

Accelerated Stokesian Dynamics simulations

By ASIMINA SIEROU AND JOHN F. BRADY

Division of Chemistry and Chemical Engineering, California Institute of Technology, Pasadena,
CA 91125, USA

(Received 5 January 2001 and in revised form 26 April 2001)

A new implementation of the conventional Stokesian Dynamics (SD) algorithm, called accelerated Stokesian Dynamics (ASD), is presented. The equations governing the motion of N particles suspended in a viscous fluid at low particle Reynolds number are solved accurately and efficiently, including all hydrodynamic interactions, but with a significantly lower computational cost of $O(N \ln N)$. The main differences from the conventional SD method lie in the calculation of the many-body long-range interactions, where the Ewald-summed wave-space contribution is calculated as a Fourier transform sum and in the iterative inversion of the now sparse resistance matrix. The new method is applied to problems in the rheology of both structured and random suspensions, and accurate results are obtained with much larger numbers of particles. With access to larger N , the high-frequency dynamic viscosities and short-time self-diffusivities of random suspensions for volume fractions above the freezing point are now studied. The ASD method opens up an entire new class of suspension problems that can be investigated, including particles of non-spherical shape and a distribution of sizes, and the method can readily be extended to other low-Reynolds-number-flow problems.

1. Introduction

Numerical simulations of the behaviour of suspensions of particles provide a valuable tool for understanding many complex rheological phenomena. Through simulations both the macroscopic suspension properties and the suspension microstructure have been studied, and insight into structure–property relationships has been obtained (Brady & Bossis 1988; Foss & Brady 2000). Determining the hydrodynamic interactions among particles in Stokes flow (small Reynolds number), however, can be a complicated and computationally expensive task, mainly owing to the long-range nature of the interactions and the presence of strong lubrication effects when particles are in close proximity to each other. The fluid velocity disturbance caused by a particle on which a net external force acts decays as $1/r$, where r is the distance from the particle, and therefore the hydrodynamic interactions cannot be truncated and no simple pairwise-additive approximation can be made. In addition, the presence of lubrication effects makes conventional numerical techniques (such as the boundary-integral technique) expensive computationally when two particles approach each other.

Durlafsky, Brady & Bossis (1987) developed a method that successfully accounts for both the many-body interactions and the near-field lubrication effects by splitting the hydrodynamic interactions into a far-field mobility calculation and a pairwise additive resistance calculation. The main advantage of the method is that a relatively small number of unknowns per particle is sufficient to solve many dynamic simulation problems adequately. The main disadvantage of the method, however, is that it

requires the inversion of a far-field mobility matrix with at least $(11N)^2$ elements (N being the number of particles in the system), with a computational cost that limits the method to N of the order of a hundred.

The method of Durlafsky *et al.* (1987) and its extension to infinite suspensions by Phillips, Brady & Bossis (1988) is known as Stokesian Dynamics (SD) (Brady & Bossis 1988) and has been used successfully over the last decade to give accurate results for many problems where the system size is of relatively little importance. For a variety of problems however, it is desirable to simulate systems containing a much larger number of particles in order to eliminate any system size effects. For example, in the simulation of very dense suspensions, large-scale simulations are often needed to capture the microstructure correctly; in addition, most commonly used Monte Carlo algorithms cannot even generate particle configurations above the freezing point of hard spheres ($\phi > 0.49$) for $N < 500$ – 1000 as the small system size causes very rapid crystallization. Similarly, problems involving fibre suspensions require large-scale simulations to ensure that the simulation box size is significantly larger than the length of each fibre.

In order to address problems where the system size is important and larger systems are therefore necessary, we present a new method that maintains the same level of accuracy as conventional Stokesian Dynamics, but with a computational cost that scales only as $N \ln N$. Our main objective is to avoid both the costly construction of the far-field mobility matrix and its inversion. Before going into the details of our new approach, it is worth making some general observations to see how we might construct a more efficient algorithm.

Iterative schemes can replace $O(N^3)$ inversions with potentially less costly $O(N^2)$ multiplications, and when applied to ‘special’ matrices – sparse, positive-definite, well-conditioned – iterative schemes can result in $O(N)$ operations. As we shall see in the next section, for most applications of hydrodynamic interactions in Stokes flow, the far-field resistance matrix usually appears as a product with a velocity and the knowledge of the far-field hydrodynamic force (and not the elements of the matrix itself) is sufficient for the iterative calculation of the particle velocities. It is therefore feasible to attempt a method that would only calculate $6N$ far-field hydrodynamic force/torques as opposed to the full far-field mobility matrix. In addition, with this approach, only the near-field contribution to the hydrodynamic interactions remains in a matrix-like form; a matrix that is now sparse and can be manipulated easily in $O(N)$ operations. This approach leads again to a system of linear equations of the form $\mathbf{A} \cdot \mathbf{x} = \mathbf{b}$, where now \mathbf{x} is the vector representing the translational and rotational velocities of all particles, \mathbf{A} is a sparse matrix representing the near-field part of the resistance matrix with only $O(N)$ non-zero elements, and \mathbf{b} represents the shearing and non-hydrodynamic force/torques on the particles and, in addition, includes all the far-field physics. We use a particle-mesh approach (Hockney & Eastwood 1988) in combination with the Ewald summation technique to generate the far-field part of \mathbf{b} in $O(N \ln N)$, an approach based on newly developed particle-mesh-Ewald algorithms (Darden, York & Pedersen 1993). Lubrication interactions are treated as in conventional SD, taking special care to store the now sparse resistance matrix in a computationally efficient manner. The linear system is solved iteratively and, with proper preconditioning, very few iterations are necessary to achieve convergence.

We should also mention here the existence of other $O(N)$ algorithms, developed by Ladd (1994*a, b*) and by Sangani & Mo (1996), that attempt to address similar problems. Ladd’s method is based on the lattice-Boltzmann technique – the fluid continuum is replaced by a lattice-Boltzmann gas, while the behaviour of the rigid

particles is simulated with the use of suitable rules for the momentum and position exchange between the lattice-gas particles. Ladd (1994*a, b*) also used continuum Stokes flow lubrication results to reduce the number of lattice gas points between near touching particles. Sangani & Mo's algorithm follows a more traditional approach by calculating the full resistance matrix through a fast multipole summation technique and inverting the resulting matrix iteratively. The fast multipole method is a widely used approach to performing fast summations in a number of fields. This method is in principle $O(N)$, although the iterative solution technique employed by Sangani & Mo appears to perform poorly, at least in their published results. The calculation of the far-field hydrodynamic interactions presented in the following sections was inspired by the work of J. J. L. Higdon and coworkers (personal communication 2000; Guckel 1999). Higdon applied a particle-mesh-Ewald-sum (PME) technique (Darden *et al.* 1993) for the calculation of the far-field hydrodynamic force/torques; the problem of solving the resulting system of equations efficiently, however, was never fully addressed.

The basic method is presented in §2 where we describe in detail both the calculation of the far-field interactions and the iterative scheme for the solution of the resulting linear system. In §3, we test the accuracy of our method by performing calculations for simple cubic arrays—cases that have been studied extensively in the past and for which both analytical and simulation results are available for comparison. In §4, we calculate properties of random suspensions and document results where the system size is indeed important, e.g. for very concentrated suspensions. We conclude in §5 with remarks on how this method can be extended to other problems in Stokes flow and the analogous problems in electrostatics and elasticity theory.

2. Method

Before proceeding with the presentation of the new method, we give a brief overview of the conventional Stokesian Dynamics technique. For N rigid particles suspended in an incompressible Newtonian fluid of viscosity η and density ρ , the motion of the fluid is governed by the Navier–Stokes equations, while the motion of the particles is described by the coupled N -body equation of motion:

$$\mathbf{m} \cdot \frac{d\mathbf{U}_p}{dt} = \mathbf{F}^h + \mathbf{F}^p, \quad (2.1)$$

which simply states that the mass times the acceleration equals the sum of the forces. In (2.1), \mathbf{m} is the generalized mass/moment-of-inertia matrix of dimensions $6N \times 6N$, \mathbf{U}_p is the particle translational/rotational velocity vector of dimension $6N$, and \mathbf{F}^h and \mathbf{F}^p are the hydrodynamic and external force-torque vectors acting on the particles. Although the Stokesian Dynamics method can also be applied to problems where Brownian motion is important, here we only consider non-Brownian systems (infinite Péclet number); extension to include Brownian motion in ASD is possible, however (Banchio & Brady 2001). When the motion on the particle scale is such that the particle Reynolds number is small, the fluid equation of motion becomes linear (Stokes equation) and the hydrodynamic forces and torques acting on the particles in a suspension undergoing a bulk linear flow are given by:

$$\mathbf{F}^h = -\mathbf{R}_{FU} \cdot (\mathbf{U}_p - \mathbf{u}^\infty) + \mathbf{R}_{FE} : \mathbf{E}. \quad (2.2)$$

Here, \mathbf{u}^∞ is the velocity of the bulk linear flow evaluated at the particle centre, \mathbf{E} is the externally imposed rate of strain tensor, and $\mathbf{R}_{FU}(\mathbf{x})$ and $\mathbf{R}_{FE}(\mathbf{x})$ are the configuration-

dependent resistance matrices that give the hydrodynamic force/torques on the particles owing to their motion relative to the fluid and owing to the imposed shear flow, respectively; the vector \mathbf{x} denotes the configuration – position and orientation – of the particles.

The combination of the resistance matrices is denoted the grand resistance matrix:

$$\mathcal{R} = \begin{bmatrix} \mathbf{R}_{FU} & \mathbf{R}_{FE} \\ \mathbf{R}_{SU} & \mathbf{R}_{SE} \end{bmatrix}, \quad (2.3)$$

where \mathbf{R}_{SU} and \mathbf{R}_{SE} are similar to \mathbf{R}_{FU} and \mathbf{R}_{FE} and relate the particle stresslet, \mathbf{S}^h – the symmetric first moment of the force density on a particle – to the velocity and the rate of strain. The inverse of the resistance matrix is known as the mobility matrix \mathbf{M} and gives the particle velocities ($\mathbf{U}_p - \mathbf{u}^\infty, -\mathbf{E}$) in terms of the forces (\mathbf{F}^h and \mathbf{S}^h). Conventional SD exploits the fact that hydrodynamic interactions among particles can be decomposed into long-range mobility interactions and short-range lubrication interactions. The long-range interactions are computed by expanding the force density on the surface of each particle in a series of moments. The zeroth moment is simply the net force acting on a particle (plus, for a spherical particle, a potential dipole), the first moment can be decomposed into the torque and the stresslet, while higher moments are neglected. There is no fundamental reason that higher moments cannot and should not be included, and indeed there are problems where higher moments are significant (Ladd 1990); but, the first two moments, when combined with the near-field lubrication interactions, are the minimum set needed and require the least computational effort. Furthermore, this level of truncation has been shown to give very accurate results for many hydrodynamic problems.

This truncated multipole expansion, in combination with Faxén’s laws, is used to form the grand mobility matrix \mathcal{M}^∞ . Once constructed, the grand mobility matrix is inverted to yield a far-field approximation to the grand resistance matrix. While \mathcal{M}^∞ is pairwise additive, on inversion infinite reflections among all moments and all particles are computed, and thus the far-field resistance matrix is a true many-body interaction. Finally, since the many-body approximation to the resistance matrix still lacks lubrication, which would only be reproduced upon inversion of the mobility matrix if all multipole moments were included, the near-field interactions are introduced into the resistance tensor in a pairwise additive fashion. The exact two-body resistance interactions (Jeffrey & Onishi 1984), \mathcal{R}_{2B} , are added to $(\mathcal{M}^\infty)^{-1}$, but since the far-field two-particle resistance interactions have already been included upon the inversion of \mathcal{M}^∞ special care is needed in order not to count these interactions twice. Thus, the two-body interactions already included in $(\mathcal{M}^\infty)^{-1}$, denoted as \mathcal{R}_{2B}^∞ , are subtracted (Durlinsky *et al.* 1987), and the approximation to the grand resistance matrix becomes:

$$\mathcal{R} = (\mathcal{M}^\infty)^{-1} + \mathcal{R}_{2B} - \mathcal{R}_{2B}^\infty. \quad (2.4)$$

Once the grand resistance matrix is known, from (2.1) and (2.2) the particle velocities can be obtained if the forces are known or vice versa. From the particle velocities new configurations are obtained, the resistance tensors are computed anew and the procedure repeated. This method captures both the near- and far-field physics and has given excellent results for many problems. Unfortunately, the direct solution of equation (2.1) as implemented in conventional Stokesian Dynamics is computationally expensive since it involves the costly $O(N^2)$ calculation of the far-field mobility matrix and its costly $O(N^3)$ inversion. We now present an alternative approach in an attempt

to minimize the computational cost and devise a method with a more favourable scaling.

As in conventional SD, we split the hydrodynamic force into a far-field and a near-field part:

$$\begin{aligned} \mathbf{F}^h &= \mathbf{F}_{ff}^h + \mathbf{F}_{nf}^h \\ &= -\mathbf{R}_{FU,ff} \cdot (\mathbf{U}_p - \mathbf{u}^\infty) + \mathbf{R}_{FE,ff} : \mathbf{E} - \mathbf{R}_{FU,nf} \cdot (\mathbf{U}_p - \mathbf{u}^\infty) + \mathbf{R}_{FE,nf} : \mathbf{E}. \end{aligned} \quad (2.5)$$

The near-field resistance matrix in the above equation corresponds simply to the proper component of $\mathcal{R}_{2B} - \mathcal{R}_{2B}^\infty$ in (2.4), while the far-field resistance matrix is the corresponding part of $(\mathcal{M}^\infty)^{-1}$. Instead of calculating $(\mathcal{M}^\infty)^{-1}$ directly as was done before, however, we now calculate the far-field hydrodynamic force instead. Although this approach does result in some loss of information (compared to the calculation of the full matrix), it allows us to calculate only $6N$ hydrodynamic force/torques, a procedure with a significantly smaller computational cost. In the following subsections, we discuss this calculation in detail, along with a discussion of iterative techniques that further reduce the computational cost of the inversion of the remaining near-field matrix.

2.1. Calculation of the far-field interactions

As was already mentioned, in order to avoid the expensive construction of the far-field resistance matrix we calculate the far-field hydrodynamic force directly—that is, the product of the resistance matrix with a known velocity. In order to clarify the most important features of the new approach, we first use a simple schematic description. In Stokes flow the velocity at any point in the fluid can be expressed in a number of equivalent ways—integral representations of Green functions, multipole expansions of force moments, etc. In the following paragraphs we will present the multipole expansion approach in detail, but for the moment we use a very general functional form:

$$u_i(\mathbf{x}) - u_i^\infty(\mathbf{x}) = \sum_n G_{ij}(\mathbf{x}, \mathbf{x}_n) F_j^n, \quad (2.6)$$

where $u_i(\mathbf{x}) - u_i^\infty(\mathbf{x})$ is the i -component of the velocity disturbance in the fluid at any position \mathbf{x} , F_j^n is the j -component of the hydrodynamic force, torque, stresslet (or even higher moments) on each particle n , and G_{ij} corresponds to the appropriate solution function. To determine the motion of a particle immersed in a flow field given by (2.6), we make use of the Faxén formulae that relate the force (and higher moments) on a particle n to the particle velocity $U_{p,i}^n$ and the fluid velocity at the particle centre, $u_i(\mathbf{x}_n)$:

$$F_i^n = (U_{p,i}^n - u_i^\infty(\mathbf{x}_n)) + H(u_i(\mathbf{x}_n) - u_i^\infty(\mathbf{x}_n)), \quad (2.7)$$

where H represents a known functional operation. (A scalar operator multiplying the particle velocity, e.g. $6\pi\eta a$ for a sphere of radius a in a fluid with viscosity η , can be included in the non-dimensionalization of the force and therefore has been omitted here.) It should be clear that by combining equations (2.6) and (2.7) we can eliminate the fluid velocity $u_i(\mathbf{x}_n)$ and construct a mobility matrix relating the velocity of each particle to the forces on all of the particles. It should also be obvious that the mobility formulation is the most straightforward to calculate, since $U_{p,i}^n$ is only present in equation (2.7); such a mobility formulation must then be inverted to give the resistance formulation.

An alternative approach is not to calculate the mobility matrix and calculate F_i^n directly in an iterative manner. Assuming that all the particle velocities are known

(either as a nested iterative procedure, or more simply as the velocities of the previous time step), an initial value is assumed for the hydrodynamic forces, equation (2.6) is solved for the corresponding fluid velocities, u_i , and then with application of Faxén's laws (2.7) the initial forces are corrected and the process repeated until convergence. The computational cost of such a procedure is dominated by the $O(N^2)$ calculation of $\sum_n G_{ij}(\mathbf{x}, \mathbf{x}_n) F_j^n$ – a sum over the forces of all the other particles. However, the calculation of these sums can be accelerated significantly with the use of recently developed methods, such as the PME technique. The ideas of particle-mesh techniques (Hockney & Eastwood 1988) are used to assign particles to a mesh according to their positions and then fast Fourier transform (FFT) techniques are used to evaluate the wave-space part of the Ewald sum on this mesh. The fast evaluation of the wave-space sum is then used to set the parameters to allow an $O(N)$ evaluation of the real-space sum. PME algorithms have been successfully applied to the calculation of Coulombic interactions (Darden *et al.* 1993, 1997; Essmann *et al.* 1995; Petersen 1995) and recently have been formulated for the case of hydrodynamic interactions by Higdon and coworkers (Guckel 1999); this formulation will be described here in detail.

Our starting point is Hasimoto's solution (1959) of the Stokes equations for the flow past a periodic array of spheres. Hasimoto's solution is exactly equivalent to that given by Beenaker (1986); however, because a regular grid of points is used for the calculation of the Fourier space sum, Hasimoto's approach is preferred. Note also that all of the following analysis refers to periodic systems; the subject of convergence of the resulting infinite sums will not be discussed since it has been resolved in the past (see Brady *et al.* 1988). Following Hasimoto, the fluid velocity in the presence of a periodic array of N suspended particles at positions \mathbf{r}_n , represented as point forces \mathbf{F}^n (with F_j^n , the j -component of the force that the fluid exerts on particle n), in a periodic unit cell of volume V_0 can be expressed as:

$$u_j(\mathbf{x}) = \frac{1}{4\pi\eta} \left(S_j^1 - \frac{\partial^2 S_l^2}{\partial x_l \partial x_j} \right), \quad (2.8)$$

where

$$S_l^2 = -\frac{1}{4\pi^3 V_0} \sum_{\mathbf{k} \neq 0} \frac{e^{-2\pi i(\mathbf{k} \cdot \mathbf{r})}}{k^4} \hat{F}_l^k, \quad (2.9)$$

and

$$S_j^1 = \nabla^2 S_j^2. \quad (2.10)$$

Here, \mathbf{k} is the reciprocal lattice vector (corresponding to the real space vector \mathbf{r}), and \hat{F}_j^k is the j -component of the Fourier transform of the real-space lattice point force \mathbf{F}^n .

The above sum, which simply represents the contribution to the fluid velocity of all the 'particle-point-forces', can be evaluated efficiently with the introduction of Ewald's summation technique. We start with an integral representation for $1/k^{2m}$:

$$\frac{1}{k^{2m}} = \frac{\pi^m}{\Gamma(m)} \int_0^\infty e^{-\pi k^2 \beta} \beta^{m-1} d\beta. \quad (2.11)$$

Then

$$\begin{aligned} \sigma_j^m &= \sum_{\mathbf{k} \neq 0} \frac{e^{-2\pi i(\mathbf{k} \cdot \mathbf{r})}}{k^{2m}} \hat{F}_j^k = \frac{\pi^m}{\Gamma(m)} \sum_{\mathbf{k} \neq 0} \int_0^\infty e^{-\pi k^2 \beta - 2\pi i(\mathbf{k} \cdot \mathbf{r})} \hat{F}_j^k \beta^{m-1} d\beta \\ &= \frac{\pi^m}{\Gamma(m)} \int_0^\infty \beta^{m-1} \left[\sum_{\mathbf{k}} \hat{F}_j^k e^{-\pi k^2 \beta - 2\pi i(\mathbf{k} \cdot \mathbf{r})} - \hat{F}_j^0 \right] d\beta. \end{aligned} \quad (2.12)$$

A splitting parameter α is introduced and the integral in (2.12) is split into two parts, one from 0 to α , and the other from α to ∞ . Ewald's theta transformation formula

$$\sum_k \hat{F}_j^k e^{-\pi k^2 \beta - 2\pi i(\mathbf{k}\cdot\mathbf{r})} = \frac{V_0}{\beta^{3/2}} \sum_n e^{-\pi(\mathbf{r}-\mathbf{r}_n)/\beta} F_j^n \quad (2.13)$$

is then applied to the integral from 0 to α .

The general formula for the evaluation of S_j^1 or S_j^2 is then:

$$\begin{aligned} \sigma_j^m &= \sum_{k \neq 0} \frac{e^{-2\pi i(\mathbf{k}\cdot\mathbf{r})}}{k^{2m}} \hat{F}_j^k \\ &= \frac{\pi^m \alpha^m}{\Gamma(m)} \left[V_0 \alpha^{-3/2} \sum_n \phi_{-m+1/2} \left(\frac{\pi(\mathbf{r}-\mathbf{r}_n)^2}{\alpha} \right) F_j^n - \frac{\hat{F}_j^0}{m} \right] \\ &\quad + \frac{\pi^m \alpha^m}{\Gamma(m)} \left[\sum_{k \neq 0} e^{-2\pi i(\mathbf{k}\cdot\mathbf{r})} \phi_{m-1}(\pi \alpha k^2) \hat{F}_j^k \right], \end{aligned} \quad (2.14)$$

where we have replaced $\beta = \alpha/\xi$ in the first integral and $\beta = \alpha\xi$ in the second. The first sum and the constant term correspond to the real-space sum contribution from the presence of all particles (including the self-term), while the second sum is the wave-space sum excluding the $\mathbf{k} = 0$ term. The function ϕ_ν is the incomplete Γ -function:

$$\phi_\nu(x) = \int_1^\infty \xi^\nu e^{-x\xi} d\xi, \quad (2.15)$$

and can easily be evaluated since it satisfies simple recurrence formulae (see Hasimoto 1959 for more detail).

Using (2.14) and substituting the expressions (2.9) and (2.10) for S_1 and S_2 into (2.8) for the fluid velocity, we obtain:

$$\begin{aligned} u_j(\mathbf{x}) &= \frac{1}{4\pi\eta} \left(\sum_n F_l^n \left[\left(-\frac{\pi}{\alpha^{3/2}} \right) r^2 \phi_{1/2} \delta_{jl} + \alpha^{-1/2} \delta_{jl} + \left(\frac{\pi}{\alpha^{3/2}} \right) x_l x_j \phi_{1/2} \right] \right) \\ &\quad + \frac{1}{4\pi\eta} \left(\frac{1}{\pi V_0} \sum_{k \neq 0} e^{-2\pi i(\mathbf{k}\cdot\mathbf{r})} \left(\frac{\hat{F}_j^k}{k^2} - \frac{\hat{F}_l^k k_l k_j}{k^4} \right) e^{-\pi \alpha k^2} (1 + \pi \alpha k^2) \right). \end{aligned} \quad (2.16)$$

This is the exact result for Stokes flow in a periodic array of point forces. We shall now show how this result can be generalized for the case of a random distribution of particles, to include higher moments (torque, stresslet, etc.) in order to account for the finite size of the particles in different types of flows, and we shall demonstrate an efficient way to calculate the above sums numerically.

The incomplete Γ -functions $\phi_\nu(x)$ tend rapidly to zero as $x \rightarrow \infty$. Therefore as $\alpha \rightarrow 0$,

$$\phi_{-m+1/2}(\pi(\mathbf{r}-\mathbf{r}_n)^2/\alpha) \rightarrow 0$$

and the velocity is represented as a pure wave-space sum. On the other hand, for large values of α the wave-space sum contribution vanishes and the velocity is represented as a real-space sum. The choice of α and the optimum efficiency of balancing the computational work for the evaluation of the two sums is one of the main issues of the algorithm.




FIGURE 1. The transition from a force acting on a particle's centre to a set of forces acting on the grid points γ .

2.2. Wave-space sum contribution

The wave-space sum contribution to the fluid velocity given by

$$WS = \frac{1}{4\pi\eta} \left[\frac{1}{\pi V_0} \sum_{k \neq 0} e^{-2\pi i(\mathbf{k} \cdot \mathbf{r})} \left(\frac{\hat{F}_j^k}{k^2} - \frac{\hat{F}_l^k k_l k_j}{k^4} \right) e^{-\pi \alpha k^2} (1 + \pi \alpha k^2) \right], \quad (2.17)$$

can be evaluated efficiently as the inverse Fourier transform of the expression

$$\left(\frac{\hat{F}_j^k}{k^2} - \frac{\hat{F}_l^k k_l k_j}{k^4} \right) e^{-\pi \alpha k^2} (1 + \pi \alpha k^2). \quad (2.18)$$

In order to do so, some further analysis of the meaning and calculation of \hat{F}_j^k is necessary. Let us consider again the flow past a periodic array of spheres. If the spheres are sufficiently small, then the force acting on the surface of each sphere can be simply assumed to be a point force acting on its centre. The position of sphere n can then be specified either in the real-space (\mathbf{r}) or in the reciprocal-space (\mathbf{k}). For the case of a real-space simple cubic lattice, the reciprocal space lattice is also simple cubic. Then if F_j^n is the j component of the point-force representing particle n at position \mathbf{r}_n , \hat{F}_j^k is simply its Fourier transform (on the same position, since both the real and reciprocal lattices are simple-cubic). For the case of a cubic array of point forces, the above description is all that is needed to evaluate the wave sum using conventional FFT methods. In order to address more realistic cases, e.g. finite-sized particles in random configurations and in different external fields, while still maintaining the 'lattice-point-force' approach, we must implement a more general particle and grid representation.

Let us now consider the case of a random distribution of spheres in a periodic computational domain where the particles are no longer part of a periodic array. Although the positions of the spheres no longer define a lattice, we can still define an artificial rectangular grid over the computational domain (see figure 1); the positions of the particles, however, no longer coincide with the grid-points. It is nevertheless possible to transfer the force acting on the centre of each sphere to a collection of forces acting on the grid-points located in the vicinity of the sphere. This is done by following well-known particle-mesh algorithms (Eastwood & Hockney 1988). It is desirable that the transition from the particle centre to the grid should involve only

a relatively small number of grid points, be a smooth function of the location of the particle and capture the correct long-range physics for distances far away from the force centre. Here, we perform the force assignment by the use of matched Taylor series expansions. For example, for the point force we require:

$$F_j^0 e^{-2\pi i(\mathbf{k}\cdot\mathbf{r}^0)} = \sum_{\text{nodes } \gamma} f_j^\gamma e^{-2\pi i(\mathbf{k}\cdot\mathbf{r}^\gamma)}, \quad (2.19)$$

or, after Taylor series expansion:

$$\begin{aligned} F_j^0 (1 + 2\pi i k_l x_l^0 - 2\pi^2 k_l k_m x_l^0 x_m^0 + \dots) \\ = \sum_{\text{nodes } \gamma} f_j^\gamma (1 + 2\pi i k_l x_l^\gamma - 2\pi^2 k_l k_m x_l^\gamma x_m^\gamma + \dots), \end{aligned} \quad (2.20)$$

where x_l^0 is the l -coordinate of the actual position of the particle and x_l^γ is the l -coordinate of the grid-point positions γ . Equations (2.19) and (2.20) simply state that when a force is transferred to the grid from a particle centre, all its moments are also preserved. It is of course not possible to retain infinite terms in (2.20) and a truncated version is used.

The above analysis was limited to the case of point-forces; in most practical applications, however, we must take into account the particle size and represent the particles more accurately. Following the conventional Stokesian Dynamics approach, each particle is represented as an expansion of the force density on its surface in a series of moments at its centre. Using the same level of accuracy as in conventional SD, each spherical particle is represented as a point force \mathbf{F} , a doublet \mathbf{D} (corresponding to the torque and stresslet), a potential dipole \mathbf{d} (which is simply equal to $-a^2 \mathbf{F}/6\eta$, where a is the particle radius), and a potential quadrupole \mathbf{L} (which simply equals $-a^4 \mathbf{S}/10\eta$, where \mathbf{S} is the particle stresslet), while higher moments are neglected. The redistribution of the potential doublet, dipole, etc., follows the exact same procedure as for the force since they are just higher moments of the force singularity, and formulae equivalent to (2.20) can be derived. These formulae simply state, as was the case with (2.20), that the net force, torque, stresslet, etc. is conserved when transferred from the particle to the grid. The number of moments retained depends on the desired accuracy, and for this work an $O(k^3)$ matching is used, which means that the potential quadrupole is distributed to the grid with a leading error of $O(1)$. To obtain that level of accuracy, the force moments must be distributed as equivalent forces on a $5 \times 5 \times 5$ set of grid points centred at the node closest to the particle centre. Higher accuracy can be obtained with an increased number of matched terms (corresponding to increased number of grid points upon which the force is distributed), but the computational cost increases accordingly. The formulae for the moment distribution on the grid can be found in Sierou (2001) or obtained from the editorial office or the authors on request.

To summarize, the wave sum is calculated through the following steps: (i) each particle's force/torque/stresslet/etc. is distributed on a regular mesh over the computational domain, (ii) the FFT of the distributed forces is calculated and the expression in equation (2.18) is evaluated, (iii) the inverse FFT of step (ii) is calculated, which simply gives the wave-space sum of equation (2.17). Thus, at the end, for a given set of forces and higher moments on the particles, the contribution to the fluid velocity owing to the wave-space sum is calculated. This is a far-field contribution. The distribution of the forces on the grid and the simultaneous solution for all grid points includes effects from all particles.

This velocity corresponds to the fluid velocity on the given grid of nodes and not to the velocity at the particle centres, which is required for Faxén's laws. The fluid velocity at a particle's centre is obtained by interpolating from the grid to the particle position by a simple Lagrangian interpolation with a $5 \times 5 \times 5$ set of grid points for each particle. The general form of the Lagrangian interpolation is:

$$\mathbf{u}(\mathbf{x}_0) = \sum_{i=1}^5 \sum_{j=1}^5 \sum_{k=1}^5 h_i(x)h_j(y)h_k(z)\mathbf{u}_{i,j,k}, \quad (2.21)$$

where $\mathbf{u}(\mathbf{x}_0)$ is the interpolated velocity at position \mathbf{x}_0 , and $\mathbf{u}_{i,j,k}$ are the velocities at a given grid point (i, j, k) (the indices i, j and k now correspond to local nodes around each particle's centre). The interpolating polynomials are of the form:

$$h_i(x_0) = \frac{\prod_{j,j \neq i}(x_0 - x_j)}{\prod_{j,j \neq i}(x_i - x_j)}, \quad (2.22)$$

where x_i are the positions of the interpolation grid points in one dimension. The derivatives of the interpolants are then used to calculate the derivatives of the fluid velocity at the particle's centre.

2.3. Real-space sum contribution

In addition to the wave-space sum contribution, the real-space sum expression,

$$RS = \frac{1}{4\pi\eta} \sum_n F_l^n \left[\left(-\frac{\pi}{\alpha^{3/2}} \right) \phi_{1/2} r^2 \delta_{jl} + \alpha^{-1/2} \delta_{jl} + \left(\frac{\pi}{\alpha^{3/2}} \right) \phi_{1/2} x_l x_j \right], \quad (2.23)$$

must be evaluated. This equation gives the velocity at any point in the fluid owing to the presence of an array of point forces at distances \mathbf{r}_n ; similar expressions can be obtained in a straightforward manner for the velocity and any number of higher derivatives of the velocity due to the presence of a torque, stresslet, etc. Equation (2.23) suggests that the calculation of the real-space sum requires $O(N^2)$ operations, since the calculation of the velocity at each particle centre requires the sum over all particles (including a self-term). The incomplete Γ -function, ϕ_v , however, decays very fast, so that for small enough α the contribution of particle pairs that are not near neighbours will be very small and can be neglected without significant error. The elementary way to evaluate this sum is to sweep through all the particles, test whether their separation is less than a 'cutoff separation', r_c , and if so compute the two-particle contribution according to (2.23). Such an approach is clearly impractical, however, since it also gives an operation count that scales as N^2 .

To reduce the computational cost we arrange the particles in such a way that the tests for locating the neighbouring particles are only performed over a small subset N_n of the total number of particles N . This is accomplished by introducing a 'chaining mesh' (Hockney & Eastwood 1988), which is a regular lattice of $(M_x \times M_y \times M_z)$ cells covering the computational periodic box; the number of cells in each direction is such that the lengths of each side of each cell are always greater or equal to the aforementioned cutoff radius. Those particles that have non-zero contributions to the real-space sum of the hydrodynamic force must lie either in the same cell as any given particle, or in one of the 27 neighbouring cells, and therefore we only need to sweep through those neighbouring cells in search of a particle's nearest neighbours. The calculation of the infinite real-space sum is thus reduced to the calculation of the sum locally over only a small number of neighbouring particles – those particles

that are closer than the chosen value of the cutoff radius, r_c . (The choice of r_c and its effect on the accuracy of the calculation will be discussed in detail in the following section.) The real-space sum calculation can thus be done analytically in a straightforward manner; the detailed formulae (Sierou 2001) corresponding to the contributions of higher moments may be obtained upon request from the editorial office or the authors. Note that according to (2.23), when calculating the real-space correction to the fluid velocity at the centre of a particle, the real-space sum is performed over all neighbouring particles, including the reference particle, and thus a self-term contribution must also be calculated. After the completion of the real-space sum step for every neighbouring particle pair, a correction to the fluid velocity at the particle centre is evaluated, which is added to the already evaluated fluid velocity from the wave-space sum.

As a final note, in a sheared suspension, the unit cell and FFT grid points must deform along with the shear rate until the lattice repeats itself, typically at the end of one strain for simple shear flow.

2.4. Force laws

With the fluid velocity determined at each particle centre, the far-field force, torque and stresslet exerted on each particle are calculated from Faxén's laws:

$$\mathbf{F}_{ff} = -6\pi\eta a (\mathbf{U}_p - \mathbf{u}^\infty(\mathbf{x})) + 6\pi\eta a (1 + \frac{1}{6}a^2\nabla^2)(\mathbf{u}_{ff}), \quad (2.24a)$$

$$\mathbf{T}_{ff} = -8\pi\eta a^3 (\boldsymbol{\Omega}_p - \boldsymbol{\omega}^\infty(\mathbf{x})) + 4\pi\eta a^3 \nabla \times (\mathbf{u}_{ff}), \quad (2.24b)$$

$$\mathbf{S}_{ff} = \frac{20}{3}\pi\eta a^3 \mathbf{E} + \frac{20}{3}\pi\eta a^3 (1 + \frac{1}{10}a^2\nabla^2)(\mathbf{e}_{ff}), \quad (2.24c)$$

where \mathbf{u}_{ff} corresponds to the far-field fluid velocity evaluated at the centre of the particle, \mathbf{U}_p and $\boldsymbol{\Omega}_p$ are the particle translational and angular velocities, \mathbf{u}^∞ , $\boldsymbol{\omega}^\infty$ and \mathbf{E} are the bulk velocity, angular velocity and rate of strain, respectively, and \mathbf{e}_{ff} is the rate of strain of the far-field velocity.

The scheme described above corresponds to a prediction–correction method where the hydrodynamic forces on the particles are first used as input to calculate both the wave- and real-space contributions to the fluid velocity and then the forces are calculated again from the far-field velocities with the assistance of Faxén's laws. Determining the forces is the resistance formulation of the problem—for a given configuration of particles and particle velocities, the particle forces and force moments are calculated. It has recently been shown (Ichiki & Brady 2001) that the method of reflections (which corresponds to the simplest iterative scheme for the inversion of the mobility matrix) can be divergent when more than two particles are present.

An alternative way to invert the mobility matrix, which is still largely based on the method of reflections, is to use a simple under-relaxation scheme and replace (2.24) by:

$$\mathbf{F}_{ff} = [-6\pi\eta a (\mathbf{U}_p - \mathbf{u}^\infty(\mathbf{x})) + 6\pi\eta a (1 + \frac{1}{6}a^2\nabla^2)(\mathbf{u}_{ff})]\omega + \mathbf{F}_{ff}^{old}(1 - \omega), \quad (2.25)$$

where ω is the relaxation parameter, and \mathbf{F}_{ff}^{old} is the previous guess for \mathbf{F}_{ff} . In practice, values of ω as low as 0.05 may be necessary for very large, or very dense systems, which leads to a large number of iterations when the initial guess is poor. The relaxation parameter ω appears to have no significant effect on the dynamics of the system after the system is relaxed to a quasi-steady state (typically after the particles have moved the length of the unit cell). For a fixed geometry, as in the lattice calculations described in § 3, the relaxation parameter has no effect. As suggested by Ichiki & Brady (2001), the breakdown of the method of reflections can also be

overcome by the use of a conjugate-gradient-type iterative method. The generalized minimum residual method (GMRES) was used here, and it was found that very accurate results can always be obtained with very few iterations. (It should be noted that the fact that the mobility matrix is never constructed does not introduce any further difficulties since in most iterative methods matrices only appear as products with vectors.)

We also briefly note here that the particle pressure (Morris, Jeffrey & Brady 1993) can be readily calculated following the same procedure. The wave-space contribution is obtained as part of the FFT procedure since it simply corresponds to the fluid pressure at the centre of each particle, while the real-space sum part is calculated analytically.

2.5. Near-field interactions

Lubrication interactions are included in a manner very similar to conventional SD. The two-body resistance matrix \mathbf{R}_{2B} of equation (2.4) is calculated from the known exact results, and the part of the two-body interactions already included in the far field (\mathbf{R}_{2B}^∞) is subtracted as in conventional SD. The chaining mesh is again used so that only interactions between neighbouring particles will be included, thus reducing the computational cost from N^2 to NN_n , where N_n is now the number of particles closer to the cutoff radius for lubrication, which is $4a$. To further reduce both the computational cost and the memory requirements the near-field resistance matrices are stored in a sparse form, i.e. for each particle only the non-zero contributions from neighbouring particles are stored, reducing the dimensions of \mathbf{R}_{FU} from $6N \times 6N$ to $6N \times 6N_n$. This procedure allows every multiplication with any of the near-field resistance matrices to be an $O(N)$ operation.

2.6. Time integration

After both the far-field force and the near-field resistance matrices are calculated, the new particle velocities must be calculated from (2.1), which, in the absence of particle inertia, gives

$$0 = -\mathbf{R}_{FU,nf} \cdot (\mathbf{U}_p - \mathbf{u}^\infty(\mathbf{x})) + \mathbf{R}_{FE,nf} : \mathbf{E}^\infty + \mathbf{F}^p + \mathbf{F}_{ff}, \quad (2.26)$$

where now $\mathbf{R}_{FU,nf}$ is a sparse, symmetric, positive-definite matrix. Note that since the far-field interactions are not included in the resistance matrix, the unit tensor \mathbf{I} is added to $\mathbf{R}_{FU,nf}$ and subtracted from equation (2.26) to assure positive definiteness. The most efficient way to solve a sparse symmetric positive-definite system is by an iterative method, where the computational times can scale linearly with the matrix size since the sparse matrix multiplications required for an iterative method scale as N . However, a carefully designed iterative method must be applied since the number of iterations can increase significantly with increasing N . It is also worth noting that in (2.26) the far-field force is considered known, implying that the previous step velocities have been used for its calculation; alternatively, an additional iterative scheme can be employed such that the velocities used in the far-field calculation are also updated. In most problems, either the far-field force is small in magnitude (e.g. sheared systems) or it does not change rapidly with time (e.g. sedimenting systems), and in both cases the two approaches give results that are statistically indistinguishable; the simpler approach is therefore used hereinafter.

Conjugate gradient methods provide a quite general means of solving the $N \times N$ linear system

$$\mathbf{A} \cdot \mathbf{y} = \mathbf{b}. \quad (2.27)$$

The advantage of these methods is that they reference matrix \mathbf{A} only through its product with a vector, an operation that can be very efficient for properly stored matrices. The basic idea behind them for the simplest case of a symmetric positive-definite matrix is minimizing the function

$$f(\mathbf{y}) = \frac{1}{2} \mathbf{y} \cdot \mathbf{A} \cdot \mathbf{y} - \mathbf{b} \cdot \mathbf{y}, \quad (2.28)$$

which is equivalent to its gradient $\nabla f = \mathbf{A} \cdot \mathbf{y} - \mathbf{b}$ being zero, thus giving the solution to (2.27). The minimization proceeds by generating a succession of search directions and improved minimizers until the required accuracy is reached. The ordinary conjugate gradient method works well—the number of iterations required for convergence is reasonably small—for matrices that are well-conditioned, i.e. matrices with eigenvalues that are tightly clustered. This suggests using these methods on a preconditioned form of equation (2.27):

$$(\tilde{\mathbf{A}}^{-1} \cdot \mathbf{A}) \cdot \mathbf{y} = \tilde{\mathbf{A}}^{-1} \cdot \mathbf{b}. \quad (2.29)$$

The matrix $\tilde{\mathbf{A}}$ is called a preconditioner and its purpose is that now the matrix $(\tilde{\mathbf{A}}^{-1} \cdot \mathbf{A})$ will be well-conditioned, i.e. ‘close’ to the identity matrix.

Torres & Gilbert (1996) applied the conjugate gradient method for the sparse positive-definite near-field hydrodynamic resistance matrix $\mathbf{R}_{FU,nf}$ for different preconditioners and established that good convergence rates are possible when an incomplete Cholesky preconditioner is used. The incomplete Cholesky factor \mathbf{L}_0 is constructed following the same algorithm as for the calculation for the complete Cholesky factor, except that a specific element $L_{0,ij}$ is calculated only when the corresponding element $R_{FU,ij}$ is non-zero and is otherwise set equal to zero. Thus, for each $i = 1, 2, \dots, 6N$

$$L_{0,ii} = \left(R_{FU,ii} - \sum_{k=1}^{i-1} L_{0,ik}^2 \right)^{1/2}, \quad (2.30)$$

and then for $j = i + 1, i + 2, \dots, 6N$

$$L_{0,ji} = \begin{cases} \frac{1}{L_{0,ii}} \left(R_{FU,ij} - \sum_{k=1}^{i-1} L_{0,ik} L_{0,jk} \right) & \text{for } R_{FU,ij} \neq 0, \\ 0 & \text{for } R_{FU,ij} = 0. \end{cases} \quad (2.31)$$

The main advantage of the incomplete Cholesky preconditioner is that it captures the essential physics of the lubrication forces between two nearly touching particles while remaining as sparse as the initial \mathbf{R}_{FU} . If only the non-zero elements of \mathbf{L}_0 are computed and unnecessary operations with zero-valued elements are avoided, the preconditioning step remains an $O(N)$ operation.

Following Torres & Gilbert (1996), a variation of the incomplete Cholesky factor is used where ‘blocks’ of the matrix \mathbf{L}_0 are handled as opposed to individual elements. Since the resistance matrix consists of 6×6 blocks that map the velocity of one particle to the force/torque on another particle, the entire corresponding block of \mathbf{L}_0 is computed if any of the elements in the corresponding \mathbf{R}_{FU} block are non-zero. Furthermore, they suggest a particle re-ordering according to the proximity of the particles to one another so that a more ‘ordered’ form of the resistance matrix can be used. The reverse Cuthill–McKee method (Cuthill & McKee 1969; George & Liu 1981) was found to achieve the best overall results, and we use this approach in our method. The benefit of particle reordering is both a more rapid convergence and a minimization of the breakdowns of the Cholesky preconditioner. Since an




FIGURE 2. The number of iterations required for convergence of 10^{-4} , ---, with and —, without preconditioning ($N = 512$, $\phi = 0.45$, sheared system). The importance of an appropriate preconditioner becomes increasingly important for large strains as particle clusters form and the resulting resistance matrix, $\mathbf{R}_{FU,nf}$, becomes increasingly ill-conditioned.

incomplete version of the Cholesky factor is calculated, there is no guarantee, even for a symmetric positive-definite matrix, that square roots of negative numbers will not occur. To solve this problem qI , where q is a positive number large enough so that positive definiteness is obtained, is added to \mathbf{R}_{FU} . Although this addition solves the problem, it increases the number of iterations required for convergence, and it was empirically observed by Torres & Gilbert that the reordering of the particles reduces significantly the number of factorization breakdowns.

We should emphasize here that the use of the preconditioner has a very important effect on the number of iterations required for convergence, especially for dense sheared systems. In the absence of preconditioning, the number of iterations can be of the order of 100–200 when particle clusters are forming and the interparticle separations become very small. Sangani & Mo (1996) also report iterations of the same order for their non-preconditioned system. The use of the preconditioner, however, decreases the number of iterations dramatically and most systems, including dense sheared suspensions, can be solved in less than 10 iterations. Figure 2 demonstrates the effect of the preconditioner for a sheared system of $N = 512$ particles at volume fraction $\phi = 0.45$; the necessary number of iterations to achieve an accuracy of 10^{-4} in the residual calculation is significantly smaller for the preconditioned case. The effect of the preconditioner becomes even more significant with increasing strain; as the particles move closer together and form clusters, the non-preconditioned iterative procedure requires consistently more than 100 iterations for convergence.

The iterative procedure works well when the near-field lubrication matrix is reasonably well-behaved—in practice, when the interparticle distances are no less than $(10^{-5}–10^{-6})a$. In order to assure these minimum separations in a dynamic simulation, a repulsive interparticle force is generally required. When shearing a dense suspension in the absence of Brownian and/or interparticle forces, particle overlaps can occur even for extremely small time steps (Dratler & Schowalter 1996). The presence of an interparticle force can have a significant effect on the properties of some systems (e.g. sheared dense suspensions) and its impact is a matter for a separate study. The main point that we wish to make here is that our new method does

not impose any more restrictions on the minimum interparticle separation compared to conventional SD, because the implicit inversion of the near-field resistance matrix captures the same physics as the inversion of the full resistance matrix. Since we shall validate the method only with non-dynamic results where the interparticle force will always be zero, the issue of the interparticle force will not be mentioned further.

2.7. Total operations count for each time step

The most time consuming parts of our algorithm are the calculation of the far-field and near-field interactions and the inversion of the resistance matrix. The inversion of the resistance matrix, when done iteratively as described in the previous section, can be an $O(N)$ operation depending on the preconditioner. The near-field interactions and the real-space contribution to the far-field, as was already discussed, can be calculated in $O(N)$ operations with the introduction of a chaining mesh. The wave-space contribution of the far-field interactions, on the other hand, is calculated in $O(N_m^3 \ln N_m)$, where N_m is the total number of mesh points in each direction. In order for this part of the calculation to be $O(N)$, N_m must scale as $N^{1/3}$. The number of mesh points N_m determines the accuracy of the wave-space sum calculation, and therefore our choice of N_m should be a compromise between desired accuracy and computational cost. This is where the choice of the splitting parameter α becomes important.

For small values of α , the wave-space sum constitutes the largest part of the contribution, and therefore a large number of FFT points per particle are necessary for an acceptable degree of accuracy. However, if a large value of α is used, then the wave-space sum contribution is small and even low accuracy in its calculation leads to an acceptable overall error. The real-space sum on the other hand, requires a small value of α to minimize its error since only a few neighbouring particles are used for the calculation of the real-space sum. By balancing those two requirements an optimum value of α can be found that will give both acceptable accuracy and reasonable computational cost.

To demonstrate the effect of the method parameters—splitting parameter α , cutoff radius for the calculation of the real-space sum r_c , and number of mesh-points N_m for the FFT calculation—on the Ewald-sum calculation, we study the following problem. A small number of particles ($N = 30$) is placed randomly in a unit cell of length $L = 16a$, corresponding to a volume fraction of 3%. An external force is imposed on all particles and the resulting far-field fluid velocities at the centre of each particle are calculated according to (2.16). The same calculation is repeated by using the exact formulation of the Ewald sum of (2.16); this is done by including the interactions between all the particles not only in the unit cell, but also in 10 neighbouring cells in each direction. The relative error between the exact Ewald-sum calculation and the PME calculation is then determined as $(u_{fluid,PME} - u_{fluid,exact})/u_{fluid,exact}$ and averaged over all the positions of the particles. Figure 3 shows this relative error for different values of α , wave-space sum discretization dx ($= L/N_m$, where L is the length of the unit cell and N_m is the number of FFT points in each direction) and cutoff radius r_c . It can be seen clearly that for every choice of dx and real-space sum cutoff radius, an optimum α exists. This is expected since, as was already mentioned, for increasing α the error in the wave-space sum calculation decreases, while the error in the real-space sum calculation increases. For a given dx , the error decreases with increasing r_c , while the optimum choice of α increases with increasing r_c . In other words, for a fixed accuracy of the wave-space sum calculation, corresponding to a




FIGURE 3. The relative error for the calculation of the fluid velocity for a system of $N = 30$ particles in a gravitational field. The error is presented for different values of the wave-space sum discretization, dx , and different values of the real-space sum cutoff radius, r_c , as a function of the splitting parameter α . (a) $dx = 0.125a$, (b) $0.25a$, (c) $0.5a$, (d) $1.0a$.

fixed number of FFT points, the overall accuracy will increase as the accuracy of the real-space sum calculation increases (cutoff radius for the real-space sum increases). In addition, the increased accuracy of the real-space sum calculation is more important for larger values of the splitting parameter α ; for small values of α the real-space-sum contribution is very small and its accuracy does not affect the overall accuracy significantly. On the other hand, for a fixed cutoff radius and increasing number of FFT points, the optimum value of α now decreases, since increased accuracy in the wave-space sum calculation now allows smaller values of the splitting parameter. It is also worth noting that the minimum error is completely determined by the choice of dx and r_c , and for some cases no choice of α gives acceptable behaviour; for example, for the case of $r_c = 2a$, $dx = 1.0a$, the minimum error is close to 30%. For two FFT points per particle radius ($dx = 0.5a$, figure 3c) and $r_c = 6 - 7a$, accuracy within 1% can be obtained with a choice of $\alpha \approx 12 - 15$, and for most applications this range of parameters will be used since it was found also to give an acceptable computational cost.

Figure 4 shows the computational time for one time step as a function of the number of particles N . The times correspond to a shearing problem at volume fraction $\phi = 0.45$; one iteration is performed for the convergence of the far-field interactions, and the inversion of the near-field matrix takes 8–10 iterations for all numbers of particles shown. All the runs were performed on a single DEC Alpha AXP 21164 processor and the times are given in seconds. The data are in good agreement with the proposed $N \ln N$ scaling set by the scaling of the FFT calculations.




FIGURE 4. The CPU time (in seconds) for one time step of a shearing simulation as a function of the number of particles N . Symbols, ASD; \cdots , $N \ln N$ scaling.

3. Results – simple cubic arrays

In this section, the accuracy of the method is tested for the case of simple cubic arrays of spheres, a case for which a number of analytical and simulation results are available. Different choices for the method parameters are presented, which further demonstrates – this time for a problem with a known solution – that the PME calculation of the far-field contribution can give acceptable accuracy with low computational cost.

3.1. Sedimentation of cubic arrays: sensitivity to method parameters

The first case we consider is a set of 8 particles sedimenting in a simple cubic lattice. Since the system is periodic, the number of particles is not important and our main goal is to check the accuracy of the far-field calculations. It is easy to show that for identical spherical particles on a cubic lattice subject to a constant external force, the pairwise additive near-field forces will always sum to zero regardless of the volume fraction; therefore, the accuracy of the calculation of the far-field is the only issue. The exact infinite sums presented in § 2.1 are expected to give, within the number of moments included, results in exact agreement with the multipole moments analytical results of Zick & Homay (1982) and the conventional SD results, since an equivalent approach is used for their derivation. The introduction of an FFT method, however, truncates the wave-space sum to the number of mesh points used, while the introduction of a near-field-like approach to the real-space sum truncates it to only the contributions of a few neighbouring particles. Nevertheless, with a careful choice of the splitting parameter α , the infinite sums decay rapidly and acceptable accuracy can be obtained.

Figure 5 presents the sedimentation velocity for a relatively low volume fraction, $\phi = 0.064$, as a function of both the splitting parameter α and the wave-space sum discretization dx ; the cutoff radius for the calculation of the real-space sum is defined so that only one neighbour in each direction is included, $r_c = \sqrt{3} \times 4a \approx 7a$. Good accuracy for the sedimentation velocity is obtained for a wide range of α and a desired plateau is observed for α near the optimum. The plateau in figure 5 corresponds to the range of α where the inclusion of only one neighbour in the real-space sum and the discretization of the wave-space sum provide a nearly exact evaluation of the sums in




FIGURE 5. The dependence of the sedimentation velocity of an SC array of spheres on α and dx ; $N = 8$, $\phi = 0.064$, $L = 8.0$, $r_c = \sqrt{3} \times 4a$.

equation (2.16). Obviously as the wave-space sum discretization (dx) becomes finer, a wider range of α (starting from smaller values) can be used with no significant loss of accuracy. Near the optimum value of α , however, values of dx near 0.5, or two points per particle radius, are sufficient for satisfactory accuracy, which allows us to treat large systems using a relatively small number of FFT discretization points.

Figures 6 and 7 present the same analysis for the case of $\phi = 0.5236$, a close packed simple cubic array. Figure 6 corresponds to the inclusion of only one neighbour in the real-space sum, a cutoff radius of $r_c = \sqrt{3} \times 2a \approx 3.46a$, while figure 7 corresponds to the inclusion of two neighbours, $r_c = \sqrt{3} \times 4a \approx 7a$. The case of $r_c \approx 3.46a$ gives satisfactory results only when a fine wave-space discretization is used, in agreement with the observations of figure 3. In figure 7 on the other hand, the larger value of r_c allows fewer FFT points to be used and permits a greater range of α , in this case going to larger values since more neighbouring particles have already been included. It should also be noted that for dense systems the same value of r_c corresponds to a larger number of near neighbours and therefore a larger computational cost in the real-space sum. On the other hand (for the same number of particles), dense systems correspond to much smaller unit cells and therefore the extra computational cost for the inclusion of more near neighbours is counteracted by the need for fewer FFT points for a given discretization (dx) and vice versa.

Figure 8 presents the sedimentation velocity for volume fraction $\phi = 0.064$ and $\alpha = 16$ as a function of the wave-space sum discretization parameter, dx , and for two different positions of the particles – on the grid points, and positioned midway between two grid points. This tests both the significance of the discretization in the wave-space sum and the accuracy of the force distribution/interpolation scheme. Again, for the given value of α corresponding to a value near the optimum, no significant loss of accuracy occurs for a wide range of dx . Also, the positions of the particles on or off the FFT grid is of little importance even for the largest values of dx . The same parametric analysis was also conducted for the case of sheared suspensions; the dependence on the far-field method parameters was even less apparent in this case since lubrication is now important and the calculation of the near-field interactions dominates the accuracy of the method.




FIGURE 6. The dependence of the sedimentation velocity of an SC array of spheres on α and dx ; $N = 8$, $\phi = 0.5236$, $L = 4.0$, $r_c = \sqrt{3} \times 2a$.




FIGURE 7. The dependence of sedimentation velocity of an SC array of spheres on α and dx ; $N = 8$, $\phi = 0.5236$, $L = 4.0$, $r_c = \sqrt{3} \times 4a$. —, $dx = 0.0625$; ---, 0.125; \cdots , 0.25; - - -, 0.5; — · —, 1.0.

3.2. Sedimentation of cubic arrays: volume fraction dependence

We now discuss the average sedimentation velocity of a simple cubic array of spherical particles as a function of the volume fraction ϕ . As was already mentioned, lubrication plays no role in this problem and the accuracy of the resulting sedimentation velocity is solely determined by the accuracy of the far-field calculation. Since the mobility matrix is never calculated, the sedimentation velocity can be found by simply imposing an external force (gravity) on all the particles and calculating the resulting steady-state (converged) velocities. Also, since all particles are identical in a periodic lattice, they all have the same velocity and therefore the relative distances between particles always remain constant; in other words, there is no difference between a dynamic and a static approach to the problem. Figure 9 shows a comparison of the non-dimensional sedimentation velocity of a simple-cubic array of spheres obtained by the ASD method with the conventional SD results of Brady *et al.* (1988) and the theoretical calculations by Saffman (1973) and Zick & Homay (1982). Saffman's calculations are




FIGURE 8. The dependence of the sedimentation velocity of an SC array of spheres on particle position and dx ; $N = 8$, $\phi = 0.064$, $L = 8.0$, $r_c = \sqrt{3} \times 4a$, $\alpha = 16$. \circ —, Particles on grid points; \blacktriangle —, off grid points.




FIGURE 9. The dependence of the sedimentation velocity of an SC array of spheres on the volume fraction. Δ , Accelerated Stokesian Dynamics (ASD) results; \blacklozenge , Stokesian Dynamics (SD) results (Brady *et al.* 1988); \times , exact results of Zick & Homsy (1982); \cdots , point-force (PF) solution of Saffman (1973).

for point-force particles and result in a sedimentation velocity $U/U_0 = 1 - 1.738\phi^{1/3}$, where U_0 is the settling velocity of an isolated particle under the same conditions. Zick & Homsy's results are exact, as they used sufficient moments, to achieve convergence to the exact limit. For the same number of moments, the SD and ASD results are identical for all volume fractions—the new method has the exact same accuracy as conventional SD.

For low volume fractions all methods give identical results. For high volume fractions ASD overestimates the sedimentation velocities in the exact same manner that conventional SD does. Only the point force and the second moment of that point force have a contribution in this formulation. Odd moments, although included, have no effect on the sedimentation velocity of cubic arrays, and as a result only the second-order method results of Zick & Homsy can be reproduced. For high-volume fractions, higher moments are needed; the SC results of Zick & Homsy and the




FIGURE 10. The dependence of the shear viscosity functions α and β of an SC array of spheres on the volume fraction. Δ , Accelerated Stokesian Dynamics (ASD) results; \times , ASD results in the absence of lubrication; \blacklozenge , exact analytical results of Hoffman (1999); \cdots , $-\cdot-$, asymptotic results of Hoffman for the $\phi \rightarrow 0$ and; $\phi \rightarrow \phi_{max}$ limits.

random array results of Ladd (1990) indicate that at least four moments are required for reasonable accuracy near maximum packing. This can easily be done in principle following the methodology outlined above, but at an increased computational cost per particle; the overall method will still scale as $N \ln N$, however.

3.3. Shear viscosity of cubic arrays

To calculate the shear viscosity of a cubic array of spheres, the relationship between the bulk stress and the imposed rate of strain is required. Again, since the far-field mobility or resistance matrices are never calculated, we simply impose an external rate of strain on the particles and calculate the resulting induced far-field hydrodynamic stresslets. To those values, the near-field contribution calculated from the near-field resistance matrix, $\mathbf{R}_{SE,nf}$, is added; for a cubic lattice there is no \mathbf{R}_{SU} contribution because all the particles move with the velocity of the bulk flow ($\mathbf{U}_p = \mathbf{u}^\infty$). It should also be noted that the viscosity is calculated for an instantaneous configuration corresponding to a non-distorted cubic lattice. At a later instant in time, the cubic lattice will be distorted as the particles move with the bulk flow and a different instantaneous viscosity can be evaluated. It is a straightforward matter to calculate the stresslets at each instant; for both simple shear and planar pure straining motion, the unit cell returns to its non-distorted shape periodically and only a finite number of configurations need to be sampled. Since all analytical results available correspond to the non-distorted cubic cell we present only this case here. The particle stress for a simple cubic array can be described by two independent parameters α and β , which are only functions of the lattice geometry and the particle volume fraction; these two functions correspond to pure straining and simple shear flow, respectively, and the viscosity for any other linear flow can be obtained as a linear combination of these two cases (Zuzovsky, Adler & Brenner 1983; Nunan & Keller 1984). (We retain the notation α and β for consistency with previous authors; the viscosity function α is not to be confused with the splitting parameter of the same symbol.)

Figure 10 shows a comparison of α and β for a simple cubic lattice obtained by our ASD simulations with the asymptotic and exact results of Hoffman (1999). Hoffman's results reproduce the well-established analytical results of Nunan & Keller for low to intermediate volume fractions, whereas the high volume fraction asymptotes are generated with higher accuracy. Specifically, Hoffman's asymptotic expansions for

high volume fractions can be expressed in terms of $\epsilon = 1 - (\phi/\phi_{max})^{1/3}$, as follows:

$$\left. \begin{aligned} \alpha &= \frac{3}{16}\pi\epsilon^{-1} + \frac{27}{80}\pi \ln \epsilon^{-1} - 2.85 - 1.3\epsilon \ln \epsilon^{-1} + O(\epsilon), \\ \beta &= \frac{1}{4} \ln \epsilon^{-1} - 0.604 - 0.30\epsilon \ln \epsilon^{-1} + O(\epsilon), \end{aligned} \right\} \quad (3.1)$$

where the constant term has been corrected from the Nunan & Keller result (Hoffman claims there is a sign error in those results) and the coefficient of the next term ($\epsilon \ln \epsilon^{-1}$) is also evaluated. The far-field only (absence of lubrication) results are also presented and are in exact agreement with the low volume fraction asymptotic behaviour. In contrast to sedimentation, the high-volume-fraction asymptote is also reproduced with very good accuracy; for shearing motion, lubrication is important and it precisely captures the two-body singular effects described by the high- ϕ asymptotes. The inclusion of higher moments in the far field is not as important as it was in the case of sedimentation, although for intermediate volume fractions our viscosities are slightly larger than the analytical results of Hoffman, as was also the case for the conventional SD algorithm.

3.4. Spin viscosity of cubic arrays

The spin viscosity of a cubic array gives the relationship between the torque exerted on each particle \mathbf{T} and its angular velocity $\boldsymbol{\omega}$. (Since all particles in a cubic array are identical, no averaging is necessary.) Symmetry of the cubic lattice reduces the calculation of the spin viscosity to the calculation of one scalar, ζ , such that:

$$T_i = -\zeta \omega_i. \quad (3.2)$$

Zuzovsky *et al.* (1983) determined the high- and low- ϕ asymptotes for ζ , whereas Hoffman (1999) calculated ζ for all volume fractions and improved upon the asymptotic expansions by including higher-order terms. Hoffman's asymptotic expansion is: $\zeta = \pi \ln((\phi/\phi_{max})^{-1/3} - 1)^{-1} - 3.15$; the evaluation of the constant term improves the expansion significantly since it is now valid for much lower volume fractions. Figure 11 compares the results obtained for the spin viscosity ζ as a function of volume fraction to the exact results of Hoffman. Excellent agreement is observed for all volume fractions, and the low- and high-volume-fraction asymptotic results are reproduced very accurately.

4. Results – random suspensions

In the following sections we present the results of a series of Monte Carlo simulations from which the hydrodynamic properties of random hard-sphere suspensions can be obtained. The Monte Carlo approach consists of generating several samples, calculating the transport properties of interest for each sample, and then averaging, in order to obtain values that describe the macroscopic behaviour of the system. Different sampling techniques were used for different volume fractions; for $\phi < 0.49$ the particles were placed in an arbitrary initial configuration and then moved using a random-stepping routine in order to ensure that the sample was disordered—for $\phi > 0.49$ the sampling technique is discussed in detail in the following sections. We first present the high-frequency dynamic viscosity of random suspensions for volume fractions ranging from infinite dilution up to random close packing along the metastable fluid branch of the hard-sphere system. For volume fractions below the freezing point ($\phi = 0.494$), analytical and other simulation results are available for comparison. For volume fractions above the freezing point, the system is maintained




FIGURE 11. The dependence of the spin viscosity ζ of an SC array of spheres on the volume fraction. Δ , Accelerated Stokesian Dynamics (ASD) results; —, exact calculation of Hoffman (1999); ---, Hoffman's low concentration analytical result; - · -, singular form as $\phi \rightarrow \phi_{max}$.

in the disordered metastable liquid state, and the singular behaviour of the viscosity as ϕ approaches random close packing ($\phi_{rcp} \approx 0.64$) is studied. We next turn briefly to the sedimentation problem. As has been shown in the past for conventional SD, the order of approximation is only accurate for low to intermediate volume fractions, and we therefore simply present the N -dependence of the sedimentation velocity for a given low volume fraction. Finally, we study the short-time self-diffusion coefficient—the mobility of a single particle in a suspension of force-free particles. This is an example of a case where the knowledge of the resistance matrix, and not just the far-field hydrodynamic force, is of importance, and we demonstrate how our method can still determine the short-time self-diffusion coefficient in $O(N \ln N)$ operations.

4.1. Viscosity of random suspensions: below the freezing point

The effective viscosity of a random dispersion of hard spheres, which is known as the high-frequency dynamic viscosity, has been studied in the past for volume fractions below the freezing point, and essentially exact, as well as low ϕ asymptotic, results are available (Batchelor & Green 1972; Ladd 1990). The effective viscosity is readily calculated by imposing an external shear flow and converging the far-field contributions and the particle velocities, while keeping the particles at fixed positions. In contrast to the cubic array, the externally imposed stresslet produces non-zero particle velocities that also need to be determined as part of the iterative procedure. The total particle stresslet is then calculated as a sum of the converged far-field contribution, \mathbf{S}_{ff} , and the near-field contribution $\mathbf{S}_{nf} = \mathbf{R}_{SE,nf} : \mathbf{E} - \mathbf{R}_{SU,nf} \cdot (\mathbf{U}_p - \mathbf{u}^\infty)$, which now also contains a non-zero velocity contribution.

Figure 12 shows the hydrodynamic viscosity of a random suspension of identical spheres as a function of the volume fraction for different numbers of particles in the range $N = 125$ –2000, and averaged over a number of independent configurations in the range 10–100. Since a much larger number of particles can now be used, fewer independent configurations are required giving statistical errors that are less than 2% for all volume fractions below the freezing point and all number of particles shown. (The error bars in figure 12 are omitted since they are always smaller than the size of the symbols.) The results of Ladd (1990), where up to 7 moments have been




FIGURE 12. The relative viscosity is plotted as a function of the volume fraction for different numbers of particles ($N = 125$ – 2000) and volume fractions. —, multipole-moment simulation results of Ladd (1990); +; \oplus , experimental results of van der Werff *et al.* (1989) and Shikata & Pearson (1994).

included in the calculations, are taken to be exact. As was the case for the cubic array, the importance of lubrication is such that very accurate results are obtained even when using a low-order moment approximation to the far-field interactions. (Ladd also observed that the effective viscosity is not sensitive to the inclusion of higher moments.) It is also apparent that the viscosity is insensitive to the system size, since statistically indistinguishable results are obtained for N ranging from 125 to 2000. Also shown are the experimental results of Van der Werff *et al.* (1989) and Shikata & Pearson (1994) obtained from oscillating Couette viscometry with a frequency sufficiently high that the distribution of the solid particles is unaffected by the shear flow and corresponds to the equilibrium hard-sphere structure. Very good agreement is observed, especially for the lower volume fractions (see below for results above the freezing point).

4.2. Viscosity of random suspensions: above the freezing point

Although a large number of analytical results exist for the effective viscosity of hard-sphere suspensions for volume fractions in the stable fluid region (below $\phi = 0.494$), very few theoretical or experimental results are available for systems along the metastable extension of the fluid branch. For volume fraction below $\phi = 0.494$, only a single stable fluid (disordered) phase exists; above this point, the phase diagram splits into a metastable fluid phase (leading to random close packing) and a stable ordered phase (leading to FCC crystals). This work will only be concerned with the metastable disordered branch of the phase diagram. This is an example where a large number of particles are necessary to capture the correct behaviour, as a macroscopic system at these high volume fractions will undoubtedly have clusters in some regions and freely mobile particles in others, which is difficult to model unless the unit cell is sufficiently large. In addition, even generating such a dense configuration can be




FIGURE 13. The relative viscosity is plotted as a function of the volume fraction for volume fractions above the melting point. —·—, extension of the empirical fit to Ladd's simulation results below the freezing point; ···, the same extension for the curve fitted to the van der Werff *et al.* (1989) results; — · —, a simple best fit curve of the form $A_1 \ln(1/\epsilon) + A_2$, which accurately represents the results for $\phi > 0.60$. +, van der Werff *et al.* (1989); ⊕, Shikata & Pearson (1994).

challenging since metastable systems have a tendency to crystallize and to introduce order into the structure.

Standard hard-sphere molecular dynamics or Monte Carlo algorithms cannot be used to generate random distributions above the freezing transition. Using a different approach with a large number of particles, however, hard-sphere microstructures at volume fractions up to random close packing can be generated. Although the term random close packing is widely used, it is not always clear whether it is a universal quantity, or whether it depends on the method used to generate any given configuration; the exact value of ϕ_{rcp} is also discussed and the value $\phi_{rcp} = 0.64$ was used here. (See Torquato, Truskett & Debenedetti (2000) for a discussion on whether random close packing is a well-defined state.) For the purposes of this work, the computational technique suggested by Rintoul & Torquato (1996*b*) was used to generate the random hard-sphere configurations. Starting from an initial set of random overlapping spheres, the spheres are expanded and simultaneously moved to reduce overlap; if the system becomes jammed (overlapping cannot be reduced) the spheres are shrunk and moved until the system becomes unjammed (see Clarke & Wiley 1987). This process of expansion and contraction is repeated until a desirable volume fraction is obtained. After that, the system is equilibrated using standard hard-sphere molecular dynamics. For volume fractions above the freezing point, the equilibration process must be closely monitored because there are two phenomena occurring simultaneously – the system moves from a non-equilibrium to an equilibrium state and at the same time it moves from the metastable branch to the stable ordered branch of the phase diagram. By monitoring the system pressure and by using the differences in the time scales between the equilibration and crystallization processes (it usually takes longer for the system to crystallize), it is possible to generate the desired configurations of random hard spheres for volume fractions very close to random close packing (Rintoul & Torquato 1996*a*, Speedy 1994). It was observed that small systems had a tendency to crystallize sooner, and therefore for some volume fractions up to 2000 particles were used. It is worth noting that it is most difficult to generate non-ordered configurations for volume fractions between 0.52 and 0.58 as




FIGURE 14. The relative viscosity is plotted as a function of $\ln(1/\epsilon)$, where $\epsilon = 1 - \phi/\phi_{rcp}$, for $\phi > 0.60$. A linear best fit is also shown.

the system tends to crystallize very fast. The higher volume fractions correspond to an amorphous glassy state, and the tendency for freezing is less severe. (See Speedy 1994, 1997 for more details on sampling and phase transitions.)

In figure 13, we present the high-frequency dynamic viscosity for volume fractions up to 0.635, along with the experimental results of van der Werff *et al.* (1989) and Shikata & Pearson (1994). Two extrapolated expressions, one from the experimental results of van der Werff *et al.* and one derived from the results of Ladd, are also presented. There is significant spread in the experimental data at high volume fraction that might be an indication of some freezing in some of the experimental data or simply represent the difficulty of measuring the volume fraction at high volume fractions. Within the error bars of the experimental data, however, our results are in reasonable agreement with the experiments. It is interesting to note that our results for $\phi < 0.60$ are in good agreement with the fitted curve given by Ladd (1990), which is a simple extrapolation of a semi-empirical fit of Ladd's results from lower volume fractions. This is not surprising; since the metastable branch is the continuous extension of the fluid branch, it is reasonable that the viscosity would also be a continuous extension of the viscosity from below the freezing point and can potentially be described by the same empirical equation. The empirical relation given by Ladd, however, does not predict any singular behaviour as maximum packing is approached, and therefore fails at very high volume fractions.

A very sharp increase in the viscosity is observed for volume fractions above 60%, suggesting a singular behaviour in the limit of $\phi \rightarrow \phi_{rcp}$. The exact form of this singular behaviour is not known. Results from lubrication theory for cubic lattices would suggest that the singular form should consist of both $1/\epsilon$ and $\ln \epsilon$ terms (corresponding to α and β of figure 10), where $\epsilon = 1 - (\phi/\phi_{rcp})^{1/3}$, but the relative amount of each term is unknown. We found the data to be well fit by $\eta_r = 15.78 \ln(1/\epsilon) - 42.47$ as seen in figure 14. A similar fit with both $1/\epsilon$ and $\ln \epsilon$ gave a very small coefficient (of order 10^{-2}) for the $1/\epsilon$ term. As far as we are able to tell at this point, the $\ln \epsilon$ behaviour accurately describes the numerical data.

4.3. Sedimentation velocity: dependence on system size

It is straightforward to calculate the average sedimentation velocity of a random suspension of spheres: after imposing an external force on each particle, the velocity of each particle is converged while the particles remain fixed at their random locations. As was discussed for the sedimentation of the cubic arrays, our method can predict




FIGURE 15. The sedimentation velocity of a random suspension of spheres at $\phi = 0.05$ and for different number of particles N . ASD results for $N = 125$ – 512 (\triangle) are plotted as a continuation of Ladd's (1990) results (∇) for smaller systems. The corrected infinite system sedimentation velocity (equation (4.1)) gives a constant value for all N .

accurate velocities only for low to moderate volume fractions; for higher volume fractions more moments are required to represent the particles correctly. Since this behaviour has been discussed and analysed in the past for the conventional SD method (Phillips *et al.* 1988), here we restrict ourselves to low volume fractions and demonstrate the dependence of the sedimentation velocity on the system size. Phillips *et al.* (1988) showed that the sedimentation velocity has a strong $N^{-1/3}$ dependence owing to the long-range effects of the periodic images of particles outside the unit cell. The motion is in essence a superposition of the sedimentation velocity of the dilute periodic array of images, $(\phi/N)^{1/3}$, with that for the random suspension (which is $O(\phi)$ at low ϕ). Figure 15 shows clearly the $N^{-1/3}$ dependence of the sedimentation velocity for a volume fraction of $\phi = 0.05$.

Mo & Sangani (1994) calculated the difference in the velocity induced at the centre of a test particle in a periodic suspension and a random suspension to be:

$$U_s = U_s(N) + 1.7601(\phi/N)^{1/3}\eta_0/\eta S(0)U_0 + O(\phi/N), \quad (4.1)$$

where $S(0)$ is the structure factor, and η and η_0 are the suspension and pure fluid viscosities, respectively (values for the suspension viscosity η are the high-frequency dynamic viscosity and were calculated in the previous section and given in figure 12). The structure factor can be estimated for the hard-sphere dispersion from the Carnahan–Starling approximation:

$$S(0) = \frac{(1 - \phi)^4}{1 + 4\phi + 4\phi^2 - 4\phi^3 + \phi^4}, \quad (4.2)$$

and the corrected sedimentation velocities can be calculated directly from (4.1). The corrected velocities are also presented in figure 15 and they do indeed give the same value for all N .

4.4. Short-time self-diffusion coefficient

The short-time self-diffusion coefficient is defined through the Stokes–Einstein relation:

$$\mathbf{D}_0^s = kT \langle \text{tr}(\mathbf{R}_{FU}^{-1}) \rangle, \quad (4.3)$$




FIGURE 16. The short-time self-diffusion coefficient for a random suspension of spheres at $\phi = 0.05$ and for different number of particles N . ASD results for $N = 125$ – 512 (\triangle) are plotted as a continuation of Ladd's (1990) results (∇) for smaller systems. The corrected infinite system self-diffusivity (equation (4.7)) gives a constant value for all N .

where the *trace* operator picks out only the diagonal elements of \mathbf{R}_{FU}^{-1} , and the angle brackets imply a sum over all particles and an average over all configurations. Since the far-field part of the resistance matrix \mathbf{R}_{FU} is never calculated, we must find a different approach to calculate the self-diffusion coefficient. In addition, we must isolate the diagonal elements of the resistance matrix for each particle, and therefore the knowledge of only $\langle \mathbf{R}_{FU}^{-1} \rangle$ is not adequate as it was for the calculation of the sedimentation velocity.

A straightforward way to perform this calculation in order N operations is to impose an external force, \mathbf{F}^g , with a Gaussian distribution on the particles, i.e. $\langle \mathbf{F}_g \rangle = 0$ and $\langle F_i^{g,m} F_j^{g,n} \rangle = \delta_{ij} \delta_{mn}$, where $F_i^{g,m}$ is the i -component of the force acting on particle m (superscripts m, n denote particles, while subscripts i and j denote Cartesian coordinates). Following the procedure described above, the resulting particle velocities can be calculated readily in order N . Although the mobility matrix is never calculated, the resulting particle velocities, U_i^n , still satisfy:

$$U_i^n = R_{FU,ijnm}^{-1} F_j^{g,m}. \quad (4.4)$$

Now form the product

$$U_i^n F_i^{g,n} = R_{FU,ijnm}^{-1} F_j^{g,m} F_i^{g,n}, \quad (4.5)$$

which gives after ensemble-averaging

$$\langle U_i^n F_i^{g,n} \rangle = \langle R_{FU,ijnm}^{-1} \rangle, \quad (4.6)$$

from which the short-time self-diffusion coefficient can be calculated directly. This approach was successfully employed to calculate the diffusion coefficient and study its dependence on both the number of particles and the volume fraction.

Like the sedimentation velocity, the translational self-diffusivity also shows a strong $N^{-1/3}$ dependence (Phillips *et al.* 1988). Ladd (1990) proposed the following expression to extract the infinite system diffusivity from the N -particle periodic diffusion coefficient:

$$D_0^s = D_0^s(N) + (\eta_0/\eta)(kT/6\pi\eta_0 a)[1.7601(\phi/N)^{1/3} - \phi/N], \quad (4.7)$$

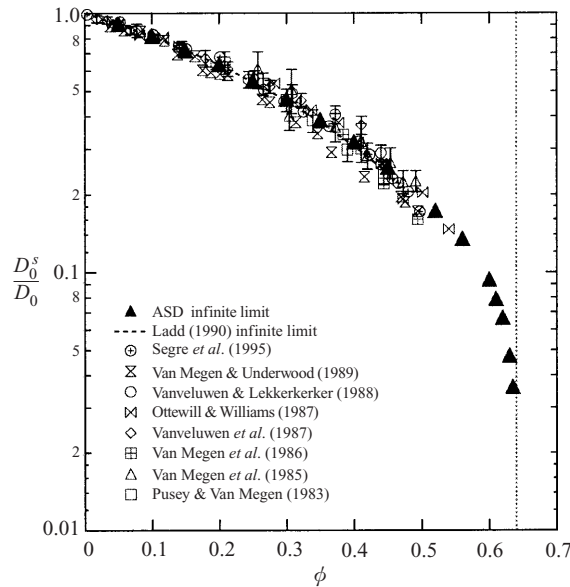


FIGURE 17. The short-time self-diffusion coefficient for an infinite random suspension as a function of the volume fraction ϕ (the values have been corrected according to equation (4.7) from values for $N = 512$). The results are compared with a variety of experimental results for volume fractions below the freezing point.




FIGURE 18. The short-time self-diffusion coefficient is plotted as a function of the inverse of the relative viscosity, for $\phi > 0.60$. Solid symbols, simulation results; ---, linear best fit.

where η and η_0 are the suspension and pure fluid viscosities, respectively. The Stokes–Einstein diffusivity of a single isolated particle is $D_0 = kT/6\pi\eta_0 a$. Figure 16 shows the self-diffusion coefficients for $\phi = 0.05$ and $N = 16$ –512. The corrected diffusivities, after applying the correction given by (4.7), are also shown verifying both the strong $N^{-1/3}$ dependence of the self-diffusion coefficient on the number of particles and the validity of equation (4.7), since a constant value can indeed be extracted for all N .

Figure 17 presents the short-time self-diffusion coefficients for volume fractions up to 0.635, along with a number of available experimental and simulation data. The self-diffusion coefficients are calculated from 20 realizations of $N = 512$ particles and the values are adjusted to correct for the $N^{1/3}$ dependence (limit for infinite N). Very good agreement is observed between experimental and simulation results for all

volume fractions below the freezing point (where experimental results are available). A sharp decrease in the self-diffusion coefficient, in agreement with the sharp increase in the viscosity, is observed for volume fractions above 60%. The self-diffusivity is expected to vanish at random close packing in a manner inversely proportional to the high-frequency dynamic viscosity. In figure 18, the self-diffusivity is plotted as a function of the inverse high-frequency dynamic viscosity, resulting in a clear linear scaling and therefore verifying that $D_0^s \approx 1/\eta$, in the limit of $\phi \rightarrow \phi_{rcp}$.

We note here that the accuracy of our simulation results is limited by the small number of moments used. It has been suggested in the past (Ladd 1990) that higher moments can influence the value of the self-diffusivity, especially for high volume fractions. No attempt has been made to correct for the inclusion of higher moments, although the very good agreement between our results and experimental results would suggest that the magnitude of the correction is indeed small.

5. Conclusions

We have described in detail a new method for calculating the hydrodynamic interactions among particles in a suspension at small Reynolds number based on the Stokesian Dynamics method, but with a significantly more favourable computational cost of $N \ln N$. The new method avoids the expensive calculation of the far-field mobility matrix in favour of the direct calculation of the far-field hydrodynamic force, and uses a carefully chosen preconditioning scheme to reduce the computational cost of any iterative matrix inversions. The results of the method are in excellent agreement with those obtained from conventional Stokesian Dynamics, and much larger systems can now be simulated with the same accuracy. The power of the new method is demonstrated in the calculation of the high-frequency dynamic viscosity and short-time self-diffusivity of suspensions for volume fractions above the freezing point, a case where a large number of particles is necessary to capture the correct behaviour. This work was limited to the validation and evaluation of the method and to some characteristic results corresponding to static (not evolving in time) systems. This is by no means a restriction of the method; in fact, the method is ideally suited for dynamic simulations where good initial guesses from the previous step in time exist for all the iterative procedures. Results from dynamical studies will be the subject of future publications.

It should also be noted that the methodology used here can be applied to other problems in Stokes flow – non-spherical particles, deformable drops, etc. – with the same performance characteristics. There are also analogous problems governed by Laplace's equation (e.g. effective conductivity) or the biharmonic equation (e.g. linear elasticity) that could also be addressed in $O(N \ln N)$ operations using the ASD methodology.

This work was supported in part by grants NAG3-2166 and NAG8-1661 from NASA. The authors benefited greatly from discussions with Professor J. J. L. Higdon on the PME method for Stokes flow. An anonymous referee is thanked for the suggestion of the random forces to determine the short-time self-diffusivity.

REFERENCES

- BANCHIO, A. & BRADY, J. F. 2001 Accelerated Stokesian Dynamics: Brownian motion (in preparation).
- BATCHELOR, G. K. & GREEN, J. T. 1972 The determination of the bulk stress in a suspension of spherical particles to order c^2 . *J. Fluid Mech.* **56**, 401–427.

- BEENAKER, C. W. J. 1986 Ewald sum of the Rotne-Prager tensor. *J. Chem. Phys.* **85**, 1581–1582.
- BRADY, J. F. & BOSSIS, G. 1988 Stokesian Dynamics. *J. Fluid Mech.* **20**, 111–157.
- BRADY, J. F., PHILLIPS, R. J., LESTER, J. C. & BOSSIS, G. 1988 Dynamic simulation of hydrodynamically interacting suspensions. *J. Fluid Mech.* **195**, 257–280.
- CLARKE, A. S. & WILEY, J. D. 1987 Numerical simulation of the dense random packing of a binary mixture of hard spheres: Amorphous metals. *Phys. Rev. B* **35**, 7350–7356.
- CUTHILL, E. & MCKEE J. 1969 Reducing the bandwidth of sparse symmetric matrices. *Proceedings 24th National Conference of the Association for Computing Machinery*, Brandon Press, NJ, 157–172.
- DARDEN, T., PEDERSEN, L., TOUKMAJI, A., CROWLEY, M. & CHEATHAM, T. 1997 Particle-mesh based methods for fast Ewald summation in molecular dynamics simulations. *Ninth SIAM Conference on Parallel Processing for Scientific Computing*, 1997.
- DARDEN, T., YORK, D. & PEDERSEN, L. 1993 Particle Mesh Ewald: An $n\log(n)$ method for Ewald sums in large systems. *J. Chem. Phys.* **98**, 10089–10092.
- DRATLER, D. I. & SCHOWALTER, W. R. 1996 Dynamic simulations of suspensions of non-Brownian hard spheres. *J. Fluid Mech.* **325**, 53–77.
- DURLOFSKY, L. J., BRADY, J. F. & BOSSIS G. 1987 Dynamic simulations of hydrodynamically interacting particles. *J. Fluid Mech.* **180**, 21–49.
- ESSMANN, U., PERERA, L., BERKOWITZ, M. L., DARDEN, T., LEE, H. & PEDERSEN, L. G. 1995 A smooth particle mesh Ewald method. *J. Chem. Phys.* **103**, 8577–8593.
- FOSS, D. R. & BRADY, J. F. 2000 Structure, diffusion and rheology of Brownian suspensions by Stokesian Dynamics simulations. *J. Fluid Mech.* **407**, 167–200.
- GEORGE, A. & LIU, J. W. 1981 Computer solution of large sparse positive definite systems. Prentice-Hall, Englewood Cliffs, NJ.
- GUCKEL, E. K. 1999 Large simulations of particulate systems using the PME method. PhD thesis, University of Illinois at Urbana-Champaign.
- HASIMOTO, H. 1959 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**, 317–328.
- HOCKNEY, R. W. & EASTWOOD, J. W. 1988 Computer simulation using particles. I.O.P. Publishing, Philadelphia.
- HOFFMAN, J. M. A. 1999 Hydrodynamic Properties of Dense Colloidal Crystals. PhD thesis, Technische University Eindhoven.
- ICHIKI, K. & BRADY, J. F. 2001 Many-body effects and matrix inversion in low-Reynolds-number hydrodynamics. *Phys. Fluids* **13**, 350–353.
- JEFFREY, D. J. & ONISHI, Y. 1984 Calculation of the resistance and mobility functions for two unequal rigid spheres in low-Reynolds-number flow. *J. Fluid Mech.* **139**, 261–290.
- LADD, A. J. C. 1990 Hydrodynamic transport coefficients of random dispersions of hard spheres. *J. Chem. Phys.* **93**, 3484–3494.
- LADD, A. J. C. 1994a Numerical simulations of particulate suspensions via a discretized Boltzmann equation. Part 1. Theoretical foundation. *J. Fluid Mech.* **271**, 285–309.
- LADD, A. J. C. 1994b Numerical simulations of particulate suspensions via a discretized Boltzmann equation. Part 2. Numerical results. *J. Fluid Mech.* **271**, 311–339.
- MO, G. & SANGANI, A. S. 1994 A method for computing Stokes flow interactions among spherical objects and its application to suspensions of drops and porous materials. *Phys. Fluids* **6**, 1637–1652.
- MORRIS, J. F., JEFFREY, J. D. & BRADY, J. F. 1993 The pressure moments for two rigid spheres in low-Reynolds-number flow. *Phys. Fluids A* **5**, 2317–2325.
- NUNAN, K. C. & KELLER, J. B. 1984 Effective viscosity of a periodic suspension. *J. Fluid Mech.* **142**, 269–287.
- OTTEWILL, R. H. & WILLIAMS, N. S. S. 1987 Study of particle motion in concentrated dispersions by tracer diffusion. *Nature* **325**, 232–234.
- PETERSEN, H. G. 1995 Accuracy and efficiency of the particle mesh Ewald method. *J. Chem. Phys.* **103**, 3668–3679.
- PHILLIPS, R. J., BRADY, J. F. & BOSSIS, G. 1988 Hydrodynamic transport properties of hard-sphere dispersions. I. Suspensions of freely mobile particles. *Phys. Fluids* **31**, 3462–3472.
- PUSEY, P. N. & VAN MEGEN, W. 1983 Measurement of the short-time self-mobility of particles in concentrated particle dispersions. *J. Phys. Paris* **44**, 285–291.

- RINTOUL, M. D. & TORQUATO, S. 1996a Metastability and crystallization in hard-sphere systems. *Phys. Rev. Lett.* **77**, 4198–4201.
- RINTOUL, M. D. & TORQUATO, S. 1996b Computer simulations of dense hard-sphere systems. *J. Chem. Phys.* **105**, 9258–9265.
- SAFFMAN, P. G. 1973 On the settling speed of free and fixed suspensions. *Stud. Appl. Maths.* **52**, 115–127.
- SANGANI, A. S. & MO, G. 1996 An $O(N)$ algorithm for Stokes and Laplace interactions of particles. *Phys. Fluids* **8**, 1990–2010.
- SEGRE, P. N., BEHREND, O. P. & PUSEY, P. N. 1995 Short-time Brownian-motion in colloidal suspensions: Experiment and simulation. *Phys. Rev. E* **52**, 5070–5083.
- SHIKATA, T. & PEARSON, D. S. 1994 Viscoelastic behavior of concentrated spherical suspensions. *J. Rheol.* **38**, 601–616.
- SIEROU, A. 2001 Accelerated Stokesian Dynamics: Development and application to sheared non-Brownian suspensions. PhD thesis, California Institute of Technology.
- SPEEDY, R. J. 1994 On the reproducibility of glasses. *J. Chem. Phys.* **100**, 6684–6691.
- SPEEDY, R. J. 1997 Pressure of the metastable hard-sphere fluid. *J. Phys. Condensed Matter* **9**, 8591–8599.
- TORQUATO, S., TRUSKETT, T. M. & DEBENEDETTI, P. G. 2000 Is random close packing of spheres well defined? *Phys. Rev. Lett.* **84**, 2064–2067.
- TORRES, F. E. & GILBERT, J. R. 1996 Large scale Stokesian Dynamics simulations of non-Brownian suspensions. Unpublished results.
- VAN MEGEN, W., OTTEWILL, R. H., OWENS, S. M. & PUSEY, P. N. 1985 Measurement of the wave-vector dependent diffusion coefficient in concentrated particle dispersions. *J. Chem. Phys.* **82**, 508–515.
- VAN MEGEN, W., UNDERWOOD, S. M. & SNOOK, I. 1986 Tracer diffusion in concentrated colloidal dispersions. *J. Chem. Phys.* **85**, 4065–4072.
- VAN MEGEN, W. & UNDERWOOD, S. M. 1989 Tracer diffusion in concentrated colloidal dispersions. 3. Mean squared displacements and self-diffusion coefficients. *J. Chem. Phys.* **91**, 552–559.
- VANVELUWEN, A., LEKKERKERKER, H. N. W., KRUIF DE, C. G. & VRIJ, A. 1987 Measurement of the short-time self-diffusion coefficient in dilute and concentrated suspensions – influence of direct particle interactions. *J. Chem. Phys.* **87**, 4783–4880.
- VANVELUWEN, A. & LEKKERKERKER, H. N. W. 1988 Non-Gaussian behavior of the displacement statistics of interacting colloidal particles. *Phys. Rev. A* **38**, 3758–3763.
- VAN DER WERFF, J. C., KRUIF DE, C. G., BLOM, C. & MELLEMA, J. 1989 Linear viscoelastic behavior of dense hard-sphere dispersions. *Phys. Rev. A* **39**, 795–807.
- ZICK, A. A. & HOMS, G. M. 1982 Stokes flow through periodic arrays of spheres. *J. Fluid Mech.* **115**, 13–26.
- ZUZOVSKY, M., ADLER, P. M. & BRENNER, H. 1983 Spatially periodic suspensions of convex particles in linear shear flows. III. Dilute arrays of spheres suspended in Newtonian fluids. *Phys. Fluids* **26**, 1714–1723.