydrodynamic interactions, and short-range interactions are calculated using the lubrication approximation. For more details and discussion the reader is referred to the original paper.

A. Motion of the fibers: Slender-body theory

The position and orientation of each fiber α is given by the location of its center of mass \mathbf{x}_α and a unit vector \mathbf{p}_α parallel to its major axis. The time evolution of \mathbf{x}_α and \mathbf{p}_α can be tracked using Batchelor's slender-body theory,³¹ which represents the disturbance due to the presence and motion of the fiber in the fluid by a line distribution of point forces or Stokeslets. To a leading order approximation in $\ln(2A)$ where A is the fiber aspect ratio, the translational and rotational velocities $\dot{\mathbf{x}}_\alpha$ and $\dot{\mathbf{p}}_\alpha$ are related to the fluid disturbance velocity and force distribution along the fiber by

$$\dot{\mathbf{x}}_\alpha + s\dot{\mathbf{p}}_\alpha - \mathbf{u}'_\alpha(\mathbf{x}_\alpha + s\mathbf{p}_\alpha) = \frac{\ln(2A)}{4\pi}(\mathbf{I} + \mathbf{p}_\alpha \otimes \mathbf{p}_\alpha)\mathbf{f}_\alpha(s), \quad (1)$$

where \mathbf{u}'_α is the disturbance velocity of the fluid surrounding fiber α due to the motion of the other fibers, s is an abscissa along the fiber length, and \mathbf{f}_α is the line distribution of point forces. Equation (1) has been nondimensionalized using the following characteristic velocity, length, and time scales:

$$u_c = \frac{gV\Delta\rho}{4\pi\mu l} \ln(2A), \quad l_c = l, \quad t_c = l_c/u_c = \frac{4\pi\mu l^2}{gV\Delta\rho \ln(2A)},$$

where V and $\Delta\rho$ are the volume and relative density of a fiber, μ is the viscosity of the solvent, and g is the acceleration of gravity. The characteristic length scale l is the fiber half-length, so that positions s along the fibers range from -1 to $+1$. The characteristic time scale, or Stokes time, is the time required for an isolated vertical fiber to sediment over its half-length.

Integrating Eq. (1) along the fiber axis gives an expression for the translational velocity,

$$\dot{\mathbf{x}}_\alpha = \frac{1}{2} \int_{-1}^1 \mathbf{u}'_\alpha(\mathbf{x}_\alpha + s\mathbf{p}_\alpha) ds + \frac{\ln(2A)}{8\pi}(\mathbf{I} + \mathbf{p}_\alpha \otimes \mathbf{p}_\alpha)\mathbf{F}_\alpha, \quad (2)$$

while multiplying Eq. (1) with s , integrating, and taking two cross products with \mathbf{p}_α yields the rotational velocity

$$\dot{\mathbf{p}}_\alpha = \frac{3}{2}(\mathbf{I} - \mathbf{p}_\alpha \otimes \mathbf{p}_\alpha) \int_{-1}^1 s \mathbf{u}'_\alpha(\mathbf{x}_\alpha + s\mathbf{p}_\alpha) ds - \frac{3 \ln(2A)}{8\pi} \mathbf{p}_\alpha \times \mathbf{T}_\alpha. \quad (3)$$

We have introduced the total force and torque on fiber α ,

$$\mathbf{F}_\alpha = \int_{-1}^1 \mathbf{f}_\alpha ds, \quad \mathbf{T}_\alpha = \int_{-1}^1 \mathbf{p}_\alpha \times s \mathbf{f}_\alpha ds.$$

Given \mathbf{F}_α , \mathbf{T}_α , and the fluid velocity \mathbf{u}'_α along the axis due to the motion of the other fibers, Eqs. (2) and (3) can be used to integrate the motion of the fiber in time.

B. Forces and torques on the fibers

In addition to the long-range hydrodynamic interactions which will be discussed in the following section, the forces and torques on the fibers can come from three contributions: gravity which causes sedimentation, lubrication forces when two fibers get close to each other (typically when the distance between their surfaces falls below one fiber diameter), and strong repulsive contact forces when this distance becomes of the order of the roughness length scale of the fiber surfaces. Gravity poses no special difficulty and only creates a force at the center of mass of the particle:

$$\mathbf{F}_\alpha^G = V\Delta\rho\mathbf{g}, \quad \mathbf{T}_\alpha^G = 0.$$

The slender-body formulation of Sec. II A, which represents the fibers as line distributions of point forces, is a far-field approximation and therefore does not capture the near-field effects which are dependent on the exact geometry of the fiber surfaces. These effects can be accounted for through additional lubrication forces, as is commonly done in Stokesian dynamics simulations.²⁸ As sedimenting fiber suspensions are known to form inhomogeneities, correctly capturing these short-range interactions can be critical, and indeed Butler and Shaqfeh¹⁸ observed that including lubrication forces in their simulations modified the sedimentation rates and pair probabilities. The modeling is quite straightforward and is derived from the formulas of Claey's and Brady.³² Four types of lubrication interactions can occur: body-body interactions (parallel and nonparallel cases), end-body interactions, and end-end interactions. In each case the explicit form of the lubrication force will differ, and the various formulas can be found in the work by Butler and Shaqfeh.¹⁸ For the most general case where the fibers interact along their lengths and are nonparallel, the force between two fibers α and β is given by

$$\mathbf{F}_{\alpha\beta}^L = \pm \frac{6\pi\dot{h}}{A^2|\mathbf{p}_\alpha \times \mathbf{p}_\beta|h} \mathbf{n},$$

where \mathbf{n} is a unit vector normal to both fiber surfaces, h is the minimum separation distance between the two surfaces, and \dot{h} is the relative velocity projected along \mathbf{n} . The appropriate sign is chosen such that the force is repulsive for approaching fibers ($\dot{h} > 0$) and attractive for separating fibers ($\dot{h} < 0$). In all cases these lubrication forces also create a torque.

As argued by Harlen, Sundararajakumar, and Koch³³ in their similar simulations of neutrally buoyant fibers, lubrication

tion interactions between high-aspect ratio particles are quite weak and will usually not be sufficient to prevent mechanical contact. Harlen *et al.* treated these contacts by solving a constraint problem and determining exactly the normal reaction forces that prevent fiber crossing. An alternate approach consists in adding strong and very short-range repulsive forces that act at the same points and in the same directions as the lubrication forces:³⁴

$$\mathbf{F}_{\alpha\beta}^R = \pm a_0 \frac{\tau e^{-\tau h}}{1 - e^{-\tau h}} \mathbf{n}.$$

The values of the parameters a_0 and τ used in the simulations are 1×10^{-4} and 1×10^3 , respectively, and are chosen such that the repulsive force remains small over most of the range of the lubrication force; the actual values are typically shown to have little impact on the simulation results.^{18,35}

C. Hydrodynamic interactions

The hydrodynamic interactions are accounted for through the disturbance velocity \mathbf{u}'_α in Eqs. (2) and (3). Since the fibers are approximated as line distributions of point forces, the disturbance velocity is given by

$$\mathbf{u}'(\mathbf{x}_\alpha + s_\alpha \mathbf{p}_\alpha) = \sum_{\beta=1}^N \int_{-1}^1 \mathbf{J}(\mathbf{x}_\alpha + s_\alpha \mathbf{p}_\alpha - \mathbf{x}_\beta - s_\beta \mathbf{p}_\beta) \mathbf{f}_\beta(s_\beta) ds_\beta, \quad (4)$$

where \mathbf{J} is the Green's function, or fundamental solution of the Stokes equation. For a finite number of point forces, \mathbf{J} is the Stokeslet, or Oseen–Burgers tensor \mathbf{K} :

$$\mathbf{K}(\mathbf{x}) = \frac{1}{8\pi} \left(\frac{\mathbf{I}}{r} + \frac{\mathbf{x} \otimes \mathbf{x}}{r^3} \right),$$

where $r = |\mathbf{x}|$. However, when the suspension is infinite, summing an infinite number of Stokeslets becomes intractable and convergence is not guaranteed.³⁶ This issue can be circumvented by using a periodic suspension obtained by replicating a unit cell in all three space dimensions. In the case of a periodic distribution of point forces an exact solution \mathbf{K}_p was derived by Hasimoto.³⁷ The details of the Hasimoto solution are given in Sec. III along with a discussion on the method used to calculate it. Hence for a periodic suspension

$$\mathbf{J}(\mathbf{x}_\alpha + s_\alpha \mathbf{p}_\alpha - \mathbf{x}_\beta - s_\beta \mathbf{p}_\beta) = \begin{cases} \mathbf{K}_p(\mathbf{x}_\alpha + s_\alpha \mathbf{p}_\alpha - \mathbf{x}_\beta - s_\beta \mathbf{p}_\beta) & \text{if } \alpha \neq \beta, \\ \mathbf{K}_p(s_\alpha - s_\beta) - \mathbf{K}(s_\alpha - s_\beta) & \text{if } \alpha = \beta. \end{cases} \quad (5)$$

When $\alpha = \beta$, the Oseen–Burgers tensor must be subtracted from the periodic solution so that the disturbance velocity due to fiber α is not included. When $\alpha = \beta$ and $s_\alpha = s_\beta$ the limit of $\mathbf{K}_p(s_\alpha - s_\beta) - \mathbf{K}(s_\alpha - s_\beta)$ is used.

D. Linearization of the force distribution

The system of equations, as presented in the preceding sections, is an integral system for the force distribution $\mathbf{f}_\alpha(s)$ along each fiber α . In a method similar to the boundary integral method, one may choose to discretize the fibers and solve for the value of \mathbf{f}_α at each of the discretization points.

This was the approach adopted by Mackaplow and Shaqfeh¹⁷ in their Monte Carlo simulations. This method, however, is very expensive as $3M + 5$ unknowns must be solved per fiber, where M is the number of discretization points and can be of the order of 10.

A more tractable approach was suggested by Harlen *et al.*,³³ and used successfully by Butler and Shaqfeh.¹⁸ It consists in expanding the force distribution in a Legendre polynomial and retaining only the first two moments:

$$\mathbf{f}_\alpha(s) \approx \frac{1}{2} \int_{-1}^1 \mathbf{f}_\alpha(s) ds + \frac{3s}{2} \int_{-1}^1 s \mathbf{f}_\alpha(s) ds.$$

Observing that $\mathbf{f}_\alpha = (\mathbf{p}_\alpha \cdot \mathbf{f}_\alpha) \mathbf{p}_\alpha - \mathbf{p}_\alpha \times (\mathbf{p}_\alpha \times \mathbf{f}_\alpha)$ and using the projection of Eq. (3) along \mathbf{p}_α allows to rewrite the linearized force distribution as

$$\mathbf{f}_\alpha(s) \approx \frac{1}{2} \mathbf{F}_\alpha + \frac{3s}{2} (\mathbf{T}_\alpha \times \mathbf{p}_\alpha + \mathcal{S}_\alpha \mathbf{p}_\alpha), \quad (6)$$

where \mathcal{S}_α is the particle stresslet, which is a scalar due to the one dimensionality of the fibers:

$$\mathcal{S}_\alpha = - \frac{2\pi}{\ln 2A} \int_{-1}^1 s \mathbf{p}_\alpha \cdot \mathbf{u}'_\alpha ds. \quad (7)$$

This method is analogous in essence to that used by Claeys and Brady in their Stokesian dynamics simulations of prolate spheroids,³⁸ although the path followed is different. Instead of using a spectral expansion of the force distribution as in our case, Stokesian dynamics is based on a multipole expansion of the Green's function into centered moments, which is typically truncated after the dipole term. Claeys and Brady then argue that in the case of spheroids the multipole moments are equivalent to distributed singularities along the focal axis of the particles. The equations they obtain are slightly more general as they also involve distributions of potential dipoles and their derivatives, which arise from the finite thickness of the particles but become negligible for high-aspect ratios.

E. Method of solution

Substituting the linearized force, Eq. (6), into the expression for the disturbance velocity, Eq. (4), and then Eq. (4) into the equations for the motion of the fibers, Eqs. (2) and (3), yields a general expression for the translational and rotational velocities as functions of the gravity force \mathbf{F}^G (which is the same for all fibers), the stresslets \mathcal{S}_β , and the lubrication forces \mathbf{F}_l^L , where the index $l = 1, \dots, M$ refers to a specific lubrication interaction between two fibers. It can be written in the general form

$$\dot{\mathbf{x}}_\alpha = \sum_{\beta=1}^N [\mathbf{L}_{\alpha\beta} \mathbf{F}^G + \mathbf{M}_{\alpha\beta} \mathcal{S}_\beta] + \sum_{l=1}^M \mathbf{N}_{\alpha l} \mathbf{F}_l^L, \quad (8)$$

$$\dot{\mathbf{p}}_\alpha = \sum_{\beta=1}^N [\mathbf{P}_{\alpha\beta}\mathbf{F}^G + \mathbf{Q}_{\alpha\beta}\mathcal{S}_\beta] + \sum_{l=1}^M \mathbf{R}_{\alpha l}\mathbf{F}_l^L. \quad (9)$$

For a given pair (α, β) , $\mathbf{L}_{\alpha\beta}$, $\mathbf{N}_{\alpha l}$, $\mathbf{P}_{\alpha\beta}$, and $\mathbf{R}_{\alpha l}$ are second-order tensors, and $\mathbf{M}_{\alpha\beta}$ and $\mathbf{Q}_{\alpha\beta}$ are vectors (since the stresslets are scalar quantities).

The tensors $\mathbf{L}_{\alpha\beta}$ and $\mathbf{P}_{\alpha\beta}$ give contribution to the velocities due to gravity, and can be shown to be

$$\mathbf{L}_{\alpha\beta} = \frac{1}{4} \int \int_{-1}^1 \mathbf{K}(s_\alpha, s_\beta) ds_\alpha ds_\beta + \delta_{\alpha\beta} \frac{\ln(2A)}{8\pi} (\mathbf{I} + \mathbf{p}_\alpha \otimes \mathbf{p}_\alpha),$$

$$\mathbf{P}_{\alpha\beta} = \frac{3}{4} (\mathbf{I} + \mathbf{p}_\alpha \otimes \mathbf{p}_\alpha) \int \int_{-1}^1 \mathbf{K}(s_\alpha, s_\beta) s_\alpha ds_\alpha ds_\beta,$$

where in each case $\mathbf{K}(s_\alpha, s_\beta)$ stands for $\mathbf{K}(\mathbf{x}_\alpha + s_\alpha \mathbf{p}_\alpha - \mathbf{x}_\beta - s_\beta \mathbf{p}_\beta)$ and takes the form given in Eq. (5). The vectors $\mathbf{M}_{\alpha\beta}$ and $\mathbf{Q}_{\alpha\beta}$ used to relate the fiber stresslets \mathcal{S}_β to the velocities are also easily obtained:

$$\mathbf{M}_{\alpha\beta} = \frac{3}{4} \int \int_{-1}^1 \mathbf{K}(s_\alpha, s_\beta) \mathbf{p}_\beta s_\beta ds_\alpha ds_\beta,$$

$$\mathbf{Q}_{\alpha\beta} = \frac{9}{4} (\mathbf{I} + \mathbf{p}_\alpha \otimes \mathbf{p}_\alpha) \int \int_{-1}^1 \mathbf{K}(s_\alpha, s_\beta) \mathbf{p}_\beta s_\alpha s_\beta ds_\alpha ds_\beta.$$

The two remaining tensors $\mathbf{N}_{\alpha l}$ and $\mathbf{R}_{\alpha l}$ give the contribution of lubrication forces to the velocity, which is of two types: a direct contribution, through the second term on the right-hand sides of Eqs. (2) and (3) and an indirect contribution through the disturbance velocity. This second contribution is a multibody interaction in the sense that each lubrication force affects all the fibers. Given a lubrication interaction l ($1 \leq l \leq M$), let a and b be the two fibers between which the interaction takes place, and let λ_{la} and λ_{lb} be the corresponding abscissas where the force is applied. If \mathbf{F}_l^L is the force on fiber a , the force on fiber b will be $-\mathbf{F}_l^L$. Introduce the two matrices

$$A_{\alpha l} = \delta_{\alpha a} - \delta_{\alpha b},$$

$$B_{\alpha l} = \delta_{\alpha a} \lambda_{la} - \delta_{\alpha b} \lambda_{lb}.$$

The two tensors $\mathbf{N}_{\alpha l}$ and $\mathbf{R}_{\alpha l}$ then take on the following form:

$$\begin{aligned} \mathbf{N}_{\alpha l} &= \frac{1}{4} \int \int_{-1}^1 \mathbf{K}(s_\alpha, s_\beta) [\mathbf{I} + 3\lambda_{la} s_\alpha (\mathbf{I} - \mathbf{p}_a \otimes \mathbf{p}_a)] ds_\alpha ds_\beta \\ &\quad - \frac{1}{4} \int \int_{-1}^1 \mathbf{K}(s_\alpha, s_\beta) [\mathbf{I} + 3\lambda_{lb} s_\beta (\mathbf{I} - \mathbf{p}_b \otimes \mathbf{p}_b)] ds_\alpha ds_\beta \\ &\quad + \frac{\ln(2A)}{8\pi} (\mathbf{I} + \mathbf{p}_\alpha \otimes \mathbf{p}_\alpha) A_{\alpha l}, \end{aligned}$$

$$\begin{aligned} \mathbf{R}_{\alpha l} &= \frac{3}{4} (\mathbf{I} - \mathbf{p}_\alpha \otimes \mathbf{p}_\alpha) \left[\int \int_{-1}^1 \mathbf{K}(s_\alpha, s_\beta) [\mathbf{I} + 3\lambda_{la} s_\alpha (\mathbf{I} - \mathbf{p}_a \otimes \mathbf{p}_a)] ds_\alpha ds_\beta \right. \\ &\quad \left. - \int \int_{-1}^1 \mathbf{K}(s_\alpha, s_\beta) [\mathbf{I} + 3\lambda_{lb} s_\beta (\mathbf{I} - \mathbf{p}_b \otimes \mathbf{p}_b)] ds_\alpha ds_\beta + \frac{\ln(2A)}{2\pi} B_{\alpha l} \right]. \end{aligned}$$

The repulsive contact forces, which are not included in the above equations, are treated in exactly the same way as the lubrication forces.

The method of solution proceeds as follows. In Eqs. (8) and (9) the tensors $\mathbf{L}_{\alpha\beta}$, $\mathbf{N}_{\alpha l}$, $\mathbf{P}_{\alpha\beta}$, and $\mathbf{R}_{\alpha l}$ and vectors $\mathbf{M}_{\alpha\beta}$ and $\mathbf{Q}_{\alpha\beta}$ are only functions of the positions and orientations of the fibers and are hence known. The gravitational force \mathbf{F}^G and the repulsive forces \mathbf{F}_l^R are also known, so that the unknowns are the translational and rotational velocities $\dot{\mathbf{x}}_\alpha$ and $\dot{\mathbf{p}}_\alpha$, of the fibers, as well as the stresslets \mathcal{S}_α and lubrication forces \mathbf{F}_l^L . The first step consists in solving for the latter two in the following manner. Substituting the linearized force distribution [Eq. (6)] (where the total force and torque are functions of the lubrication interactions and stresslets which are unknown) into the disturbance velocity, Eq. (4), and in turn substituting Eq. (4) into the definition of the stresslets, Eq. (7), yields a linear system for the stresslets and magnitudes of the lubrication forces. Once this system is inverted, using an iterative solver discussed in Sec. IV, Eqs. (8) and (9) can be used to obtain the velocities of the fibers. The positions of the fibers can then be advanced using a fourth-order Runge–Kutta time-marching method. The time step is chosen so as to avoid collisions or fiber overlap, and so that no fiber moves by more than half a fiber diameter at each step. For a more extensive discussion on the time integration method, the reader is referred to Butler and Shaqfeh.¹⁸

III. THE SMOOTH PARTICLE-MESH EWALD ALGORITHM

A. Periodic fundamental solution and Ewald summation formula

The disturbance velocity Eq. (4) must be evaluated along each fiber. In practical simulations the integrals over the lengths of the fibers are computed numerically using Gauss–Legendre quadrature, so that we are left with calculating the disturbance field created by a distribution of point forces at each point force location. Consider a distribution of N point forces $\mathbf{F}_1, \mathbf{F}_2, \dots, \mathbf{F}_N$ at positions $\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_N$ in a unit cell of volume τ_0 ; to alleviate the notations, assume that the quadrature weights resulting from the discretization of the integrals have been included in the force vectors \mathbf{F}_n . Denote by \mathbf{a}_i , $i=1, 2, 3$, the lattice vectors forming the edges of the unit cell. To simulate an infinite suspension we use periodic boundary conditions, so that each point force \mathbf{F}_n at position \mathbf{x}_n has periodic images at all the locations $\mathbf{x}_n + p_1 \mathbf{a}_1 + p_2 \mathbf{a}_2 + p_3 \mathbf{a}_3$ for all integers p_1, p_2, p_3 . With these notations, the disturbance velocity at location \mathbf{x}_m created by the point forces other than \mathbf{F}_m is written as

$$\mathbf{u}(\mathbf{x}_m) = \sum_{\mathbf{p}} \sum_{n=1}^N \mathbf{K}(\mathbf{x}_n - \mathbf{x}_m + \mathbf{p}) \mathbf{F}_n, \quad m = 1, \dots, N, \quad (10)$$

where \mathbf{K} is the Oseen–Burgers tensor. The term corresponding to $n=m$ and $\mathbf{p}=\mathbf{0}$ is omitted in the sum. Because of the decay of the Oseen–Burgers tensor as $1/r$, the infinite sum in Eq. (10) is generally divergent. This is typically remedied by realizing that the average force $\langle \mathbf{F} \rangle \neq \mathbf{0}$ creates a backflow, and that only the velocity relative to this backflow has a physical significance.^{28,36} Hasimoto³⁷ took advantage of the spatial periodicity of the problem and used Fourier series to derive a convergent expression for Eq. (10). His solution can be written as

$$\mathbf{u}(\mathbf{x}_m) = \frac{1}{4\pi} [\mathbf{S}^1 - \nabla(\nabla \cdot \mathbf{S}^2)], \quad (11)$$

where \mathbf{S}^1 and \mathbf{S}^2 are the following sums:

$$\begin{aligned} \mathbf{S}^2 &= -\frac{1}{4\pi^3 \tau_0} \sum_{\mathbf{k} \neq \mathbf{0}} \frac{e^{2\pi i \mathbf{k} \cdot \mathbf{x}_m}}{k^4} \hat{\mathbf{F}}(\mathbf{k}), \\ \mathbf{S}^1 &= \nabla^2 \mathbf{S}^2 = \frac{1}{\pi \tau_0} \sum_{\mathbf{k} \neq \mathbf{0}} \frac{e^{2\pi i \mathbf{k} \cdot \mathbf{x}_m}}{k^2} \hat{\mathbf{F}}(\mathbf{k}). \end{aligned} \quad (12)$$

The vectors \mathbf{k} are the reciprocal lattice unit vectors: $\mathbf{k} = k_1 \mathbf{b}_1 + k_2 \mathbf{b}_2 + k_3 \mathbf{b}_3$, where the vectors \mathbf{b}_i , $i=1, 2, 3$, define the unit cell in the reciprocal or Fourier space:

$$\mathbf{b}_1 = \frac{\mathbf{a}_2 \times \mathbf{a}_3}{\tau_0}, \quad \mathbf{b}_2 = \frac{\mathbf{a}_3 \times \mathbf{a}_1}{\tau_0}, \quad \mathbf{b}_3 = \frac{\mathbf{a}_1 \times \mathbf{a}_2}{\tau_0}.$$

$\hat{\mathbf{F}}(\mathbf{k})$ is the Fourier transform of the distribution of point forces, or structure factor:

$$\hat{\mathbf{F}}(\mathbf{k}) = \sum_{n=1}^N \mathbf{F}_n e^{2\pi i \mathbf{k} \cdot \mathbf{x}_n}. \quad (13)$$

The presence of this structure factor in the periodic fundamental solution suggests that the disturbance velocity could be evaluated using the fast Fourier transform algorithm; this observation will be exploited below. It is important to note the absence of the term $\mathbf{k}=\mathbf{0}$ in the Fourier representation of the velocity: this term corresponds to the mean backflow alluded to above and should be set to zero if the fluid is globally quiescent. The absence of this term arises naturally in the derivation of Eq. (11), where it is shown that the mean pressure gradient cancels exactly the mean force term $\hat{\mathbf{F}}(\mathbf{0})$, so that the latter does not contribute to the disturbance velocity.³⁷

Equations (11) and (12) constitute an exact and absolutely convergent expression for the disturbance velocity, and can be used as such in computations. However, the relatively slow decay in $1/k^2$ makes this direct method quite inefficient as many terms are needed to achieve a reasonable accuracy. The convergence can be accelerated by recasting Eqs. (11) and (12) into a slightly different form called Ewald summation formula. It was first used by Ewald for electrostatic interactions,³⁹ and the present form for hydrodynamic interactions is due to Hasimoto.³⁷ The starting point is an integral representation for $1/k^{2m}$ for $m=1, 2$, and the introduction in

the integral of a cutoff parameter α , called Ewald coefficient. The derivation is classic^{27,30,37,39} and is not repeated here. The final result expresses the disturbance velocity as follows:

$$\begin{aligned} \mathbf{u}(\mathbf{x}_m) &= \sum_{\mathbf{p}} \sum_{n=1}^N \mathbf{A}(\alpha, \mathbf{x}_m - \mathbf{x}_n + \mathbf{p}) \mathbf{F}_n \\ &\quad + \sum_{\mathbf{k} \neq \mathbf{0}} e^{-2\pi i \mathbf{k} \cdot \mathbf{x}_m} \mathbf{B}(\alpha, \mathbf{k}) \hat{\mathbf{F}}(\mathbf{k}), \quad m = 1, \dots, N. \end{aligned} \quad (14)$$

The two tensors \mathbf{A} and \mathbf{B} are given by

$$\begin{aligned} \mathbf{A}(\alpha, \mathbf{x}) &= \frac{\pi}{\alpha^{3/2}} \phi_{1/2} \left(\frac{\pi r^2}{\alpha} \right) (r^2 \mathbf{I} + \mathbf{x} \otimes \mathbf{x}) - \frac{2}{\alpha^{1/2}} e^{-\pi r^2 / \alpha} \mathbf{I}, \\ \mathbf{B}(\alpha, \mathbf{k}) &= \frac{\pi \alpha^2}{\tau_0} \phi_1(\pi \alpha k^2) (k^2 \mathbf{I} - \mathbf{k} \otimes \mathbf{k}), \end{aligned}$$

where the functions ϕ_ν are incomplete γ functions:

$$\phi_{1/2}(x) = \frac{e^{-x}}{x} + \frac{1}{2x} \operatorname{erfc}(\sqrt{x}), \quad \phi_1(x) = \frac{e^{-x}}{x^2} (1+x).$$

The Ewald coefficient α is a user-defined parameter that determines the relative importance of the two sums: its choice is typically dictated by cost considerations. The two tensors \mathbf{A} and \mathbf{B} decay exponentially in $r^2 = |\mathbf{x}|^2$ and $k^2 = |\mathbf{k}|^2$, respectively, so that both sums in Eq. (14) converge rapidly.

B. Description of the algorithm

The Ewald summation formula Eq. (14) is the basis for most simulations with periodic boundary conditions. A standard and widely used method, sometimes called Ewald summation technique, consists in truncating both sums after a few terms and choosing the coefficient α so as to minimize the overall cost.^{17,18,36} In spite of the rapid convergence of the sums, this direct method can still be quite expensive for very large systems. Also, it does not typically exploit the presence of the Fourier transform in the second sum. The smooth particle-mesh Ewald method proceeds quite differently: the parameter α is chosen so as to reduce the cost of the first sum (or real sum), and a fast algorithm based on the fast Fourier transform on a Cartesian grid is used for the second sum (or Fourier sum). More details are provided in the next sections.

1. Real sum

The calculation of the real sum at all the point force locations \mathbf{x}_m is *a priori* an $O(N^2)$ operation as it involves summing over all the other point forces and their images. However, the exponential decay of the tensor \mathbf{A} can be exploited to restrict the evaluation of the sum to close particle pairs. Given a tolerance ϵ and a cutoff radius r_c , the Ewald coefficient α can be chosen to make all the coefficients of $\mathbf{A}(\alpha, \mathbf{x})$ less than ϵ whenever $|\mathbf{x}| \geq r_c$. Once α is obtained following this procedure, the real sum in Eq. (14) only needs to be performed over the point forces located within a sphere of radius r_c , i.e., over a small number independent of the system size. The evaluation of the sum at all the point force

locations \mathbf{x}_m then has an $O(N)$ cost, with a constant of proportionality that can be adjusted through the cutoff radius.

2. Fourier sum

The real gain is achieved in the evaluation of the Fourier sum. As mentioned previously, the presence of the structure factors suggests use of the fast Fourier transform algorithm. As the point forces can be located anywhere inside the unit cell, prior interpolation to a Cartesian grid is necessary. This can be done formally using Cardinal B -splines, which are introduced briefly in the Appendix. A mesh is defined inside the reciprocal unit cell by choosing three integers K_1, K_2 , and K_3 corresponding to the numbers of points along the reciprocal lattice vectors $\mathbf{b}_1, \mathbf{b}_2$, and \mathbf{b}_3 . Given a point \mathbf{x} in the unit cell, we define its scaled fractional coordinates by $\xi_i = K_i \mathbf{b}_i \cdot \mathbf{x}$ for $i=1, 2, 3$.

We wish to approximate the structure factor $\hat{\mathbf{F}}(\mathbf{k})$ using a discrete Fourier transform on the grid defined above. This is achieved by interpolating the complex exponentials in the definition of $\hat{\mathbf{F}}(\mathbf{k})$ using Cardinal B -splines. The Appendix shows that in one dimension we have

$$\exp\left(2\pi i \frac{k\xi}{K}\right) \simeq b(k) \sum_{m=-\infty}^{+\infty} M_p(\xi - m) \exp\left(2\pi i \frac{km}{K}\right), \quad (15)$$

where the coefficients $b(k)$ and M_p can be computed using a recursive definition. p is the order of interpolation, and the sum over m is in fact limited to p terms as the functions M_p have compact support. Generalizing Eq. (15) to three dimensions we obtain the following approximation for the structure factor:

$$\hat{\mathbf{F}}(\mathbf{k}) \simeq b_1(k_1)b_2(k_2)b_3(k_3)\mathcal{F}(\mathbf{Q})(\mathbf{k}), \quad (16)$$

where $\mathcal{F}(\mathbf{Q})$ is the three-dimensional discrete Fourier transform of the following array $\mathbf{Q}(\mathbf{m})$:

$$\begin{aligned} \mathbf{Q}(\mathbf{m}) = & \sum_{n=1}^N \sum_{p_1, p_2, p_3} \mathbf{F}_n M_p(\xi_1^n - m_1 + p_1 K_1) M_p(\xi_2^n - m_2 \\ & + p_2 K_2) M_p(\xi_3^n - m_3 + p_3 K_3). \end{aligned} \quad (17)$$

The algorithm for the computation of the Fourier sum can now be summarized. The first step is the assignment of the point forces \mathbf{F}_n to the Cartesian grid using Cardinal B -splines, i.e., the calculation of the array $\mathbf{Q}(\mathbf{m})$ using the interpolation formula Eq. (17). The discrete Fourier transform $\mathcal{F}(\mathbf{Q})$ is then computed using the fast Fourier transform algorithm (FFT), and is multiplied in turn by $b_i(k_i)$ to yield the approximation to the structure factor Eq. (16), and by the Fourier convolution kernel $\mathbf{B}(\alpha, \mathbf{k})$. The result is multiplied by $b_i^*(k_i)$ (where $*$ denotes the complex conjugate) and the inverse FFT is applied, yielding the Fourier sum in Eq. (14) but evaluated at the grid points. The sum can then be interpolated from the grid points to the particle locations, again using B -splines [transpose operation of the force assignment Eq. (17)].

As we explained above, the real sum need only be performed on a small number of neighbors for each point force, resulting in an $O(N)$ cost for N point forces. The cost of

computing the Fourier sum using the previous algorithm is limited by the two fast Fourier transforms, which have a cost of $O(K \ln K)$ where K is the number of points in each direction on the Cartesian mesh, and is chosen proportional to the system size N . Therefore the total cost of the method scales as $N \ln N$, which can be a significant improvement for large systems compared to the original cost of $O(N^2)$. Comparisons of the CPU times for the traditional Ewald summation technique and the smooth particle-mesh Ewald algorithm are presented in Sec. III D.

The use of the Cardinal B -splines for the force assignment and interpolation schemes is the major difference between SPME and the accelerated Stokesian dynamics method of Sierou and Brady.²⁷ Accelerated Stokesian dynamics uses a Taylor series expansion to assign the forces to the grid, after which it applies the fast Fourier transform algorithm; Lagrange interpolation is then used to go back from the grid to the particle locations. Exploiting instead the interpolation properties of Cardinal B -splines for complex exponentials gives directly a smooth approximation to the structure factors,^{30,40} which are the relevant quantities appearing in the Fourier sum of the Ewald summation formula. The final interpolation is done again using B -splines and is completely analogous to the force assignment, as it should be considering the symmetric nature of the two operations. Another advantage of using B -splines is their smooth behavior at high interpolation order, while Lagrange interpolation is known to become unstable as the order increases. Essmann *et al.*³⁰ compared both methods for electrostatic interactions, and obtained better accuracies using B -splines for both the forces and interaction energies.

C. Accuracy

The accuracy and efficiency of the method depend on several parameters: the tolerance ϵ and cutoff radius r_c for the real sum (the two of which uniquely define the Ewald coefficient α , as explained in Sec. III B 1), and the number of grid points K and the order p of the B -spline interpolation for the evaluation of the Fourier sum. These parameters are typically adjusted to maximize the accuracy while minimizing the cost.

To investigate the accuracy of the method and the influence of the parameters in more detail, tests are performed in a square box with a distribution of 100 point forces at random locations, with random orientations and unit strengths. In all the tests, the value of the real-space tolerance is set to $\epsilon=10^{-10}$. The disturbance velocity at the force locations is computed using the smooth particle-mesh Ewald method, and is compared to the solution obtained with converged Ewald sums. The measure of accuracy used here is the root-mean square (rms) error between the two solutions:

$$\Delta u = \left[\frac{1}{N} \sum_{n=1}^N [\mathbf{u}^{\text{SPME}}(\mathbf{x}_n) - \mathbf{u}^{\text{EWALD}}(\mathbf{x}_n)]^2 \right]^{1/2}.$$

Since our method is based on evaluation of sums such as Eq. (4) which only involve point force distributions, we limit our attention to this measure of accuracy. It should be noted, however, that such a measure does not guarantee that other

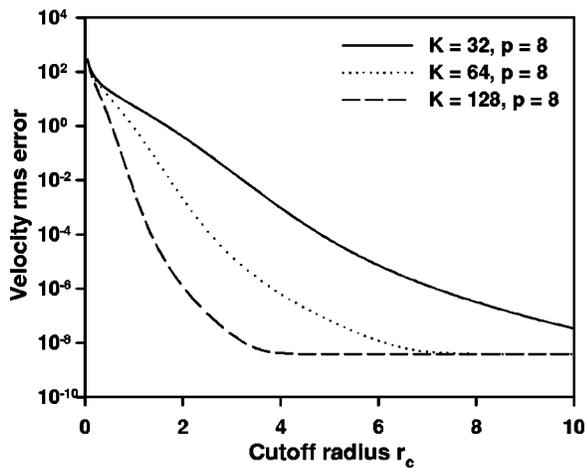


FIG. 1. Velocity rms error as a function of the cutoff radius r_c . The three curves correspond to the three different grid sizes $K=32$, 64, and 128 for the evaluation of the Fourier sum.

quantities such as higher force moments are as adequately captured.

The rms error as a function of the cutoff radius at constant values of K and p is plotted in Fig. 1. The three curves correspond to the three grid sizes $K=32$, 64, and 128. In all cases the rms error becomes large as $r_c \rightarrow 0$, but decays very rapidly when r_c increases; for large values of r_c , it reaches a plateau at around $\Delta u = 10^{-9}$. These curves are easily interpreted. Whatever the value of the cutoff radius, the error in the real sum is of order ϵ and accounts for the plateau at around 10^{-9} : indeed the Ewald coefficient is chosen precisely so as to make the real sum converge within ϵ . The remainder of the error (i.e., all the values above the plateau) stems from the evaluation of the structure factor using the discrete Fourier transform, and from the assignment of the point forces to the grid and back. The Fourier sum accounts for the smooth and long-range part of the solution and cannot capture the short-range singular part; this explains why for a given number of grid points K the error increases rapidly when the cutoff and hence the Ewald coefficient become small.

The influence of K and p on the accuracy is shown in Fig. 2. As already suggested in Fig. 1, increasing the number of points for the fast Fourier transform improves the accuracy of the Fourier sum: Fig. 2(a) shows the decay of the error with increasing K at constant cutoff radius and interpolation order. On a log-log plot the curves are essentially straight, suggesting an algebraic decay; the exponents obtained from the curves in Fig. 2(a) are, respectively, -6.7 and -8.5 and are close to the value of the interpolation order p , confirming the scaling of the error with K^{-p} for the B -spline interpolation.^{30,40} Another way of making the Fourier sum more accurate is to choose a higher order of interpolation for the force assignment scheme, and is illustrated in Fig. 2(b). Increasing K or p has the same effect on the accuracy, so that either one or both can be done and the final decision depends on the computational expense.

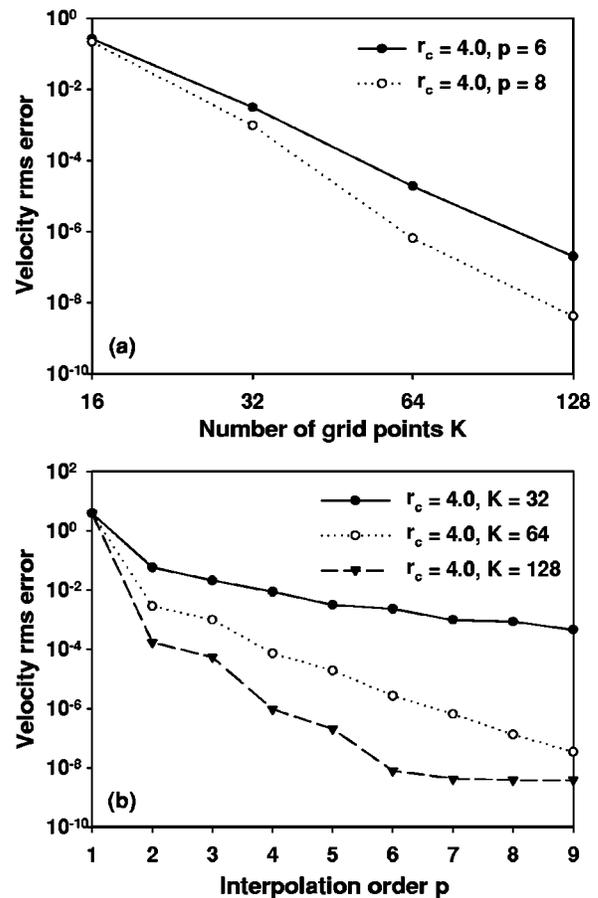


FIG. 2. Velocity rms error as a function of (a) the number of grid points K used for the evaluation of the Fourier sum and (b) the interpolation order p of the force assignment scheme.

D. Efficiency

The efficiency of the method is assessed by comparing the CPU times required by the traditional Ewald summation technique and our smooth particle-mesh Ewald algorithm. These are shown in Fig. 3, where simulations were performed for the same systems as in the preceding section (unit-strength forces with random locations and orientations), for different system sizes N and levels of accuracy. A quick observation of Fig. 3 suffices to show the superiority of the SPME method in terms of efficiency. Except for very small systems, the CPU times are all smaller for SPME than for the Ewald sums at a given accuracy. In fact, even the high-accuracy SPME is more advantageous than the low-accuracy Ewald sums. It becomes even more so for large systems: at the same level of accuracy (root-mean square error of the order of 10^{-9}), SPME is about 40 times faster than the Ewald sums for 1000 point forces and 300 times faster for 5000 points forces. While the cost scales quadratically with size for the Ewald sums, it increases almost linearly with SPME when the grid size is fixed in Fourier space (the three curves for SPME correspond to three different grid resolutions in Fourier space).

It must still be noted that the CPU times of Fig. 3 are for the calculation of the disturbance velocity only. To assess the actual cost of our simulations, we should also take into ac-

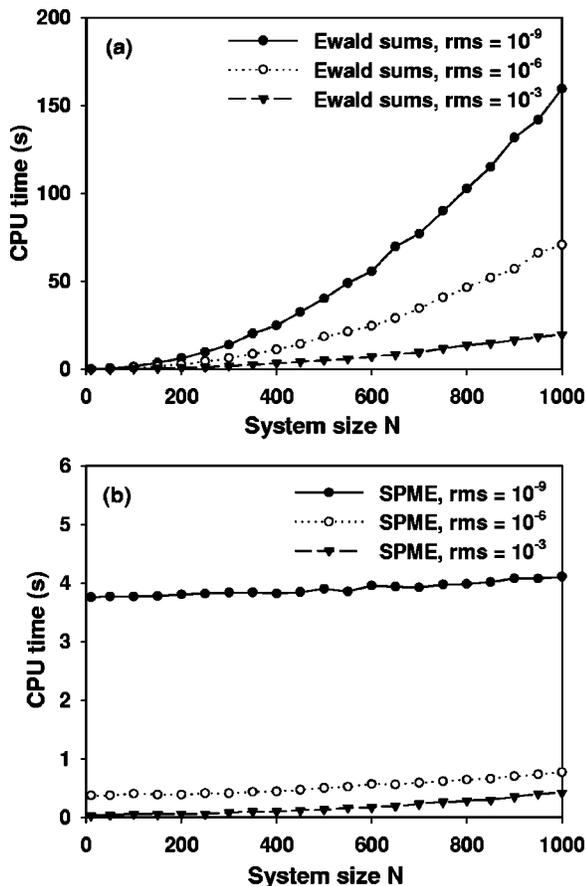


FIG. 3. CPU times for the calculation of the Ewald sums using (a) the traditional Ewald summation technique and (b) the SPME method, as a function of the system size.

count the inversion of the linear system for the stresslets and lubrication forces, which cannot be performed using standard elimination procedures when SPME is used. The additional cost of this system inversion does somewhat impede the performance of SPME, but for large systems it still remains far superior to the standard Ewald sums. A detailed discussion on the system inversion follows.

IV. THE ITERATIVE SOLVER

A. Generalities

Solving for the particle stresslets and amplitudes of the lubrication forces involves inverting a linear system. This operation, which in general can be quite costly, becomes more complicated when the SPME method is used. Indeed in SPME the coefficients in the Ewald summation formula Eq. (10) are not computed directly, and the reciprocal sum instead is calculated as a whole using the fast Fourier transform. In other words, the matrix that needs to be inverted is not explicitly accessible: the SPME algorithm provides us with a “black box” that performs matrix-vector multiplications in an efficient manner. Classical methods for inverting linear systems such as the LU factorization used by Butler and Shaqfeh¹⁸ are therefore not applicable, and we must resort to a different approach. An attractive method is the use of an iterative solver to compute an approximate solution,

TABLE I. Dependence of the condition number on the inhomogeneity of the suspension and on the presence of lubrication interactions. The results are for 200 fibers, $A=11$, $nl^2=0.1$, and a box of aspect ratio $d_x:d_y:d_z=1:1:2$.

Time t (Stokes units)	Number of lubrication interactions	Condition number
0	19	8
50	134	55
100	218	381
150	251	729

such as the generalized minimum residual (GMRES) method of Saad and Schultz,⁴¹ which applies to large nonsymmetric linear systems. At each iteration of the solver, only one matrix-vector multiplication is needed, so that this type of method can be used in conjunction with SPME.⁵⁴

In SPME, the real sum in Eq. (10) is still computed directly, and is written as the multiplication of a matrix \mathbf{A} with a vector containing the point forces. The coefficients of the matrix depend on the spatial configuration of the suspension, and therefore do not change from one solver iteration to the next within a given time step. This seemingly benign observation allows one to precompute the real-space matrix before applying GMRES, so that at each iteration of the solver the real sum is simply obtained by a matrix-vector multiplication of a stored matrix with the current iterate. This spares the expense of computing the real sum coefficients at each step and reduces the time of the algorithm by an order of magnitude, the only downside being the additional storage of the matrix.

B. Preconditioning

To benefit fully from the efficiency of the SPME method, the number of iterations required by the solver to achieve a reasonable convergence must be kept to a minimum. This number of iterations is a function of the condition number of the linear system, defined as the ratio of the largest over the smallest eigenvalue. For our problem, the condition number was shown not to increase significantly with the problem size at a given fiber concentration, but rather to depend on the spatial configuration of the fibers and on the presence of lubrication forces: for very dense or inhomogeneous suspensions with many lubrication interactions (such as the ones that occur when the sedimenting fibers form clusters), the linear system is typically quite stiff and a large number of iterations is required. This is illustrated in Table I, which shows how the condition number increases over the course of a simulation as a result of the concentration instability. Sangani and Mo¹⁹ already encountered this issue with their fast-multipole method, where they observed that the inclusion of lubrication forces in their simulations of sedimenting spheres greatly decreased their convergence rates. A similar problem was reported in accelerated Stokesian dynamics simulations.²⁷

The classic approach to accelerate convergence is the preconditioning of the system, which consists in multiplying it by an approximate inverse of the original matrix. Numer-

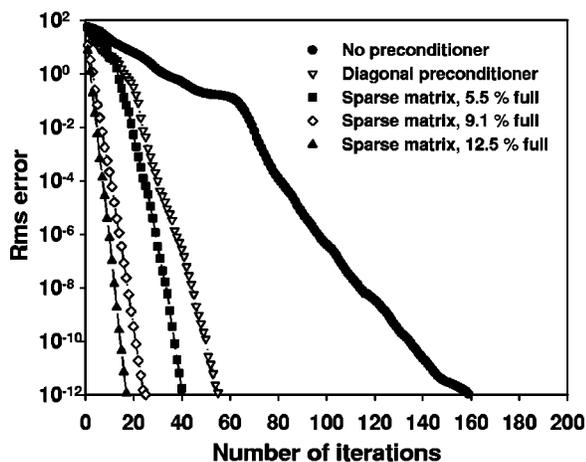


FIG. 4. Convergence of the GMRES solver depending on the preconditioner, for an inhomogeneous system of 200 fibers, $A=11$, $nl^3=0.1$, and a box aspect ratio of $d_x:d_y:d_z=1:1:2$.

ous methods are available:⁴² sparse approximate inverse, incomplete LU factorization, etc., which are typically very efficient for large sparse matrices. Our problem, however, is quite atypical in a few ways. First the coefficients of the matrix are not known if SPME is used: applying any kind of preconditioner will require calculating some of these coefficients using direct summation of the Ewald sums, which is an expensive operation as was established earlier. Second, the matrix is full, whereas most preconditioners are designed for the sparse systems arising in finite difference and finite element codes. Finally, the matrix is relatively small, from a few hundred to a few thousand rows and columns: unlike most situations where iterative solvers are used, the limiting factor is not the system size but the difficulty of constructing the matrix efficiently.

In accelerated Stokesian dynamics Sierou and Brady²⁷ showed that using an incomplete Cholesky preconditioner greatly improved the efficiency of their solver. This method, however, is only valid for symmetric matrices, which is not the case of the system in our simulation method. Other preconditioning techniques were therefore investigated, and the convergence rates are compared in Fig. 4. The first method

consists in calculating the diagonal of the matrix, which is very dominant, and preconditioning the system by this diagonal. This alone reduces the number of iterations by a factor of 2 or 3, but is not quite sufficient for very inhomogeneous systems. The next step consists in calculating a sparse approximation of the matrix, trying to compute the terms of large magnitude. The relative importance of the different terms in the matrix depends on the configuration of the suspension and cannot be known exactly *a priori*, but a good rule of thumb is to assume that the interactions between close particles dominate. A sparse approximate matrix is therefore obtained by only computing the interactions between particles within a cutoff distance d_c , which is chosen by trial and error to minimize the number of iterations while keeping the cost of computing the approximate matrix small. This sparse matrix is then inverted, either exactly using LU factorization or approximately using incomplete LU factorization, and the inverse is used to precondition GMRES. A good choice of d_c can reduce the number of iterations by an order of magnitude, and cut the total time for the system inversion by a factor of 4 for a system of 200 fibers. Table II gives more details on the efficiency of these preconditioners. In all cases the terms in the approximate matrix must be calculated directly using Ewald summation; however, since the sparse matrix is only approximate the Ewald summation need not be fully converged and the first few terms in the sums are typically sufficient to improve convergence.

V. SIMULATION RESULTS AND DISCUSSION

A. General remarks

This section presents some simulation results, all obtained on a single-processor workstation. We first show some comparisons with the Monte Carlo simulations of Mackaplow and Shaqfeh¹⁷ for random dispersions. Mackaplow and Shaqfeh used slender-body theory combined with a boundary integral formulation to compute the sedimentation rate of fixed random arrays of fibers: comparing our results to theirs is therefore a good way of evaluating the consequences of the force linearization described in Sec. II D. Figure 5 shows sedimentation rates for random arrays of high-aspect ratio

TABLE II. Efficiency of various preconditioning techniques for an inhomogeneous system of 200 fibers of aspect ratio $A=11$ at $nl^3=0.1$ in a box of aspect ratio $d_x:d_y:d_z=1:1:2$. Using fuller and fuller matrix approximations decreases the number of GMRES iterations, but increases the overhead cost of computing the sparse matrix. For this example, the optimal value for the cutoff distance is around $d_c=1.2$.

Type of preconditioner	Number of GMRES iterations	Number of terms in the sparse matrix	CPU time to compute the sparse matrix (s)	Total CPU time to solve the system (s)
No preconditioner	138	558
Diagonal preconditioner	56	451 (0.2 %)	12	242
Sparse matrix, $d_c=0.4$	46	11 237 (5.5 %)	12	232
Sparse matrix, $d_c=0.8$	26	18 432 (9.1 %)	27	154
Sparse matrix, $d_c=1.2$	19	25 260 (12.5 %)	48	144
Sparse matrix, $d_c=1.6$	14	32 047 (15.7 %)	75	148
Sparse matrix, $d_c=2.0$	13	38 580 (18.9 %)	101	171

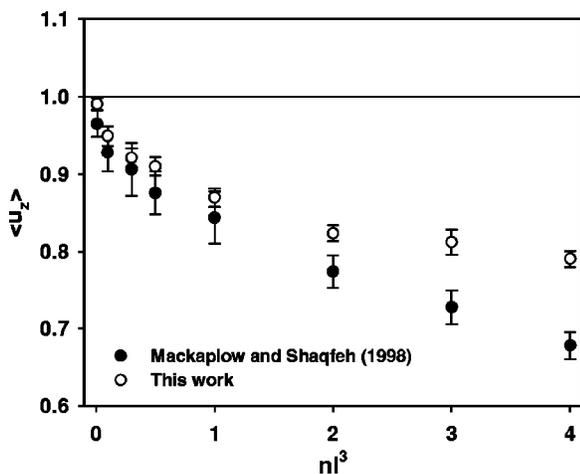


FIG. 5. Sedimentation rates of fixed random arrays of fibers as a function of the effective volume fraction nl^3 , obtained with our simulation method (open symbols) and by Mackaplow and Shaqfeh (Ref. 17) (full symbols). The sedimentation rates are scaled by their value at infinite dilution. The results were obtained by averaging over 20 random configurations of ≈ 300 fibers in a square unit cell. The error bars are 95% confidence intervals.

particles ($A=100$) as a function of the effective volume fraction nl^3 (where n is the particle number density and l is the fiber half-length); the velocities are normalized by the velocities at infinite dilution. Our method slightly overestimates the sedimentation rates, but the agreement is within error bars up to $nl^3=1$, which is quite remarkable as the force linearization cannot be expected to perform well at high concentrations. Other differences between the two methods may also contribute to the differences observed in the sedimentation rates: Mackaplow and Shaqfeh retained an additional term in the slender-body theory asymptotic expansion Eq. (1), but did not include lubrication interactions which should become more and more important at high concentrations. Given these differences, the agreement with our data shown in Fig. 5 is quite good, and we can expect our method to perform well in the dilute regime.

The code was also benchmarked by comparison with the previous work of Butler and Shaqfeh¹⁸ for small systems of up to 128 fibers: the results obtained with our implementation were in all respects similar, so that only larger simulations of 512 fibers are discussed below. For ease of comparison with the experimental results of Herzhaft and Guazzelli,¹⁵ all the simulations presented in the next sections are for a fiber aspect ratio of $A=11$ and an average effective volume fraction of $nl^3=0.05$. A systematic investigation of the effects of slenderness and volume fraction will be deferred to a subsequent study.

Unless otherwise mentioned, all the initial distributions were obtained by assigning the fibers to random positions and with random orientations. In the event of two fibers intersecting, one of them was repositioned at another random location in the box. This ensured that the initial suspensions were homogeneous, as would be the case with a well-mixed suspension in an experiment. As already observed on small systems, the local concentration fluctuations inherently present in such a random distribution cause the suspension to evolve towards very inhomogeneous states where the fibers

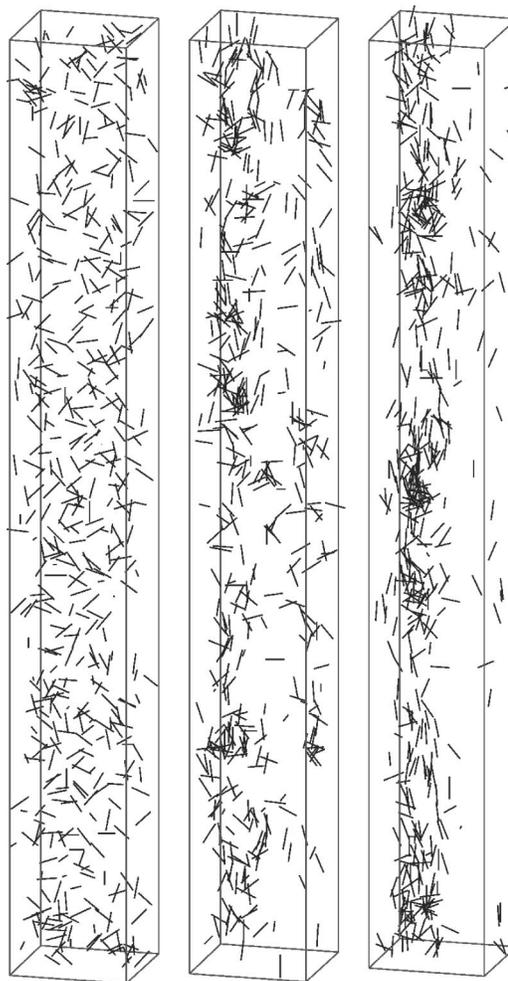


FIG. 6. Fiber distribution in the simulation box at $t=0.0$ (random homogeneous distribution), $t=60.0$, and $t=120.0$ (left to right). The simulation is for 512 fibers of aspect ratio $A=11$ with an average concentration of $nl^3=0.05$. The box aspect ratio is $d_x:d_y:d_z=1:1:8$.

form concentrated streamers surrounded by clarified fluid. By adjusting the dimensions of the periodic unit cell, we were able to observe the formation of one to two or three streamers for the system sizes considered here. We study in turn the structure of the suspension in the vertical direction (inside a given streamer) and in the horizontal direction (formation of distinct streamers).

B. Suspension microstructure and cluster formation

1. Vertical structure

Figure 6 shows the evolution of the suspension for a simulation box of high-aspect ratio ($d_x:d_y:d_z=1:1:8$). Initially small clusters of only a few fibers form at random locations in the simulation box. These clusters have an increased sedimentation rate and entrain the fluid around them, creating a backflow in the other areas, and eventually converge to form a streamer of high velocity. The correlation between the position of the streamers and the fluid velocity is very obvious, as shown in Fig. 7. Because the fluid is globally quiescent, a relatively strong backflow exists outside the streamer, capable of carrying isolated fibers upwards.

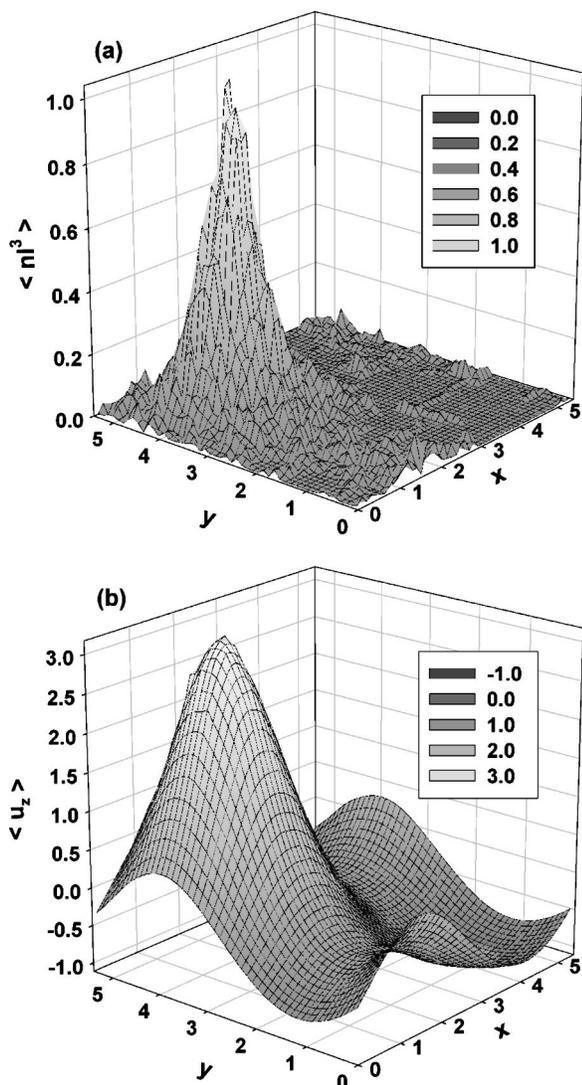


FIG. 7. Vertically averaged local number density $n l^3$ (a) and fluid vertical velocity $\langle u_z \rangle$ (b) for the simulation of Fig. 6. Positive values of $\langle u_z \rangle$ correspond to downward velocities. The position of the streamer is where the two maxima occur. Note the negative velocities (backflow) outside the streamer.

As can be observed in Fig. 6, the streamers are not uniform in the vertical direction, but composed of several distinct clusters of various sizes. Because they have different sizes, they also have different sedimentation velocities, so that their relative positions inside the streamers change: large clusters catch up with smaller ones, forming yet larger clusters and so on. The clusters do not always contain the same fibers, as fibers constantly enter and leave the clusters as they sediment. These interesting dynamics can be quantified using the following procedure, summarized in Fig. 8. A density function of the vertical positions of the fibers at a given time can be computed by dividing the simulation box along the vertical axis into equally sized bins and counting the number of fibers in each bin: such a function presents peaks at the locations of the clusters and valleys between them, and the heights and breadths of the peaks are good indicators of the cluster sizes. More precisely we define a cluster as a region where the local particle volume fraction peaks above a threshold of 1.3 times its average value $n l^3$, and the limits of

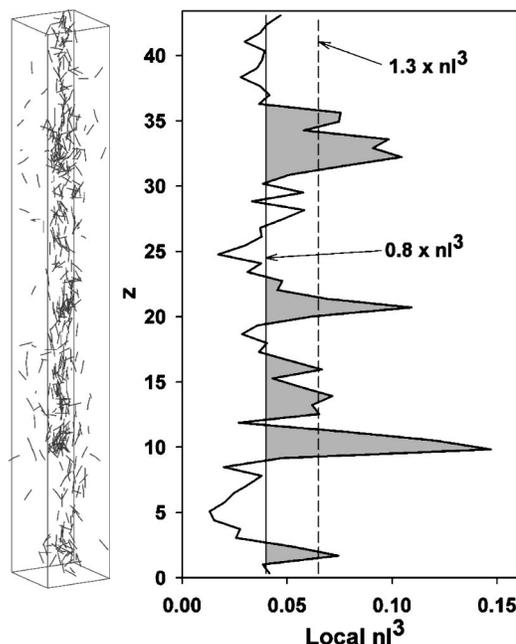


FIG. 8. Procedure for the systematic determination of cluster positions and sizes. A density function of the vertical distribution of fibers is obtained by dividing the box into bins and counting the fibers in each bin. Clusters are defined as regions where the local particle volume fraction peaks above 1.3 times its average value $n l^3$ and always remains above 0.8 times the average $n l^3$. In the example shown, the procedure finds five clusters of sizes 52, 47, 30, 17, and 12 fibers.

the cluster are taken to be the locations on each side where it decreases below 0.8 times the average $n l^3$. Integrating the particle density function over the extent of the cluster then provides an estimate of the number of fibers inside the cluster. Quite obviously the values of the two thresholds could be chosen differently: the values suggested here are such that the clusters defined by this procedure correspond to what one would define to be clusters by simply looking at the suspension.

The procedure is easily programmed on a computer and can be repeated at different times, allowing one to track the cluster positions and to perform statistics on their numbers and sizes. Figure 9 shows the evolution in time of the numbers of large (>30 fibers) and small (<30 fibers) clusters in the simulation of Fig. 6. To remove the high-frequency noise present in the original data, the curves have been smoothed by convolution with a top-hat function of width $\Delta t=8$ and unit area, which explains the noninteger values for the numbers of clusters. The early values (up to $t=15$ approximately) are not very significant as the streamer is not properly formed yet, and are therefore omitted on the graph. Both curves initially present slow oscillations at nearly the same frequency, where the peaks in the number of small clusters correlate with the valleys in the number of large clusters and vice versa. This observation has an easy interpretation and confirms the dynamics alluded to above: up to approximately $t=100$, there is a periodic build-up mechanism by which small clusters merge into larger clusters, which end up breaking up into small clusters again and so forth. Progressively the number of small clusters decreases to the benefit of larger

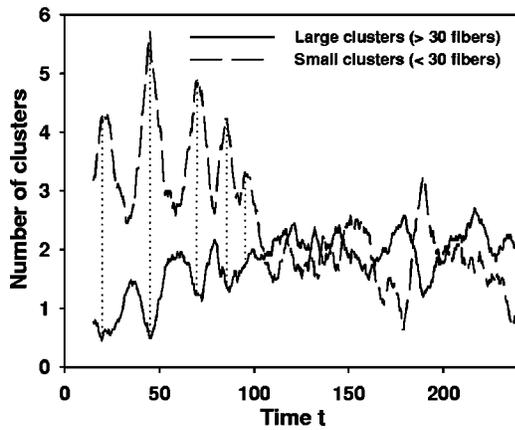


FIG. 9. Time evolution of the numbers of clusters of a given size in a streamer. The results are for the same simulation as in Fig. 6: 512 fibers, $A=11$, $nl^3=0.05$, and $d_x:d_y:d_z=1:1:8$. The estimates of the numbers of clusters were obtained using the procedure described in Fig. 8. To remove the high-frequency noise, the curves have been smoothed by convolution with a unit-area top-hat function of width $\Delta t=8$.

clusters, until a steady state is reached where a combination of clusters of different sizes cohabit (two small and two large clusters on average). Note that some oscillations can still be observed after the steady state is reached, but typically at a higher frequency and weaker amplitude.

The periodic boundary conditions in the vertical direction can be expected to have a significant influence on the observed oscillations. In fact it can be seen that the period of the oscillations is of the same order as the time that it takes for a fiber to sediment one box height at the mean sedimentation speed (see Sec. V D for a discussion on the sedimentation rate). The decrease of the period that can be observed in Fig. 9 is therefore a direct consequence of the increase in the sedimentation velocity as a result of the streamer formation. This periodicity can be understood in the following way: big clusters attaining the bottom of the simulation box reappear at the top and therefore catch up with the smaller clusters that they left behind during the previous period of oscillation. It is unlikely that such oscillations would be observed in an experiment, where a segregation would occur between the different cluster sizes (with the large clusters reaching the bottom of the vessel first). However, the fact that very big clusters break up and that the cluster size reaches a steady state in Fig. 9 bears some significance: it suggests that there is a maximum size beyond which cluster growth becomes unfavorable and clusters either break up or abandon fibers in their wakes.

Since the spatial wavelength in the vertical direction and the cluster size distribution (in terms of number of fibers) reach steady-state values, it is interesting to investigate the densification of clusters through the pair distribution function, which specifies the fraction of pairs of fibers having a center-to-center separation at a given distance. The pair distribution function in the horizontal plane, which is the most interesting, is shown in Fig. 10. Each curve corresponds to the time average of the function over a different time interval. Note that the functions are only plotted over half a box width $d_x/2$, whereas the maximum horizontal distance be-

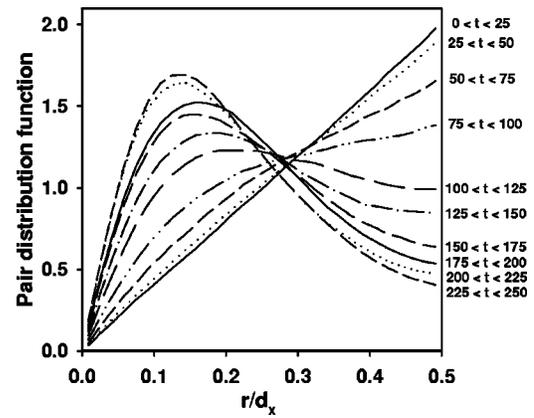


FIG. 10. Pair distribution function in the (x,y) plane for the simulation of Fig. 6, averaged over various time intervals. The pair distribution function specifies the fraction of fibers with a center-to-center separation at a given distance.

tween two fibers is $\sqrt{2}d_x/2$ (owing to the periodicity in x and y). Starting from a linear profile, corresponding to a uniform distribution of fibers, the function evolves until a maximum appears at a finite distance, which represents the most probable separation between two fibers in the horizontal plane. This peak in the distribution function slowly increases and migrates towards lower values, while the amplitude of the tail decreases. The interpretation is rather simple: the streamer progressively becomes thinner and denser, capturing more and more fibers into its core. This process appears to be extremely slow, and continues even after the cluster distribution reaches its steady state in Fig. 9. We can expect this densification to eventually slow down and stop, as excluded volume effects become more and more important as the fibers get more tightly packed. Limitations in computation time, however, did not allow us to run our simulations until a true steady state was observed.

2. Horizontal structure and wavenumber selection

Obtaining more than one concentrated streamer in the lateral direction is a challenging task. In all the previous computational studies only one streamer was obtained, and it was not clear whether this was an artefact of the periodic boundary conditions or a consequence of the limited size of the simulated systems. Neither the horizontal extent of the streamers nor the wavelength of the instability have been studied systematically in the published experiments,^{14,15} and they are not predicted either by the linear stability analysis,¹⁶ so that no information is available on the requirements for the simulation box dimensions or number of fibers.

By increasing the horizontal dimensions of our simulation box we were able to obtain more than one streamer in the lateral direction, as shown in Fig. 11 for instance. For the system sizes that we considered, the only way to achieve this was to drastically reduce the height of the box as well as its width in one horizontal direction: the aspect ratio of the box in Fig. 11 is $d_x:d_y:d_z=10:1:2$. This choice of box dimensions is likely to have an influence on the structure of the suspension, first because choosing very different values of d_x and d_y introduces an artificial anisotropy in the horizontal

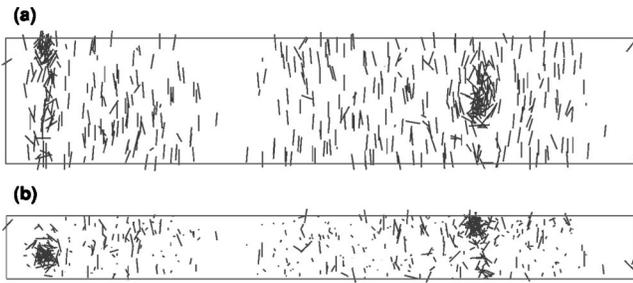


FIG. 11. Distribution of fibers at $t=140$ for a highly elongated box in the horizontal direction (box aspect ratio: $d_x:d_y:d_z=10:1:2$); the top plot (a) shows a side view of the suspension while the bottom plot (b) provides a vertical view. The simulation is for 512 fibers of aspect ratio $A=11$ with an average effective volume fraction of $nl^3=0.05$. Two distinct streamers separated by clarified regions can be observed.

plane, and second because the small values of d_y and d_z make the suspension essentially homogeneous in both the y and z directions. The precise consequences are difficult to assess until larger simulations are performed, so one should be cautious when trying to extrapolate the results presented here to full-scale suspensions.

The structure of the suspension in Fig. 11 is quite interesting, as it exhibits a broad region between the core of the streamers and the clarified fluid where most of the fibers are very well oriented with gravity, and simply sediment vertically with little horizontal motion; a more detailed observation shows that fibers in this region slowly migrate towards the streamers. The vertically averaged fluid velocity field for the simulation of Fig. 11 is shown in Fig. 12 and provides an explanation. As expected the vertical fluid velocity peaks in the core of the streamers, and presents a backflow in the clarified regions; between those two, a large region of shear exists where the velocity varies almost linearly. The alignment of the fibers in the direction of the velocity is then a simple consequence of Eq. (3) for the slender-body orientation dynamics, which predicts a stable equilibrium with no rotational velocity in linear shear when the fiber is aligned in the direction of shear. The reorientation of fiber suspensions in a simple shear flow is a well-known phenomenon.^{43–46} Unless subjected to lubrication or contact forces, the fibers located in this region can undergo horizontal motion only as

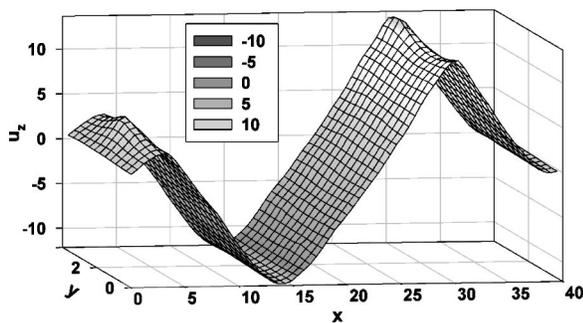


FIG. 12. Vertically averaged vertical velocity field for the simulation of Fig. 11. The positive values correspond to downward velocities. The two peaks are located at the core of the concentrated streamers, and a backflow occurs in the clarified regions. Large regions of almost linear shear exist between streamers and clarified fluid.

TABLE III. Influence of the initial distribution on the streamer formation. Simulations were performed in boxes of aspect ratio $d_x:d_y:d_z=10:1:2$, in which the initial random distribution was perturbed by the superposition of two waves of wavenumbers $k=1$ and $k=2$: $c(x,y,z)=\hat{c}_0+\hat{c}_1\exp(2\pi ix/d_x)+\hat{c}_2\exp(4\pi ix/d_x)$. The ratio \hat{c}_1/\hat{c}_2 of the amplitudes of the two waves was varied. The table reports the ratio \hat{u}_1/\hat{u}_2 of the $k=1$ and $k=2$ coefficients of the Fourier transform in the x direction of the vertical component of the initial disturbance velocity field; the number of distinct backflow regions in the initial velocity field; the number of concentrated streamers in the final distribution.

\hat{c}_1/\hat{c}_2	\hat{u}_1/\hat{u}_2	Backflow regions	Streamers
0.20	0.88	2	2
0.66	3.07	1	1
0.87	3.86	1	1
1.29	5.80	1	1
1.59	7.07	1	1

a result of the velocity fluctuations around the mean: this process, termed hydrodynamic dispersion, is diffusive in nature and therefore rather slow.⁴⁷ It is worthwhile here to consider the limitations of this simulation. First the orientation dynamics as described by slender-body theory are inexact as they cannot predict the tumbling motion of an isolated particle in shear flow: the consequences of this approximation are discussed in more detail in Sec. V E. Second and perhaps most importantly, the quasihomogeneity of the suspension in the vertical direction as created by the small box dimension and periodic boundary conditions helps make the configuration of Fig. 11 stable: inhomogeneities in the vertical direction would quite likely perturb this configuration at least during the onset of the instability, and may disrupt these large regions of shear.

The simulation of Fig. 11 was repeated for different random initial distributions: sometimes only one streamer formed, and in very rare cases three streamers were observed. This strong dependence of the structure formation on the initial distribution calls for a more systematic investigation. Simulations were run in which the initial distribution was perturbed artificially by a superposition of two planar waves in the x direction of wavenumbers $k=1$ and $k=2$ (where k is nondimensionalized by the box length d_x): $c(x,y,z)=\hat{c}_0+\hat{c}_1\exp(2\pi ix/d_x)+\hat{c}_2\exp(4\pi ix/d_x)$; the ratio \hat{c}_1/\hat{c}_2 of the amplitudes of the two waves was varied. The initial velocity fields were computed and compared to the final fiber distributions, and the results are summarized in Table III. Several observations can be made. Depending on the ratio of the two initial wave amplitudes, either one or two streamers form; only for very small values of \hat{c}_1/\hat{c}_2 do two concentrated streamers develop, and in all other cases only one streamer is observed. It is enlightening to compare these observations to the initial velocity fields, and more precisely to the ratio \hat{u}_1/\hat{u}_2 of the $k=1$ and $k=2$ coefficients of the Fourier transform in the x direction of the vertical component of the velocity, and to the number of distinct backflow regions in the velocity field. The number of concentrated streamers in the final distribution seems to be correlated to the number of backflow regions, which itself is determined by the ratio \hat{u}_1/\hat{u}_2 . More precisely, backflow regions seem to

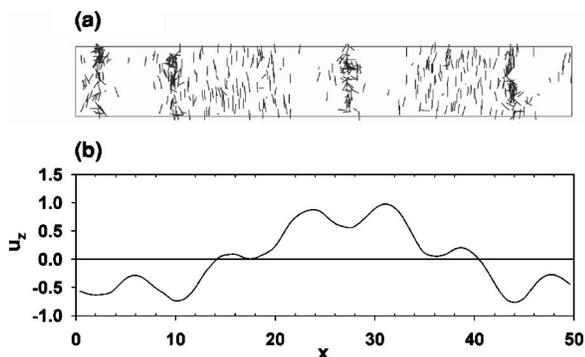


FIG. 13. Final fiber distribution (a) and initial disturbance velocity field (b) for an initial plane wave perturbation at wavenumber $k=6$: $c(x,y,z)=\hat{c}_0 + \hat{c}_6 \exp(12\pi i x/d_x)$. Only four concentrated streamers can be observed, which are located in the wells of the initial velocity field.

act as barriers that cause the fibers to rotate and migrate towards the denser “velocity wells.” Since, from the periodic fundamental solution Eq. (11), the Fourier coefficients of the velocity scale as $\hat{u}_k \sim \hat{c}_k/k^2$, the flow is typically dominated by the $k=1$ mode and only one streamer forms.

At the other end of the spectrum, it is observed that perturbations in concentration usually decay at high wave numbers. Simulations were run in which a high-frequency perturbation ($k=4$ to 7) was applied to the initial distribution; in most cases fewer streamers were obtained than initial waves. This phenomenon, that we term streamer coalescence, is illustrated in Fig. 13 for the case $k=6$. Starting from six initial waves, only four concentrated streamers are observed in the final distribution [Fig. 13(a)]. The initial vertical velocity field [Fig. 13(b)] somewhat modifies our previous criterion and shows that streamers form in the local minima or “wells” of the velocity field, although they need not be surrounded by independent backflows. Figure 14 shows the discrete Fourier transforms in the x direction of the initial concentration perturbation and of the initial velocity field. Although the $k=6$ mode dominates the initial concentration, it only creates a very small disturbance in the velocity spectrum, which is swamped by the low-wavenumber modes (mostly $k=1$) arising from random concentration fluctuations. Therefore high-wavenumber perturbations cannot survive as even tiny low-wavenumber fluctuations will typically dominate the velocity field.

The presence in most simulations of this strong $k=1$ mode in the velocity field is reminiscent of the large recirculation vortices observed in experiments on sphere suspensions. Guazzelli⁴⁸ studied the sedimentation of suspensions of spheres using particle image velocimetry, and observed that during the initial moments of her experiments the velocity field is dominated by large vortices of the size of the container; after a while these vortices decay and leave place to smaller vortices whose size is found to be independent of the container dimensions. As explained by Hinch,⁴⁹ these large convection currents in the initial times are due to the difference in weight between the two sides of the suspension; these currents have the effect of homogenizing the two sides, after which they decay. The strong $k=1$ mode observed in our simulations has the same origin, but is not allowed to

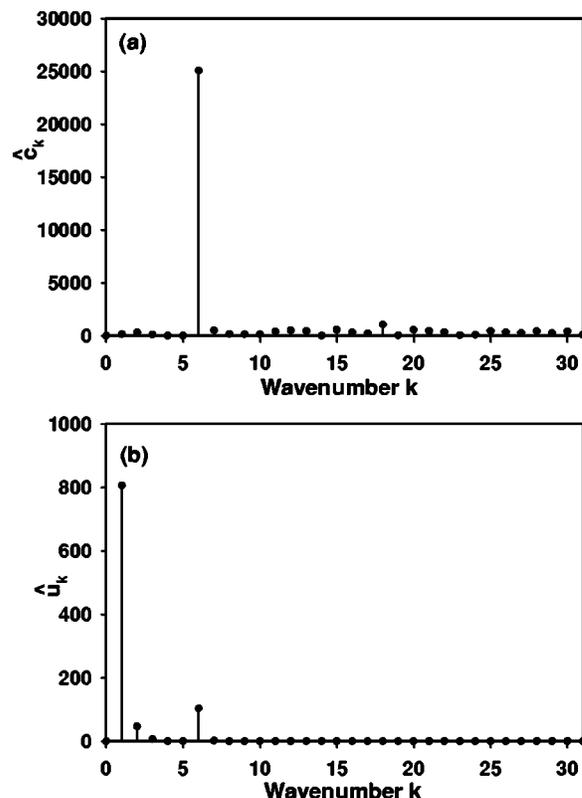


FIG. 14. Magnitude of the coefficients of the Fourier transform in the x direction of (a) the initial concentration field and (b) the vertical component of the initial disturbance field for the simulation of Fig. 13. A strong peak at $k=6$ can be observed in the initial concentration field; the velocity field, however, is dominated by the random low-wavenumber fluctuations ($k=1$).

decay because of the periodic boundary conditions in the vertical direction, which effectively prevent the horizontal convection currents that would otherwise occur in a bounded system. Fiber suspensions are more complex due to the strong coupling between the concentration and velocity fields, and even though a similar recirculation as in sphere suspensions may occur in a container with a bottom wall and may lead to some type of homogenization, the instability on the contrary acts against the homogenization of the suspension and tends to accentuate the fluctuations. The combination of these two effects may lead to a qualitatively different behavior than is observed in sphere suspensions, and more work is needed to determine whether the physical picture described by Guazzelli⁴⁸ by which the initial recirculation vortex diminishes and is replaced by smaller less intense vortices still holds; understanding this process may be the key to explaining the wavenumber selection of the instability.

C. Orientation dynamics

Both experiments¹⁵ and previous simulations¹⁸ showed a reorientation of the fibers in the direction of gravity, with occasional flippings. This is confirmed in our results and is illustrated in Fig. 15, which shows the evolution of the projected angle of a typical sedimenting fiber (the projected angle being defined as the angle between the major axis of the fiber and the horizontal plane). After a transient phase

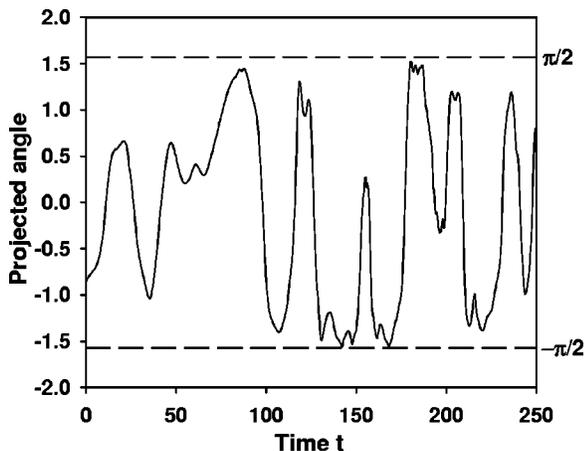


FIG. 15. Time evolution of the projected angle for a typical fiber (angle between the major axis of the fiber and the horizontal). After a transient phase, the fiber spends most of its time almost aligned with gravity (projected angle of $\pm\pi/2$) with occasional flippings.

lasting until $t=100$ approximately, the fiber ends up spending most of its time aligned in the vertical direction. The trajectory of Fig. 15 is very similar to those observed in experiments.¹⁵

The evolution of the orientation statistics can be studied by looking at the average square orientation in the direction of gravity $\langle p_3(t)p_3(t) \rangle$. Its time evolution is presented in Fig. 16 for three simulations with the same numbers of fibers, aspect ratios, and volume fractions but different periodic cell dimensions. In all three simulations $\langle p_3(t)p_3(t) \rangle$ starts from a value close to 1/3 corresponding to an initial distribution with random orientations, and increases progressively as the fibers begin to align. The three curves eventually reach a steady state, the value of which differs in each case and gives an indication of the average orientation. A lower steady state value is obtained with high-aspect ratio boxes, corresponding to fewer fibers being aligned with gravity: this had already been observed in smaller simulations,¹⁸ and is corroborated by the orientation distribu-

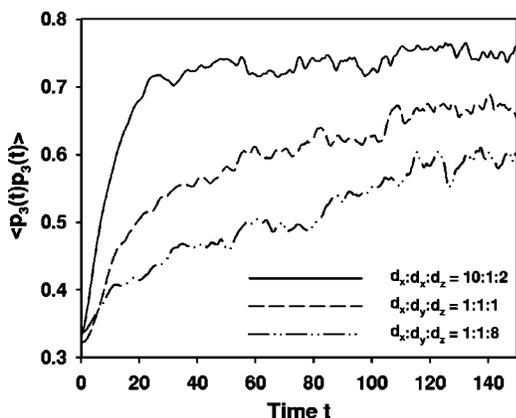


FIG. 16. Time evolution of the average square orientation in the direction of gravity. All three curves are for 512 fibers, $A=11$ and $n^3=0.05$, and different box dimensions. Note how the time to steady state and the value of the steady state depend on the periodic cell aspect ratio.

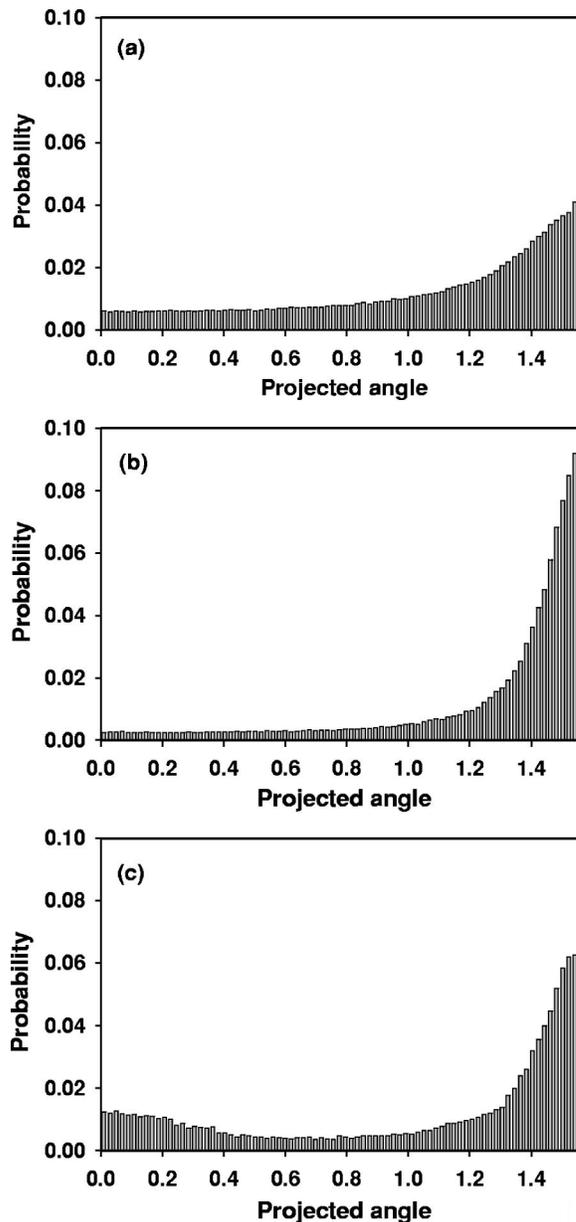


FIG. 17. Comparison of the orientation distributions of two different simulations with the experiments of Herzhaft and Guazzelli (Ref. 15). The two upper plots are for simulations of 512 fibers of aspect ratio $A=11$, with an average effective volume fraction of $n^3=0.05$ and different simulation box aspect ratios: (a) $d_x:d_y:d_z=1:1:8$ (simulation of Fig. 6), (b) $d_x:d_y:d_z=10:1:2$ (simulation of Fig. 11). The bottom plot (c) shows the experimental distribution, obtained for the same particle volume fraction and fiber aspect ratio.

tions presented below. The time to steady state also depends on the simulation box dimensions: it is much quicker in general for low-aspect ratio boxes.

Figure 17 compares the orientation distributions of the two simulations of Figs. 6 and 11 to the experimental results of Herzhaft and Guazzelli.¹⁵ In all three cases the distribution presents a clear peak slightly below $\pi/2$, which was expected and indicates that a large number of fibers are almost aligned in the vertical direction. A second local maximum near zero can also be observed in the experimental data,

showing that the horizontal direction is also a preferred fiber orientation, but this maximum is not reproduced by either of the two simulations.

The breadth and height of the main peak differs greatly between the two simulations. In the first case (high-aspect ratio periodic cell), the peak is rather broad and weak, and the tail of the distribution is quite thick, indicating that quite a large number of fibers are not aligned; in the case of the second simulation (wide and thin periodic cell), the peak is quite higher and narrower and the tail is finer, and the overall profile of the distribution compares somewhat more favorably with the experimental data. This is consistent with the findings of Fig. 16, where we observed two different values for the steady-state average square orientations. A simple inspection of the suspensions of Figs. 6 and 11 provides an explanation. In the case of the high simulation box, only the core of the streamers is really captured, where most of the fibers are entangled inside clusters and are not free to align. Figure 11, however, demonstrates that in a wider box the streamers are surrounded by large regions of shear where the fibers are isolated and can align freely with gravity. In Fig. 17(b) the height of the peak is slightly overestimated: this may be a consequence of the very short height of the box and of the periodic boundary conditions, as in a real system inhomogeneities in the vertical direction would perturb to some extent the shear regions between the streamers. Still these results suggest that these large regions should not be ignored if we want to obtain accurate orientation statistics, hence the necessity to use wide enough boxes to capture the full extent of the horizontal structures.

D. Sedimentation rate and velocity statistics

The formation of dense clusters has a strong influence on the average sedimentation speed of the suspension, which is enhanced in the dilute and semidilute regimes. This phenomenon is really a consequence of the concentration instability, as the velocity is generally hindered in dilute suspensions such as homogeneous suspensions of spheres.¹ Note that in the concentrated regime suspensions of fibers also exhibit velocity hindering,⁵⁰ presumably because strong entanglements present in concentrated suspensions prevent the instability from developing.

In their small simulations, Butler and Shaqfeh¹⁸ observed that the average sedimentation rate of the fibers slowly increases as the instability develops and eventually reaches a plateau, the value of which depends on the aspect ratio of the periodic cell. For the larger systems considered in this study, the time to steady state was generally very long (typically more than 200 time units), so that most of the simulations were stopped before reaching a steady-state velocity (Fig. 18). This is easily understood in the light of the earlier discussion on the pair probability functions: in the Stokes flow regime the sedimentation velocity is a function of the spatial configuration of the suspension only, and as was observed in Sec. V B 1 the clusters keep getting denser and denser even after the cluster distribution and orientation statistics have reached their steady state. However, we also argued in Sec. V B 1 that the densification should eventually

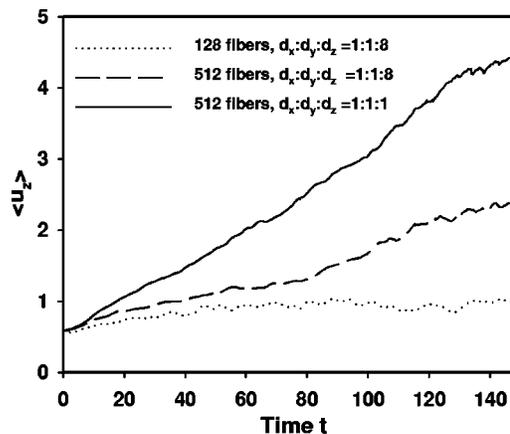


FIG. 18. Time evolution of the average sedimentation rate for simulations of 128 and 512 fibers. In each case $A=11$ and $nI^3=0.05$. While a steady state is reached rapidly with 128 fibers, the sedimentation rate keeps increasing with 512 fibers, and its value depends on the periodic cell aspect ratio.

slow down due to excluded volume effects, and it is likely that a steady sedimentation rate would be reached if the simulations were run long enough.

Figure 18 also shows a strong dependence of the sedimentation rate on the dimensions of the periodic cell: lower sedimentation rates are obtained in high-aspect ratio boxes, in better agreement with the experimental values. Figure 19 compares the velocity distributions for such a simulation to

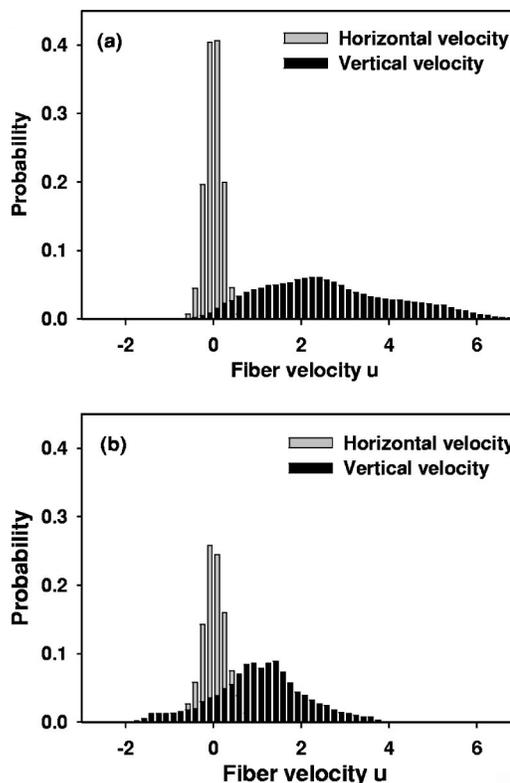


FIG. 19. Comparison of the velocity distributions with the experimental data of Herzhaft and Guazzelli (Ref. 15). The upper plot (a) corresponds to the simulation of Fig. 6 (512 fibers, $A=11$, $nI^3=0.05$, $d_x:d_y:d_z=1:1:8$). The bottom plot (b) shows the experimental distribution, obtained for the same particle volume fraction and fiber aspect ratio.

the experimental results of Herzhaft and Guazzelli:¹⁵ while the horizontal velocities are captured reasonably well, the vertical velocity distribution predicted by our simulation is slightly shifted to the right and somewhat broader than in the experiments; both the mean and the variance of the vertical velocity are overestimated, and this generally gets worse as the simulation box gets wider. Butler and Shaqfeh¹⁸ encountered the same difficulty and came up with the following heuristic rule: for the system sizes they considered, they observed that setting the aspect ratio of the periodic cell to $d_z/d_x=N/16$, where N is the number of fibers, enabled them to obtain the right sedimentation rate and fairly good velocity distributions. Their rule, however, is not easily applied in the case of very large systems, as it would require extremely high and thin boxes, and preclude the capture of more than the core of a single streamer. If we want to resolve the large shear regions between the streamers, which as we established are important for obtaining correct orientation statistics, yet larger systems may be needed, in which both the lateral and vertical box dimensions are large compared to the length of a fiber. Using a bottom wall and nonperiodic boundary conditions in the vertical direction may also provide a solution.

E. Validation of the slender-body approximation

Earlier mention was made of the limitations of the slender-body approximation. The significant feature of the model is that slender bodies in a linear shear flow simply align in the direction of the flow, while exact Stokes flow solutions suggest that an ellipsoid⁵¹ and in fact any body of revolution⁵² should undergo a periodic tumbling motion known as Jeffery's orbits. These poor orientation dynamics may influence our results in two ways: by modifying the dynamics inside the large shear regions observed around the streamers in Sec. V B 2, which in turn may have an impact on the orientation statistics presented in Sec. V C. In particular, we may expect the migration towards the streamers to be enhanced in the presence of tumbling.

To assess more precisely the consequences of the slender-body dynamics, we implemented a slightly different model in which Eq. (3) for the orientation dynamics was replaced by the following:

$$\begin{aligned} \dot{\mathbf{p}}_\alpha = & \frac{3}{4}(\mathbf{I} - \mathbf{p}_\alpha \otimes \mathbf{p}_\alpha) \int_{-1}^1 (1-s^2) \left(\frac{A^2-1}{A^2+1} \mathbf{E}(s) + \boldsymbol{\Omega}(s) \right) \mathbf{p}_\alpha ds \\ & - \frac{3 \ln(2A)}{8\pi} \mathbf{p}_\alpha \times \mathbf{T}_\alpha, \end{aligned} \quad (18)$$

where $\mathbf{E}(s) = \{\nabla \mathbf{u}'(s) + [\nabla \mathbf{u}'(s)]^T\}/2$ and $\boldsymbol{\Omega}(s) = \{\nabla \mathbf{u}'(s) - [\nabla \mathbf{u}'(s)]^T\}/2$ are, respectively, the rate of strain and rate of rotation tensors of the disturbance velocity field, evaluated at position $\mathbf{x}_\alpha + s\mathbf{p}_\alpha$ along the axis of the fiber. Equation (18) is a simplified version of the exact equation for the orientation dynamics of a spheroid,⁵³ in which we neglected an additional term involving the Laplacian of the rate of strain on the basis that its coefficient $1/8(A^2-1)$ is very small for high-aspect ratio particles. Letting $A \rightarrow \infty$ in Eq. (18) and performing an integration by parts allows one to recover Eq. (3) for the slender-body dynamics. Equation (18) can still be

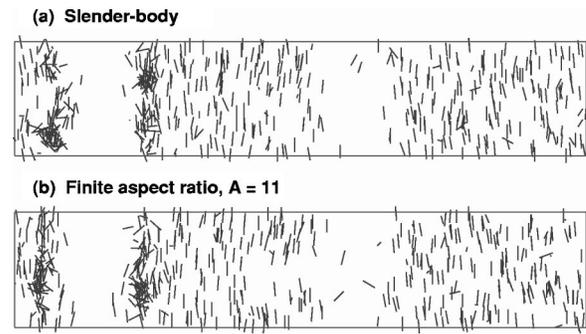


FIG. 20. Comparison of the final fiber distributions obtained using the original slender-body formulation of Eq. (3) (a), and the finite aspect ratio formulation of Eq. (18) (b). The simulation is for 512 fibers, $A=11$, $n\ell^3=0.05$, and $d_x:d_y:d_z=10:1:2$. The two simulations agree qualitatively, which justifies the use of the slender-body approximation.

implemented using SPME: indeed taking the gradient of the velocity in Eq. (14) allows one to derive similar Ewald summation formulas for the components of the rate of strain and rate of rotation tensors, which are then computed following the same algorithm as for the disturbance velocity (Sec. III B). While analytical expressions were used for the real sums, the Fourier sum of the gradient was computed numerically on the Fourier space grid using a fourth-order central finite difference scheme.

For the aspect ratios considered in this study, the modifications entailed by the use of Eq. (18) were minor in general. This is illustrated in Fig. 20, which compares the final distributions obtained with the two different formulations. In particular, the streamer formation is not affected, and neither are the shear regions surrounding the streamers where the alignment of the fibers is sensibly the same. Figure 21 compares for both simulations the time evolution of the average square orientation in the direction of gravity, which quantifies the degree of alignment of the suspension: the differences observed are within the statistical fluctuations of the data. This can be interpreted as follows: even though slender bodies in a simple shear flow align with the flow, the disturbance field between the streamers is not exactly linear and

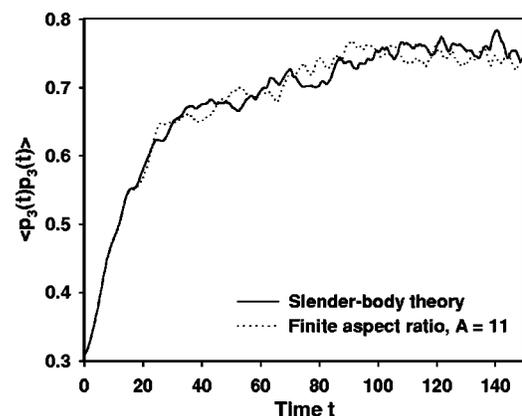


FIG. 21. Time evolution of the average square orientation in the direction of gravity for the two simulations of Fig. 20. Both models [slender-body formulation Eq. (3), and finite aspect ratio formulation Eq. (18)] give very similar results, confirming the validity of the slender-body approximation.

fluctuations in this disturbance field can still cause tumbling (and indeed such tumbling or “flipping” events can be observed in Fig. 15). Furthermore the correction entailed by Eq. (18) is quite weak: the coefficient $(A^2 - 1)/(A^2 + 1)$ only changes from 1 to 0.98 when the aspect ratio is varied from infinity to 11. This suggests that the flipping induced by the velocity fluctuations dominates the tumbling due to the finite thickness of the particles, and under the assumption that the Laplacian term that we neglected to obtain Eq. (18) is indeed small, these results confirm the validity of the slender-body approximation for simulations of high-aspect ratio particles.

VI. CONCLUSION

We have implemented a fast algorithm, called smooth particle-mesh Ewald, which allows one to compute efficiently the hydrodynamic interactions between small particles suspended in a fluid. SPME is comparable to the accelerated Stokesian dynamics method of Sierou and Brady²⁷ for spherical particles, but is characterized by different force assignment and interpolation schemes using Cardinal B -splines, which are particularly well suited for complex exponentials and provide stability at high orders of interpolation. The method, which allows one to efficiently compute the disturbance field induced by periodic distributions of point forces, has been applied to infinite suspensions of slender bodies (line distributions of point forces) in a gravity field, but could easily be applied to different particle shapes (disks or platelike particles for instance, or arbitrary solid surfaces in a boundary element simulation), as well as different flow conditions (particulate flows in an imposed velocity field, pressure-driven flows, etc.). The algorithm can also be easily parallelized, as it is based on a summation of close particle pair interactions for which a domain decomposition is suitable, and on the fast Fourier transform algorithm for which parallel implementations are already available.

When applied to the sedimentation of rigid fibers at zero Reynolds number, our fast algorithm has allowed us to simulate large systems of up to 512 fibers on a single processor, while maintaining a good level of approximation in our simulation method where the force moments on each fiber are retained up to the stresslet term. In particular, the concentration instability observed in sedimenting suspensions of nonspherical particles was captured very convincingly, and our simulations of larger systems have provided valuable information on the microstructure of this type of suspensions.

Using very elongated boxes in the direction of gravity, we were able to study the dynamics inside the core of the streamers that form as the fibers sediment. We observed that the streamers are in fact composed of distinct clusters of various sizes and velocities, which can either merge or break up. As the instability begins to develop, our simulations exhibited a mechanism by which small clusters aggregate into larger clusters, which can again break up. After a transient phase, a steady state is observed in which a combination of clusters of different sizes cohabit on average. The break-up phenomenon, along with the presence of a steady state, suggests that clusters beyond a certain size become unstable, and either break up or abandon fibers in their wakes.

More interestingly, our simulations managed to capture several concentrated streamers in the horizontal direction. This was made possible by drastically increasing one of the horizontal dimensions of our simulation box. These simulations in wide boxes highlighted the presence of broad shear regions between the core of the streamers and the clarified fluid, where most of the fibers sediment almost vertically with little horizontal motion. These regions of shear had not been captured in previous simulations owing to the insufficient size of the periodic cells, and seem to account for some discrepancies that had been observed in the orientation distributions. While the exact process leading to the wavenumber selection is still poorly understood, our simulations provided partial answers: the formation of concentrated streamers seems to be linked to the presence of velocity wells in the initial disturbance field, created by concentration fluctuations in the distribution of fibers. A simple argument shows that for a random suspension the initial disturbance field is dominated by perturbations of long wavelengths ($k=1$ mode). While experiments on the sedimentation of spheres suggest that this initial recirculation cell should decay, this is not observed in our simulations owing to the periodic boundary conditions in the vertical direction.

Obtaining accurate sedimentation rates and velocity distributions remains a very difficult task, as these two quantities depend strongly on the aspect ratio of the periodic cell. Using high-aspect ratio boxes in general gives acceptable results, but the rule defined previously by Butler and Shaqfeh¹⁸ on the box aspect ratio does not allow to capture more than the core of the streamers. In order to capture correct velocity statistics while resolving the horizontal structure of the suspension, we speculate that yet larger systems are required.

ACKNOWLEDGMENTS

The authors wish to thank Jason E. Butler and Élisabeth Guazzelli for some fruitful discussions. D.S. was supported by a Stanford Engineering Fellowship.

APPENDIX: CARDINAL B -SPLINE INTERPOLATION

This appendix presents a few basic results on the Cardinal B -splines used for the interpolation in the SPME algorithm. A similar, somewhat more detailed discussion can be found in the work of Essmann *et al.*,³⁰ or in the work of Schoenberg.⁴⁰

A Cardinal spline of order p is a function defined over the set of real numbers \mathbb{R} , of class C^{p-2} (i.e., $p-2$ times continuously differentiable), and such that its restriction to any interval $[m, m+1)$ for any integer m is a polynomial of degree less than or equal to $p-1$. The set of Cardinal splines of order p is denoted S_p and is easily shown to be a vector space. Quite obviously, S_p is stable by integer translation, i.e., if $f(u)$ is in S_p , so is $f(u-m)$ for any integer m . A linear basis for this vector space is provided by the Cardinal B -splines, which are defined as follows. For any real number $u \in \mathbb{R}$, introduce the notations $u_+ = \max(u, 0) = (u + |u|)/2$ and $u_+^p = (u_+)^p$. Define the following function over \mathbb{R} :

$$M_p(u) = \frac{1}{(p-1)!} \sum_{k=0}^p (-1)^k \frac{p!}{k!(p-k)!} (u-k)_+^{p-1}, \quad u \in \mathbb{R}. \quad (\text{A1})$$

It can be shown that $M_p(u)$ is a Cardinal spline, and that the set of its integer translates $M_p(u-m)$, called Cardinal B -splines, forms a basis of S_p . This means, in particular, that any function $f(u)$ in S_p can be written uniquely as

$$f(u) = \sum_{m=-\infty}^{+\infty} c_m M_p(u-m).$$

The function $M_p(u)$ can be shown to have the following interesting properties:

- (1) $M_p(u) > 0$ for $u \in (0, p)$ and $M_p(u) = 0$ otherwise. In particular, this shows that all the p th order Cardinal B -splines have compact support of length $p+1$.
- (2) $M_p(u)$ is symmetric with respect to the center of its support: $M_p(p-u) = M_p(u)$.
- (3) The Cardinal B -splines sum up to 1: $\sum_{m=-\infty}^{+\infty} M_p(u-m) = 1$.
- (4) The p th order B -splines can be obtained from the lower-order B -splines by a recursion formula

$$M_p(u) = \frac{u}{p-1} M_{p-1}(u) + \frac{p-u}{p-1} M_{p-1}(u-1). \quad (\text{A2})$$

In practice Eq. (A2) is used to compute the Cardinal B -splines, rather than the definition Eq. (A1).

- (5) The derivatives of the B -splines can be obtained from lower-order B -splines:

$$\frac{d}{du} M_p(u) = M_{p-1}(u) + M_{p-1}(u-1).$$

A list of the Cardinal B -splines of orders 1–7 can be found in Deserno and Holm.²⁵

Let us now proceed to show how the Cardinal B -splines can be used to interpolate complex exponentials. Consider the complex-valued function $g(u) = \exp(2\pi i k u / K)$, where u is a real number, and k and K are fixed integers. To alleviate the notations we will set $z = \exp(2\pi i k / K)$, so that now $g(u) = z^u$ where z is some fixed complex number. Quite obviously, we have the property that $g(u+1) = z g(u)$. We wish to interpolate the function g by a p th order spline $g_p \in S_p$ which satisfies the same property, $g_p(u+1) = z g_p(u)$. Since $g_p(u)$ is a Cardinal spline, it can be decomposed on the basis of B -splines,

$$g_p(u) = \sum_{m=-\infty}^{+\infty} c_m M_p(u-m), \quad (\text{A3})$$

and also

$$g_p(u+1) = \sum_{m=-\infty}^{+\infty} c_m M_p(u+1-m) = \sum_{m=-\infty}^{+\infty} c_{m+1} M_p(u-m). \quad (\text{A4})$$

Invoking the uniqueness of the decompositions Eqs. (A3) and (A4), we must have $c_{m+1} = z c_m$ for all integers m for the

property $g_p(u+1) = z g_p(u)$ to be satisfied. A simple recursion leads to $c_m = z^m c_0$ for all integers m , and hence

$$g_p(u) = c_0 \Phi_p(u), \quad \text{where} \quad \Phi_p(u) = \sum_{m=-\infty}^{+\infty} z^m M_p(u-m). \quad (\text{A5})$$

If $\Phi_p(0)$ is nonzero, we can choose the coefficient c_0 to be $1/\Phi_p(0)$. Then we have by construction $g_p(0) = 1$ and $g_p(m) = z g_p(m-1)$ for any integer m , so that by simple recursion $g_p(m) = z^m g_p(0) = z^m$. It can be shown that the only case where $\Phi_p(0) = 0$ is when p is odd and $2|k| = K$. In all other cases we have constructed a p th order spline which interpolates the function $g(u)$ at all the integers. Recalling the original definition of z , and defining

$$b(k) = \frac{1}{\Phi_p(0)} = \frac{\exp\left[2\pi i \frac{k}{K}(p-1)\right]}{\sum_{m=0}^{p-2} M_p(m+1) \exp\left(2\pi i \frac{km}{K}\right)},$$

we obtain

$$\exp\left(2\pi i \frac{ku}{K}\right) \approx b(k) \sum_{m=-\infty}^{\infty} M_p(u-m) \exp\left(2\pi i \frac{km}{K}\right), \quad (\text{A6})$$

which is the same as Eq. (15). The error in the approximation Eq. (A6) can be shown to be bounded by $(2|k|/K)^p$.

¹G. K. Batchelor, "Sedimentation in a dilute dispersion of spheres," *J. Fluid Mech.* **52**, 245 (1972).

²E. J. Hinch, "An averaged-equation approach to particle interactions in a fluid suspension," *J. Fluid Mech.* **83**, 695 (1977).

³A. B. Glendinning and W. B. Russel, "A pairwise additive description of sedimentation and diffusion in concentrated suspensions of hard spheres," *J. Colloid Interface Sci.* **89**, 124 (1982).

⁴F. Feuillebois, "Sedimentation in a dispersion with vertical inhomogeneities," *J. Fluid Mech.* **139**, 145 (1984).

⁵P. G. Saffman, "On the settling speed of free and fixed suspensions," *Stud. Appl. Math.* **52**, 115 (1973).

⁶R. E. Caflisch and J. H. C. Luke, "Variance in the sedimentation speed of a suspension," *Phys. Fluids* **28**, 759 (1985).

⁷H. Nicolai and E. Guazzelli, "Effect of the vessel size on the hydrodynamic diffusion of sedimenting spheres," *Phys. Fluids* **7**, 3 (1995).

⁸D. L. Koch and E. S. G. Shaqfeh, "Screening in sedimenting suspensions," *J. Fluid Mech.* **224**, 275 (1991).

⁹M. P. Brenner, "Screening mechanisms in sedimentation," *Phys. Fluids* **11**, 754 (1999).

¹⁰A. J. C. Ladd, "Effect of container walls on the velocity fluctuations of sedimenting spheres," *Phys. Rev. Lett.* **88**, 048301 (2002).

¹¹S.-Y. Tee, P. J. Mucha, L. Cipelletti, S. Manley, M. P. Brenner, P. N. Segre, and D. A. Weitz, "Nonuniversal velocity fluctuations of sedimenting particles," *Phys. Rev. Lett.* **89**, 054501 (2002).

¹²L. Bergougnoux, S. Ghicini, E. Guazzelli, and J. Hinch, "Spreading fronts and fluctuations in sedimentation," *Phys. Fluids* **15**, 1875 (2003).

¹³P. J. Mucha, S.-Y. Tee, D. A. Weitz, B. I. Shraiman, and M. P. Brenner, "A model for velocity fluctuations in sedimentation," *J. Fluid Mech.* **501**, 71 (2004).

¹⁴B. Herzhaft, E. Guazzelli, M. B. Mackaplow, and E. S. G. Shaqfeh, "Experimental investigation of the sedimentation of a dilute fiber suspension," *Phys. Rev. Lett.* **77**, 290 (1996).

¹⁵B. Herzhaft and E. Guazzelli, "Experimental study of the sedimentation of dilute and semidilute suspensions of fibres," *J. Fluid Mech.* **384**, 133 (1999).

¹⁶D. L. Koch and E. S. G. Shaqfeh, "The instability of a dispersion of sedimenting spheroids," *J. Fluid Mech.* **209**, 521 (1989).

- ¹⁷M. B. Mackaplow and E. S. G. Shaqfeh, "A numerical study of the sedimentation of fibre suspensions," *J. Fluid Mech.* **376**, 149 (1998).
- ¹⁸J. E. Butler and E. S. G. Shaqfeh, "Dynamic simulations of the inhomogeneous sedimentation of rigid fibres," *J. Fluid Mech.* **468**, 205 (2002).
- ¹⁹A. S. Sangani and G. Mo, "An $O(N)$ algorithm for Stokes and Laplace interactions of particles," *Phys. Fluids* **8**, 1990 (1996).
- ²⁰L. Greengard and V. Rokhlin, "A fast algorithm for particle simulations," *J. Comput. Phys.* **73**, 325 (1987).
- ²¹A. Z. Zinchenko and R. H. Davis, "An efficient algorithm for hydrodynamical interaction of many deformable drops," *J. Comput. Phys.* **157**, 539 (2000).
- ²²K. Ichiki, "Improvement of the Stokesian dynamics method for systems with a finite number of particles," *J. Fluid Mech.* **452**, 231 (2002).
- ²³A. J. C. Ladd, "Numerical simulations of particulate suspensions via a discretized Boltzmann equation. Part I. Theoretical foundation," *J. Fluid Mech.* **271**, 285 (1994).
- ²⁴R. W. Hockney and J. W. Eastwood, *Computer Simulation Using Particles* (McGraw-Hill, New York, 1981).
- ²⁵M. Deserno and C. Holm, "How to mesh up Ewald sums. I. A theoretical and numerical comparison of various particle mesh routines," *J. Chem. Phys.* **109**, 7678 (1998).
- ²⁶E. K. Guckel, "Large-scale simulations of particulate systems using the PME method," Ph.D. thesis, University of Illinois at Urbana-Champaign, 1999.
- ²⁷A. Sierou and J. F. Brady, "Accelerated Stokesian dynamics simulations," *J. Fluid Mech.* **448**, 115 (2001).
- ²⁸J. F. Brady and G. Bossis, "Stokesian dynamics," *Annu. Rev. Fluid Mech.* **20**, 111 (1988).
- ²⁹A. Sierou and J. F. Brady, "Rheology and microstructure in concentrated noncolloidal suspensions," *J. Rheol.* **46**, 1031 (2002).
- ³⁰U. Essmann, L. Perera, M. L. Berkowitz, T. Darden, H. Lee, and L. G. Pedersen, "A smooth particle mesh Ewald method," *J. Chem. Phys.* **103**, 8577 (1995).
- ³¹G. K. Batchelor, "Slender-body theory for particles of arbitrary cross-section in Stokes flow," *J. Fluid Mech.* **44**, 419 (1970).
- ³²I. L. Claeys and J. F. Brady, "Lubrication singularities of the grand resistance tensor for two arbitrary particles," *PCH, PhysicoChem. Hydrodyn.* **11**, 261 (1989).
- ³³O. G. Harlen, R. R. Sundararajakumar, and D. L. Koch, "Numerical simulations of a sphere settling through a suspension of neutrally buoyant fibres," *J. Fluid Mech.* **388**, 133 (1999).
- ³⁴L. J. Durlofsky and J. F. Brady, "Dynamic simulation of bounded suspensions of hydrodynamically interacting particles," *J. Fluid Mech.* **200**, 39 (1989).
- ³⁵P. R. Nott and J. F. Brady, "Pressure-driven flow of suspensions: Simulation and theory," *J. Fluid Mech.* **275**, 157 (1994).
- ³⁶J. F. Brady, R. J. Phillips, J. C. Lester, and G. Bossis, "Dynamic simulation of hydrodynamically interacting suspensions," *J. Fluid Mech.* **195**, 257 (1988).
- ³⁷H. 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).
- ³⁸I. L. Claeys and J. F. Brady, "Suspensions of prolate spheroids in Stokes flow. Part 1. Dynamics of a finite number of particles in an unbounded fluid," *J. Fluid Mech.* **251**, 411 (1993).
- ³⁹P. P. Ewald, "Die Berechnung optischer und elektrostatischer Gitterpotentiale," *Ann. Phys.* **64**, 253 (1921).
- ⁴⁰I. J. Schoenberg, *Cardinal Spline Interpolation* (SIAM, Philadelphia, PA, 1973).
- ⁴¹Y. Saad and M. H. Schultz, "A generalized minimum residual algorithm for solving nonsymmetric linear systems," *SIAM (Soc. Ind. Appl. Math.) J. Sci. Stat. Comput.* **7**, 856 (1986).
- ⁴²Y. Saad, *Iterative Methods for Sparse Linear Systems* (Princeton, MA, 1996).
- ⁴³C. A. Stover, D. L. Koch, and C. Cohen, "Observations of fibre orientation in simple shear flow of semi-dilute suspensions," *J. Fluid Mech.* **238**, 277 (1992).
- ⁴⁴M. Rahnama, D. L. Koch, and E. S. G. Shaqfeh, "The effect of hydrodynamic interactions on the orientation distribution in a fiber suspension subject to simple shear flow," *Phys. Fluids* **7**, 487 (1995).
- ⁴⁵Y. Yamane, Y. Kaneda, and M. Dio, "Numerical simulation of semi-dilute suspensions of rodlike particle in shear flow," *J. Non-Newtonian Fluid Mech.* **54**, 405 (1994).
- ⁴⁶X. Fan, N. Phan-Thien, and R. Zheng, "A direct simulation of fibre suspensions," *J. Non-Newtonian Fluid Mech.* **74**, 113 (1998).
- ⁴⁷M. Rahnama, D. L. Koch, Y. Iso, and C. Cohen, "Hydrodynamic, translational diffusion in fiber suspensions subject to simple shear flow," *Phys. Fluids A* **5**, 849 (1993).
- ⁴⁸E. Guazzelli, "Evolution of particle-velocity correlations in sedimentation," *Phys. Fluids* **13**, 1537 (2001).
- ⁴⁹E. J. Hinch, "Sedimentation of small particles," in *Disorder and Mixing*, edited by E. Guyon, J.-P. Nadal, and Y. Pomeau (Kluwer Academic, Dordrecht, 1988).
- ⁵⁰M. A. Turney, M. K. Cheung, R. L. Powell, and M. J. McCarthy, "Hindered settling of rod-like particles measured with magnetic resonance imaging," *AIChE J.* **41**, 251 (1995).
- ⁵¹G. B. Jeffery, "The motion of ellipsoidal particles immersed in a viscous fluid," *Proc. R. Soc. London, Ser. A* **102**, 161 (1922).
- ⁵²F. P. Bretherton, "The motion of rigid particles in a shear flow at low Reynolds number," *J. Fluid Mech.* **14**, 284 (1962).
- ⁵³S. Kim and S. P. Karrila, *Microhydrodynamics: Principles and Selected Applications* (Butterworth-Heinemann, Boston, 1991).
- ⁵⁴The implementation that we used in this work was that by V. Frayssé, L. Giraud, S. Gratton, and J. Langou, CERFACS Technical Report No. TR/PA/03/3, 2003; available from <http://www.cerfacs.fr/algos/Softs/GMRES>