



<http://www.diva-portal.org>

This is the published version of a paper published in *Journal of Computational Physics*.

Citation for the original published paper (version of record):

Broms, A., Sandberg, M., Tornberg, A-K. (2023)

A locally corrected multiblob method with hydrodynamically matched grids for the Stokes mobility problem

Journal of Computational Physics, 487: 112172-112172


<https://doi.org/10.1016/j.jcp.2023.112172>

Access to the published version may require subscription.

N.B. When citing this work, cite the original published paper.

Permanent link to this version:

<http://urn.kb.se/resolve?urn=urn:nbn:se:kth:diva-328324>



A locally corrected multiblob method with hydrodynamically matched grids for the Stokes mobility problem

Anna Broms*, Mattias Sandberg, Anna-Karin Tornberg

KTH Royal Institute of Technology, Lindstedtsvägen 25, 114 28 Stockholm, Sweden

ARTICLE INFO

Article history:

Received 5 September 2022

Received in revised form 23 February 2023

Accepted 18 April 2023

Available online 24 April 2023

Keywords:

Stokes flow

Rigid multiblob

Pair-correction

Accuracy

Grid optimisation

Axisymmetry

ABSTRACT

Inexpensive numerical methods are key to enabling simulations of systems of a large number of particles of different shapes in Stokes flow and several approximate methods have been introduced for this purpose. We study the accuracy of the multiblob method for solving the Stokes mobility problem in free space, where the 3D geometry of a particle surface is discretised with spherical blobs and the pair-wise interaction between blobs is described by the RPY-tensor. The paper aims to investigate and improve on the magnitude of the error in the solution velocities of the Stokes mobility problem using a combination of two different techniques: an optimally chosen grid of blobs and a pair-correction inspired by Stokesian dynamics. Different optimisation strategies to determine a grid with a given number of blobs are presented with the aim of matching the hydrodynamic response of a single accurately described ideal particle, alone in the fluid. It is essential to obtain small errors in this self-interaction, as they determine the basic error level in a system of well-separated particles. With an optimised grid, reasonable accuracy can be obtained even with coarse blob-resolutions of the particle surfaces. The error in the self-interaction is however sensitive to the exact choice of grid parameters and simply hand-picking a suitable geometry of blobs can lead to errors several orders of magnitude larger in size. The pair-correction is local and cheap to apply, and reduces the error for moderately separated particles and particles in close proximity. Two different types of geometries are considered: spheres and axisymmetric rods with smooth caps. The error in solutions to mobility problems is quantified for particles of varying inter-particle distances for systems containing a few particles, comparing to an accurate solution based on a second kind BIE-formulation where the quadrature error is controlled by employing quadrature by expansion (QBX).

© 2023 The Author(s). Published by Elsevier Inc. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

1. Introduction

A fluid with immersed rigid particles on the microscale can be modelled by the linear Stokes equations, along with no slip boundary conditions on all particle surfaces. The Stokes equations constitute the low Reynolds number limit of the Navier-Stokes equations, applicable under the assumption that the inertia of the fluid is negligible compared to viscous effects. In such a fluid-particle system, every particle affects every other particle, due to the long range of the hydrodynamic interactions, meaning that the motion of all particles is coupled through the fluid. Examples of fluid-particle systems of this

* Corresponding author.

E-mail address: annabrom@kth.se (A. Broms).

type are found both in biology and industry. Industrial applications are vast in materials science, with studies of flows of polymers, fibrils and fibres, and the forming of gels and crystalline phases [1–5]. Describing the dynamics on the micro level is key to understanding macro level properties of such processes, for instance to manufacture novel materials.

The mobility problem for rigid particles in a Stokesian fluid is that of computing the translational and angular velocities of a set of non-deformable bodies, given assigned net forces and torques such as e.g. gravity or electrostatic forces on every particle. We will focus on the free space problem in 3D, considering no confinements or periodicities for the particles. The mobility problem for a system of p particles can mathematically be stated as $\mathbf{U} = \mathbf{M}\mathbf{F}$, where $\mathbf{F} \in \mathbb{R}^{6p}$ denotes a vector of all (known) net forces and torques, and $\mathbf{U} \in \mathbb{R}^{6p}$ denotes a vector of all translational and angular velocities to be computed, i.e.

$$\mathbf{F}^T = [\mathbf{f}^{1T} \quad \mathbf{t}^{1T} \quad \dots \quad \mathbf{f}^{pT} \quad \mathbf{t}^{pT}], \mathbf{U}^T = [\mathbf{u}^{1T} \quad \boldsymbol{\omega}^{1T} \quad \dots \quad \mathbf{u}^{pT} \quad \boldsymbol{\omega}^{pT}], \quad (1)$$

with \mathbf{f}^i and \mathbf{t}^i the net force and torque on particle i and \mathbf{u}^i and $\boldsymbol{\omega}^i$ the translational and rotational velocities of particle i . The mobility matrix $\mathbf{M} \in \mathbb{R}^{6p} \times \mathbb{R}^{6p}$ is dense, symmetric and depends on the position and orientation of all particles in the system [6]. In addition, the mobility matrix is positive definite, which is a consequence of the dissipative nature of a Stokesian suspension [6]. The inverse of the mobility matrix is termed the *resistance* matrix, \mathbf{R} , and appears in the related resistance problem of computing forces and torques, given assigned particle velocities. We will return to the resistance matrix when discussing techniques for improving on the accuracy for a solution to the mobility problem, as presented in Section 3.

The molecules of the fluid collide with each other and with the immersed particles at a high rate, resulting in Brownian motion for small enough colloidal particles. This stochastic behaviour can be characterised by the overdamped Langevin equation, which is a stochastic differential equation incorporating not only the action of the mobility matrix, but also the action of its square root and divergence [7,8]. The last two quantities have to be approximated from matrix vector products of the form $\mathbf{M}\mathbf{F}$ for some “force” vector \mathbf{F} , using e.g. a Krylov method for approximating the square root [9] and a so-called random finite difference quotient for the divergence term [10]. The large number of such matrix-vector evaluations needed in either a dynamic simulation to determine a trajectory or for drawing statistical conclusions about some physical property of interest (such as e.g. the mean squared displacement, diffusion coefficients or equilibrium distributions) emphasises the need of a method for which the matrix vector product is fast to evaluate, also for systems with many particles. In a Brownian simulation, there are error contributions from several sources: modelling errors in the description of the geometry and in the physical assumptions for the studied particle type, a statistical error in determining physical quantities as averages of a large number of realisations or geometrical configurations, the time discretisation error and an error related to the numerical solution of the (deterministic) Stokes mobility problem. The latter is important to control also in non-Brownian simulations. The aim of this work is to understand the deterministic error related to solving the mobility problem and we assume no Brownian motion.

Specifically, this paper aims at studying the accuracy of a so called *multiblob method*, where the surface of each rigid particle is discretised with spherical blobs and the blobs belonging to one particle are restricted to move as a rigid body, constrained by net forces located at the centre of each blob. The method is simple to implement in its vanilla version and allows for a great flexibility in the particle shapes that can be studied without altering the method as such. In addition, particles of varying shapes and sizes can easily be handled within the same simulation and can be coupled to a fast method for evaluating the action of the blob-blob mobility matrix for different periodicities, which allows for a large number of particles to be studied at a low cost. In [11], large systems of $\mathcal{O}(1000)$ Brownian particles have been studied and in [12], fast approximate solution techniques are discussed, making the complexity close to linear in the number of blobs. The idea of forming larger particles from spheres was first introduced by Kirkwood in 1954 [13] and there is a large collection of work on methods of this type, including but not limited to [14–23]. Multiblob methods have recently been employed for studying Brownian motion and active slip in a number of works by Usabiaga and coauthors [24,12,11,25,26], and by Brosseau and coauthors [27,28], but also in connection to Stokesian dynamics, by the group of Swan [29–31]. A similar technique is employed in [32], where springs constrain the blobs to move as one body. Despite the multiblob method being approximate in its very nature, we present a strategy to understand, control and improve on its accuracy.

Several accurate methods exist for solving the Stokes mobility problem, among which boundary integral equation (BIE) methods form an important class [33–35]. Integral equations have successfully been employed for various particle shapes and various domains in 3D (including confinements [33] and/or periodicities [34]). A special quadrature method has to be used for accurate numerical treatment of the singularities and near-singularities appearing in any integral equation formulation for evaluation on, or close to, a boundary; one example of such quadrature for a double layer formulation is quadrature by expansion (QBX) developed by af Klinteberg & Tornberg [34] and Bagge & Tornberg [33]. This solver, based on a second kind integral equation formulation, allows for accurate computations and has been used to evaluate the Stokes flow for $\mathcal{O}(100)$ closely interacting particles. However, to do so repeatedly for many time steps and/or statistical realisations would be unfeasibly slow, and in such situations, cheaper and hence less accurate methods must be used. In this paper, the QBX-based solver will serve two purposes: provide a reference solution when studying the accuracy of the multiblob method and be used to construct precomputable corrections to the multiblob method that are inexpensive to apply.

Multiple families of methods exist for computing the hydrodynamic interactions in a particle suspension, reviewed by Maxey in [36]. The multiblob method can be viewed as a regularisation of a first kind integral equation [12]. A similar

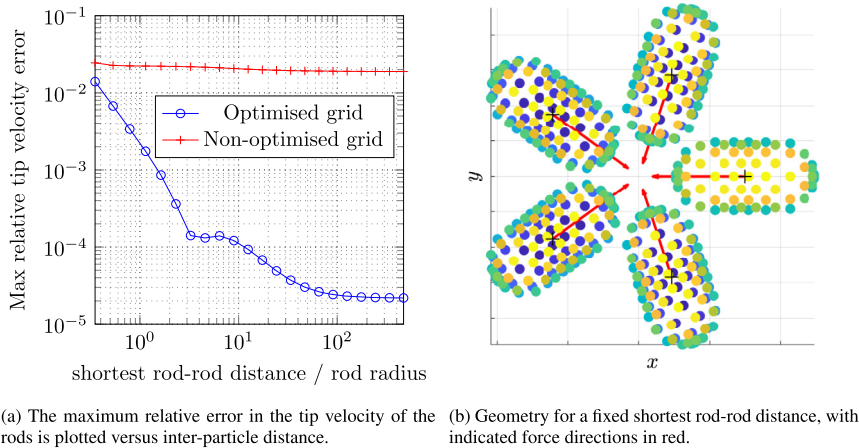


Fig. 1. Numerical example with five rods of aspect ratio $L/R = 4$ forming a circle, comparing two discretisations with the same number of blobs. We gain up to three orders of magnitude in accuracy at larger distances by using the optimised grid for which the multiblob particle matches the hydrodynamic response of the ideal particle. The non-optimised grid is obtained from the optimised grid parameters by adding a perturbation of 1%. For a colour version of this and other figures in the paper, the reader is referred to the web version of this article.

regularisation idea is considered in the method of regularised Stokeslets by Cortez et al. [37].¹ Another example of an approximate method is presented in the large collection of papers related to Stokesian Dynamics [38–40,29,41,31], first introduced in 1987 by Brady, Durlofsky and Bossis in [38], where a near-field correction for each pair of spherical (or spheroidal) particles in close proximity is added to the global far-field resistance matrix, allowing for good approximations for very close particles and for widely separated particles. The near-field correction is constructed from lubrication expressions first presented by Jeffery & Onishi [42] and also listed in [6]. However, the accuracy of Stokesian dynamics is worse for moderate particle separations than for very large or small separations [43], a potential reason being that the lubrication expressions are dominant only for closely interacting pairs – for moderately separated particles, it becomes evident that the additivity assumption of the resistance matrix does not hold [44]. A potential cure to this problem is presented by Lefebvre-Lepot et al. [44]. In their paper, the hydrodynamic interaction is described using a multipole expansion with lubrication corrections only for relative motion within a closely interacting particle pair, which however requires an evaluation of the resulting velocity field for all particles not in the pair, relying on a multivariate interpolant that is not straight-forward to compute cheaply. Despite the possible drawbacks in Stokesian Dynamics, the multiblob method is coupled to that type of pair-corrections in [31]. We will here further develop this idea and study the accuracy of the coupling. In the results section, 4, we show numerically that in this setting, the additivity assumption is not limiting the accuracy of the method. A reason for this is further discussed in the conclusions in Section 6, and in Appendix D.

In Section 1.1, the details of the multiblob method are presented. In Section 2, we describe how an optimisation problem can be solved to closely match the hydrodynamic response of a multiblob particle to that of a chosen particle type that we would like to model. We will talk about this hydrodynamic response as the self-interaction. The particle geometry and discretisations with blobs are presented, emphasising that the multiblob particle is a model of an ideal particle with a certain geometric extension. By optimising the particle grid, we can obtain much smaller errors in the self-interaction, which will set the basic error level also in a multi-particle case. To display the importance of using such a matching technique, we consider the following numerical example: rods are placed in a circle centred at the origin in the xy -plane with rod tips pointing towards the centre of the circle. Each rod is affected by a unit force towards the origin and a unit torque in the z -direction. We vary the circle radius and hence also the distance between rods. The relative error in the velocity at the tip of the rods, visualised in Fig. 1, is determined compared to a BIE reference solution. If an optimised grid is used, we gain up to three orders of magnitude in accuracy compared to using a grid with only slightly perturbed parameters. For closely interacting particles, note that there is a need to correct for interaction errors, also for the optimised grid. The specific parameters used in this example are presented in Table B.14 in Appendix B, and in the manuscript we carefully describe why two seemingly similar discretisations of a particle can yield fundamentally different results in terms of accuracy.

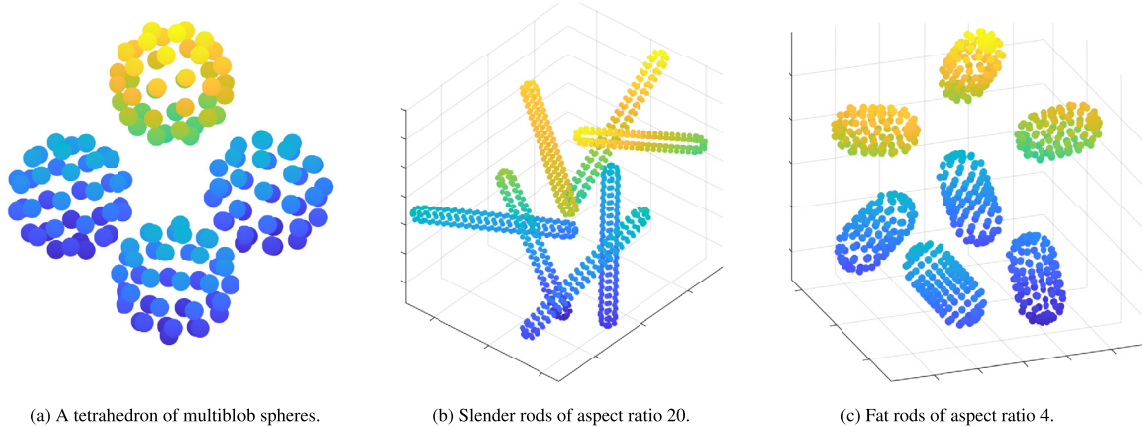
The optimisation technique is fit for use when the mobility matrix is known for a single particle from an accurate method or analytical expressions. In Section 3, we introduce the technique to correct for pair-interactions, cheap to apply to large systems of multiblob particles. In Section 4, numerical results are presented for spheres and rods for systems consisting of a small number of particles at varying inter-particle distances, for which comparisons can be made to results obtained with an accurate solver. The interplay between the self-interaction error and a pair-interaction error and how they can be improved will be discussed. We present the performance of what we refer to as the original multiblob method, with an

¹ A limitation of that method is however that positive-definiteness of the mobility matrix is not naturally obtained due to the choice of regularisation, in contrast to in the multiblob method.

Table 1

List of commonly used symbols and notation and page references to where they are first introduced.

s	Characteristic grid spacing, p. 4
a_h	Hydrodynamic radius of a blob, p. 4
R	Radius of an ideal particle, p. 7
L	Length of an ideal particle, p. 10
R_g	Geometric radius of a multiblob sphere or rod, p. 6
L_g	Geometric length of a multiblob rod, p. 6
\hat{R}_h^{trans}	Computed translational radius of a multiblob sphere, p. 7
\hat{R}_h^{rot}	Computed rotational radius of a multiblob sphere, p. 7
$\xi_r^{\parallel}, \xi_t^{\perp}, \xi_r^{\perp}, \xi_t^{\parallel}$	Resistance coefficients, p. 6
$\hat{\xi}_t^{\parallel}, \hat{\xi}_t^{\perp}, \hat{\xi}_r^{\parallel}, \hat{\xi}_r^{\perp}$	Approximate resistance coefficients computed with the multiblob method, p. 6
n_b	The number of blobs on one particle, p. 4
$n_{\text{cyl}}, n_{\text{cap}}, n_{\varphi}$	Discretisation parameters for a rod-like particle, p. 10
\mathbf{N}	The RPY-tensor, p. 4
\mathbf{M}	System mobility matrix relating given particle forces and torques and the computed particle velocities, p. 1
\mathbf{R}	System resistance matrix, $\mathbf{R} = \mathbf{M}^{-1}$, p. 1
λ	Vector of forces on blobs, p. 4
\mathbf{u}^i	Translational velocity of particle i , p. 1
ω^i	Angular velocity of particle i , p. 1
\mathbf{v}^i	Velocity at the tip of particle i , p. 20
\mathbf{f}^i	Net force on particle i , p. 1
\mathbf{t}^i	Net torque on particle i , p. 1
δ	Smallest particle-particle distance, p. 20
α	Relation between particle force and torque magnitudes in numerical experiments, p. 20
rt-grid	Optimised blob discretisation where errors in translational and rotational mobility coefficients are minimised simultaneously, p. 10
t-grid	Optimised blob discretisation where translational mobility coefficients are matched, p. 14
r-grid	Optimised blob discretisation where rotational mobility coefficients are matched, p. 14
combined solve	Solving for \mathbf{u}^i with the r-grid and ω^i with the t-grid, p. 14

**Fig. 2.** Example multiblob geometries.

optimised grid matching both translational and rotational properties of the ideal particle. This is done for particles of various resolutions, and we also display the improved accuracy when applying pair-corrections. When a coarse discretisation is used, two solves based on two differently matched grids can also be combined for improved accuracy, with one grid optimised for translational properties and one for rotational.

1.1. The multiblob method

We will use the same description of the multiblob method as in the works by Usabiaga, Donev and coauthors [24,11,12,25], by Brosseau et al. in [27,28] and by Swan and coauthors [41,30]. See especially the work by Usabiaga et al. [12], which presents the method in detail and also investigates the accuracy of the multiblob method in its deterministic setting for spheres in free space. A list of commonly used notation in this paper is collected for reference in Table 1. Each particle in a suspension is described as a set of spheres or *blobs* distributed on a surface. A few example geometries are visualised in Fig. 2. Blobs can be distributed on a surface in a magnitude of ways, as will be discussed in Section 2. Given the discretisation points, at which blobs are to be centred, we define the characteristic grid spacing s as the distance between

discretisation points in the cross-section of the particle where the discretisation grid is the coarsest. To be more specific, for an axisymmetric particle, with C a set of cross-sections perpendicular to the axis and with blob i centred at \mathbf{b}_i , let

$$s = \max_C \left(\min_{\{b_j, b_i\} \in C, i \neq j} \|\mathbf{b}_i - \mathbf{b}_j\| \right). \quad (2)$$

The blobs are associated with their *hydrodynamic* radius a_h , which serves as an effective model for how a blob interacts hydrodynamically with other blobs. Blobs with radius a_h may overlap or be separated by a gap on the particle surface. The ratio between the hydrodynamic blob radius and the characteristic blob spacing is related to the extent of these overlaps or gaps. This parameter, a_h/s , has to be chosen and is discussed in terms of the blob-blob spacing in Chapter 4.1 of [12]. In Section 2, we solve an optimisation problem to determine a_h , given s . The effect of this optimal choice is discussed in Section 4.

The hydrodynamic interaction between particles is accounted for by first computing the mobility for all the blobs in the system, interacting in a *pair-wise* manner only. The blob-mobilities are described in its most reduced and simplified form through the Rotne-Prager-Yamakawa tensor (RPY) [12,9], where the block \mathbf{N}_{ij} describes the motion on blob i resulting from a given force on blob j , governed by the far field approximation

$$\mathbf{N}_{ij} \approx \eta^{-1} \left(\mathbf{I} + \frac{1}{6} a_h^2 \nabla_{\mathbf{x}}^2 \right) \left(\mathbf{I} + \frac{1}{6} a_h^2 \nabla_{\mathbf{y}}^2 \right) \mathbb{G}(\mathbf{r}_{ij}), \quad (3)$$

with η the viscosity and

$$\mathbb{G}(\mathbf{r}) = \frac{1}{8\pi r} \left(\mathbf{I} + \frac{\mathbf{r} \otimes \mathbf{r}}{r^2} \right) \quad (4)$$

the Stokeslet, with $r = |\mathbf{r}|$ [45]. The RPY-tensor, corrected for overlapping blobs so that the resulting mobility \mathbf{N}_{ij} is positive-definite, takes the form [12,45]

$$\mathbf{N}_{ij} = \frac{1}{6\pi\eta a_h} \begin{cases} C_1(r_{ij})\mathbf{I} + C_2(r_{ij})(\mathbf{r}_{ij} \otimes \mathbf{r}_{ij})/r_{ij}^2, & r_{ij} > 2a_h, \\ C_3(r_{ij})\mathbf{I} + C_4(r_{ij})(\mathbf{r}_{ij} \otimes \mathbf{r}_{ij})/r_{ij}^2, & r_{ij} \leq 2a_h, \end{cases} \quad (5)$$

with

$$\begin{aligned} C_1(r) &= \frac{3a_h}{4r} + \frac{a_h^3}{2r^3}, & C_2(r) &= \frac{3a_h}{4r} - \frac{3a_h^3}{2r^3}, \\ C_3(r) &= 1 - \frac{9r}{32a_h}, & C_4(r) &= \frac{3r}{32a_h}, \end{aligned} \quad (6)$$

where $\mathbf{r}_{ij} = \mathbf{b}_i - \mathbf{b}_j$ is the centre-centre vector for two blobs i and j . The diagonal blocks simply reduce to $(6\pi\eta a_h)^{-1} \mathbf{I}$, i.e. the well-known translational part of the mobility matrix for a single sphere. Note that the RPY-tensor is a good approximation to the translation-translation coupling between two blobs when the separation between the blobs is sufficiently large.

The blobs belonging to one particle are restricted to move together as a rigid body, which is assured by applying a net force, $\lambda_i^l \in \mathbb{R}^3$, to every blob centre \mathbf{b}_i^l on particle l . The forces on the blobs belonging to one particle are summed to yield the net force, \mathbf{f}^l , and torque, \mathbf{t}^l , on the particle. If we let B^l be the indices of the set of n_b blobs belonging to particle l with centre coordinate \mathbf{c}^l , let the blobs be centred in $\{\mathbf{b}_i^l\}_{i \in B^l}$ and impose no-slip boundary conditions on the multiblob particles, we obtain the set of equations for particle l

$$\begin{aligned} \sum_j \mathbf{N}_{ij}^l \lambda_j^l &= \mathbf{u}^l + \boldsymbol{\omega}^l \times (\mathbf{b}_i^l - \mathbf{c}^l) \quad \text{for all } i \in B^l, \\ \sum_{i \in B^l} \lambda_i^l &= \mathbf{f}^l, \\ \sum_{i \in B^l} (\mathbf{b}_i^l - \mathbf{c}^l) \times \lambda_i^l &= \mathbf{t}^l, \end{aligned} \quad (7)$$

to be solved for the unknown velocity pair \mathbf{u}^l and $\boldsymbol{\omega}^l$ and blob force vector λ^l . The formulation can be motivated as being the regularised discretisation of the first-kind integral equation [12]

$$\mathbf{v}(\mathbf{x}) = \mathbf{u} + \boldsymbol{\omega} \times \mathbf{x} = \eta^{-1} \int_{\Gamma} \mathbb{G}(\mathbf{x} - \mathbf{y}) \boldsymbol{\psi}(\mathbf{y}) d\mathbf{y} \quad \text{for all } \mathbf{x} \in \Gamma, \quad (8)$$

where Γ is the particle boundary, and ψ is an unknown single layer density representing traction, being a continuous analogue of λ .

Defining the matrix $\mathbf{K}^l \in \mathbb{R}^{3n_b} \times \mathbb{R}^6$ from the centre coordinates of the blobs and the particle as

$$(\mathbf{K}^l \mathbf{U}^l)_I = \mathbf{u}^l + \boldsymbol{\omega}^l \times (\mathbf{b}_i - \mathbf{c}^l), \quad (9)$$

with $I = \{3(i-1) + k\}_{k=1}^3$, the system in (7) can be written on the form

$$\begin{bmatrix} \mathbf{N}^{ll} & -\mathbf{K}^l \\ -\mathbf{K}^{lT} & \mathbf{0} \end{bmatrix} \begin{bmatrix} \lambda^l \\ \mathbf{U}^l \end{bmatrix} = \begin{bmatrix} \mathbf{0} \\ -\mathbf{F}^l \end{bmatrix} \quad (10)$$

with $\mathbf{F}^{lT} = [\mathbf{f}^{lT} \mathbf{t}^{lT}]$ and $\mathbf{U}^{lT} = [\mathbf{u}^{lT} \boldsymbol{\omega}^{lT}]$ (ignoring all other particles in the system). We can now eliminate λ^l and solve for \mathbf{U}^l from

$$(\mathbf{K}^{lT} (\mathbf{N}^{ll})^{-1} \mathbf{K}^l) \mathbf{U}^l = \mathbf{F}^l. \quad (11)$$

Note that this allows us to define the single particle mobility matrix as $\mathbf{M}^l = (\mathbf{K}^{lT} (\mathbf{N}^{ll})^{-1} \mathbf{K}^l)^{-1}$ (and its inverse, the resistance matrix \mathbf{R}^l).² For a suspension of multiple particles, we obtain a system of a similar form:

$$\begin{bmatrix} \mathbf{N} & -\mathbf{K} \\ -\mathbf{K}^T & \mathbf{0} \end{bmatrix} \begin{bmatrix} \lambda \\ \mathbf{U} \end{bmatrix} = \begin{bmatrix} \mathbf{0} \\ -\mathbf{F} \end{bmatrix}, \text{ with } \mathbf{K} = \begin{bmatrix} \mathbf{K}^1 & \mathbf{0} & \dots & \mathbf{0} \\ \mathbf{0} & \mathbf{K}^2 & \mathbf{0} & \dots \\ \vdots & \vdots & \ddots & \mathbf{0} \\ \mathbf{0} & \dots & \mathbf{0} & \mathbf{K}^p \end{bmatrix} \in \mathbb{R}^{3n_b p} \times \mathbb{R}^{6p}, \quad (12)$$

where \mathbf{N}^{ll} now is a block on the diagonal of the larger global blob-blob matrix \mathbf{N} and λ contains the forces on all blobs. The mobility matrix for the system is then given by $\mathbf{M} = (\mathbf{K}^T \mathbf{N}^{-1} \mathbf{K})^{-1}$.

In general, this is not how we would solve the mobility problem, and there are efficient preconditioning strategies for solving the system in (12) [12,41]. However, when comparing the solution of the mobility problem for a small number of particles to that of a more accurate method, it is a realistic option to actually compute mobility matrices if the computational cost for doing so is manageable, and apply a large number of right hand side force/torque vectors.

Note that if the particles are subject to a background flow, the upper block of the right hand side vector in (12) would be modified to contain a vector of flow velocities at the blob centres. This would hence not affect the form of the mobility matrix and in the remainder of this work, we assume that particles are immersed in a quiescent fluid.

2. Matching the multiblob grid

Discretising a particle surface with blobs introduces an approximation on several levels. We will quantify the error in the mobility of a multiblob particle by comparing to a well-resolved model where a BIE-solver with QBX is employed for two different particle shapes: the sphere and the rod with smooth caps. Of importance is that a multiblob particle is viewed as a *model* of what we refer to as an *ideal* particle of a certain type. To match with the ideal particle, the centres of the blobs should for good results be placed on a surface interior to that of the ideal particles. We refer to this interior surface as the *geometric* surface. There are several choices to be made, and the combination of these will determine how well the hydrodynamic “response” of the multiblob particle will match that of the ideal particle. Hence, given an ideal particle, we need to consider the questions

1. How should the geometric surface be chosen?
2. What is the optimal placement of blobs (given some restriction on the number of blobs)?
3. What is the optimal hydrodynamic radius of the blobs (the parameter a_h in the RPY-tensor in (5))?

If the blob discretisation of a particle surface is coarse, it is reasonable to believe that the hydrodynamic response will be different from that of the ideal particle. One option is to view the blobs as quadrature nodes on the particle surface, which intuitively is reasonable in the limit with many blobs covering the surface. If the blobs are large and the resolution is coarse, however, the geometry of the multiblob particle is far from the geometry of the particle that we would like to model, and the blobs popping out from the particle surface will make the fluid “interpret” the particle as larger. By introducing an offset

² See the discussion of p. 229 of [12] on the invertibility of \mathbf{N} and \mathbf{M} . Special concerns are particles consisting of a single row of blobs or with infinitely many blobs covering its surface, but none of these settings will be considered here and \mathbf{M} will always be invertible for the particle types studied in the paper.

from the surface, such that blobs are placed centred on a geometric surface close to the boundary but in the interior of the ideal particle, the hydrodynamic response of a coarse multiblob model will be closer to that of the ideal particle.³

To see specifically how the geometry of the multiblob particle should be chosen, we start with a sphere in Section 2.1 and later use a similar strategy also for the rod geometry, in Sections 2.2–2.3. The geometric surface for a sphere will be a sphere with the geometric radius R_g . For an axisymmetric particle, such as a rod, the geometric surface is defined using both a geometric radius R_g and a geometric length L_g .

Given a strategy for how to place n_b blobs on a geometric surface, we have to select R_g , L_g and the hydrodynamic radius a_h . This will be done through an optimisation procedure, where we seek to match the mobility matrix for a single particle, alone in a fluid. The reason is that in a multi-particle suspension where the particles are widely separated, this *self-interaction* will be dominating. Said differently, independent of how dilute the suspension is, the error level will never be lower than the error in the self-interaction. When analytical formulas are not available, an accurate reference mobility matrix for a single particle can be obtained from an accurate numerical method such as the BIE-solver with QBX [33,34]. We will in this paper consider axisymmetric particles, for which the resistance matrix of a single particle takes a particularly simple form. A resistance matrix can generally be written as

$$\mathbf{R} = \begin{bmatrix} \mathbf{R}^{UF} & \mathbf{R}^{\Omega F} \\ \mathbf{R}^{UT} & \mathbf{R}^{\Omega T} \end{bmatrix}, \quad (13)$$

where the four blocks represent the coupling between the assigned translational and angular velocities, indicated by U and Ω , and the induced net force and torque, denoted by F and T . For a single particle with symmetry axis described by the unit vector \mathbf{s} in an unbounded domain, $\mathbf{R}^{\Omega F} = \mathbf{R}^{UT} = \mathbf{0}$, $\mathbf{R}^{UF} = \xi_t^{\parallel}(\mathbf{s}\mathbf{s}^T) + \xi_t^{\perp}(\mathbf{I} - \mathbf{s}\mathbf{s}^T)$ and $\mathbf{R}^{\Omega T} = \xi_r^{\parallel}(\mathbf{s}\mathbf{s}^T) + \xi_r^{\perp}(\mathbf{I} - \mathbf{s}\mathbf{s}^T)$, with the subscript t denoting a coefficient representing translation, the subscript r denoting rotation and the superscripts \parallel and \perp representing motion parallel or perpendicular to the axis of symmetry [8]. The four positive coefficients ξ_t^{\parallel} , ξ_t^{\perp} , ξ_r^{\parallel} and ξ_r^{\perp} are specific for a given particle shape and are known analytically for some simple geometries, such as spheres, ellipsoids and infinitely long rods. The mobility matrix, \mathbf{M} , is the inverse of the resistance matrix, relating applied forces and torques to computed velocities. Its blocks are given by $\mathbf{M}^{F\Omega} = \mathbf{M}^{TU} = \mathbf{0}$, $\mathbf{M}^{FU} = (1/\xi_t^{\parallel})(\mathbf{s}\mathbf{s}^T) + (1/\xi_t^{\perp})(\mathbf{I} - \mathbf{s}\mathbf{s}^T)$ and $\mathbf{M}^{T\Omega} = (1/\xi_r^{\parallel})(\mathbf{s}\mathbf{s}^T) + (1/\xi_r^{\perp})(\mathbf{I} - \mathbf{s}\mathbf{s}^T)$.

For what follows, let $\{\xi_t^{\parallel}, \xi_t^{\perp}, \xi_r^{\parallel}, \xi_r^{\perp}\}$ be the set of coefficients determined analytically or with an accurate method and let $\{\hat{\xi}_t^{\parallel}, \hat{\xi}_t^{\perp}, \hat{\xi}_r^{\parallel}, \hat{\xi}_r^{\perp}\}$ be another set of coefficients deduced from the approximate multiblob method.

2.1. Spheres

A sphere constitutes a simple special case of an axisymmetric particle. It is well-known [6] that the non-zero blocks of the resistance matrix for a single sphere of radius R in a fluid with viscosity η modelled in free space are given by $\mathbf{R}^{UF} = 6\pi\eta R\mathbf{I}$ and $\mathbf{R}^{\Omega T} = 8\pi\eta R^3\mathbf{I}$; this is a consequence of the fact that the translational velocity of a single sphere affected by a net force \mathbf{f} theoretically is given by $\mathbf{f}/(6\pi\eta R)$ and, similarly, that the angular velocity of a single sphere affected by a net torque \mathbf{t} is $\mathbf{t}/(8\pi\eta R^3)$ (these relations are often referred to as Stokes' first and second law). The Stokes' laws can be used to numerically identify the hydrodynamic radius \hat{R}_h of a spherical multiblob particle, given a computed translational and angular velocity, \mathbf{u} and $\boldsymbol{\omega}$, corresponding to an assigned force and an assigned torque respectively. The effective hydrodynamic radius will however in general not automatically be the same for the translational and angular motion, and we denote these two computed hydrodynamic radii by \hat{R}_h^{trans} and \hat{R}_h^{rot} . Relating to the general resistance expressions, the exact coefficients are given by $\xi_t^{\parallel} = \xi_t^{\perp} = 6\pi\eta R$ and $\xi_r^{\parallel} = \xi_r^{\perp} = 8\pi\eta R^3$, while the computed coefficients for the multiblob particle are $\hat{\xi}_t^{\parallel} = \hat{\xi}_t^{\perp} = 6\pi\eta\hat{R}_h^{\text{trans}}$ and $\hat{\xi}_r^{\parallel} = \hat{\xi}_r^{\perp} = 8\pi\eta(\hat{R}_h^{\text{rot}})^3$.

The sphere geometry can be discretised with blobs from uniform subdivisions of an icosahedron projected onto the sphere geometry,⁴ as is illustrated for three consecutive refinements in Fig. 3, resulting in models with $10 \cdot 4^{k-1} + 2$ blobs in subdivision k . This is the same sphere geometry as used for multiblob spheres in [12,47]. The blobs are placed on spheres of geometric radius R_g not necessarily equal to the hydrodynamic radius R . We choose the characteristic grid spacing, s , to be the minimum spacing between grid nodes for each sphere resolution (this corresponds to the definition given in (2)) and define the hydrodynamic radius a_h relative to s . An illustration of the parameters R , R_g and a_h is displayed in Fig. 4a. We want to select R_g and a_h to match the hydrodynamic response of the multiblob particle to that of the ideal particles as closely as possible.

The hydrodynamic particle radii \hat{R}_h^{trans} and \hat{R}_h^{rot} depend both on R_g and a_h . In general, $\hat{R}_h^{\text{trans}} \neq \hat{R}_h^{\text{rot}}$, if R_g and a_h are not chosen carefully. In Figs. 5–6, the dependence on the parameters is visualised. We can however determine R_g and a_h such

³ The idea is similar to that in the immersed boundary method in [46], where there is also a discrepancy between the geometric and hydrodynamic surface.

⁴ Code for generating these subdivisions is taken from (<https://www.mathworks.com/matlabcentral/fileexchange/50105-icosphere>), MATLAB Central File Exchange. Retrieved March 10, 2021.

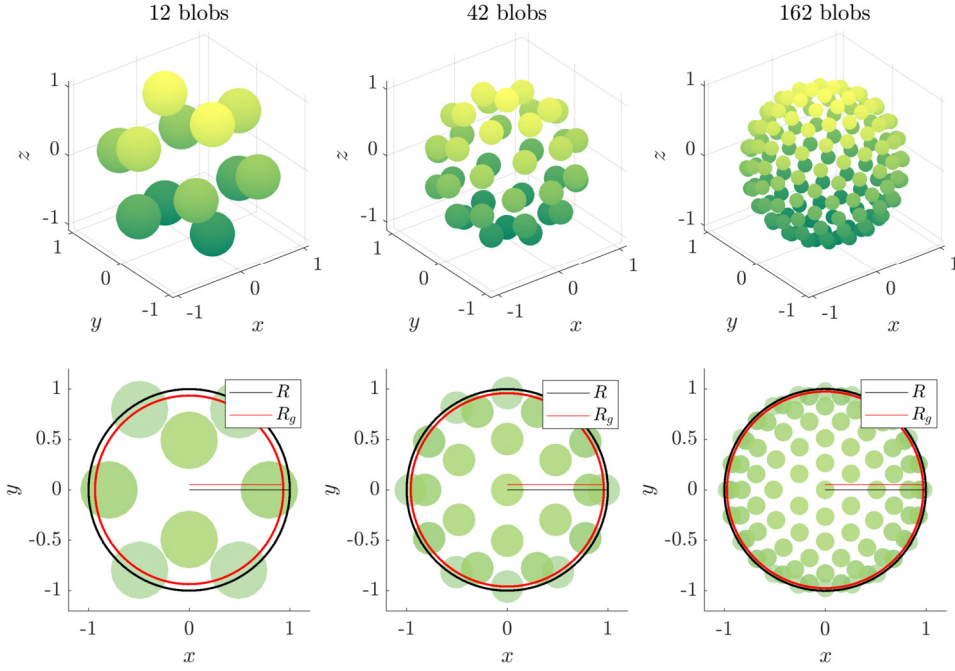


Fig. 3. Spherical particles of unit hydrodynamic radius $R = 1$ modelled with multiblobs. The geometric radius R_g sets the surface on which to place the blob centres. A sphere is constructed from uniform subdivisions of an icosahedron (the first three subdivisions are depicted from left to right). In the bottom panel, a 2D projection is displayed for each discretisation, where the geometric and hydrodynamic radius are indicated. We have solved for the optimal R_g and blob spacing a_h/s as presented in Table 2.




Fig. 4. Multiblob particles discretised with blobs placed on a geometric surface chosen with an optimisation procedure to hydrodynamically match the corresponding ideal particle.

that $\hat{R}_h^{\text{trans}} \approx \hat{R}_h^{\text{rot}} \approx R$. For this purpose, we minimise the relative error in the mobility coefficients, that is

$$\min_{R_g, a_h} \max \left\{ R \left| \frac{1}{R} - \frac{1}{\hat{R}_h^{\text{trans}}(R_g, a_h)} \right|, R^3 \left| \frac{1}{R^3} - \frac{1}{(\hat{R}_h^{\text{rot}}(R_g, a_h))^3} \right| \right\}. \quad (14)$$

The problem in (14) might have multiple local minima and we choose the minimising (R_g, a_h) such that $R_g \approx R$. The reason for this is that the boundary conditions are imposed at the blob centres, as presented in (7) in the description of the multiblob method, which is physically reasonable if $R_g \approx R$. The optimal pairs (R_g, a_h) are presented for four resolutions of the sphere in Table 2, given $R = 1$. We would like to emphasise the importance of this optimisation procedure by inspecting the optimisation landscape in Fig. 6. The relative error in the mobility coefficients is very sensitive to the values of R_g and a_h .

In the work by Usabiaga et al. [12], R_g is chosen so that $R = \hat{R}_h^{\text{trans}}$, with the blob radius ratio $a_h/s = 0.5$. The motivation is that the translational coupling is the most long-ranged in the fluid. This is a good choice if we would only care about




Fig. 5. The computed effective translational and rotational radius for a spherical multiblob particle of different resolutions plotted versus a_h/s , given $R_g = 1$. Note that in general $\hat{R}_h^{\text{trans}} \neq \hat{R}_h^{\text{rot}}$ and we need to solve for R_g and a_h/s such that the translational and rotational hydrodynamic radii agree. The optimal R_g and a_h/s given $\hat{R}_h^{\text{trans}} = \hat{R}_h^{\text{rot}} = 1$ are presented in Table 2.




Fig. 6. Optimisation landscape for the problem of determining the sphere grid in (14). The maximum relative mobility coefficient error (\log_{10}) is displayed for each $(R_g, a_h/s)$ with the optimum marked with a red circle for three resolutions.

Table 2

Optimal R_g and a_h/s computed as solutions to (14) for spherical multiblobs of varying resolutions with $R = 1$, presented together with the relative error in each of the mobility coefficients.

Resolution	R_g	a_h/s	$ 1 - (\hat{R}_h^{\text{trans}}(R_g, a_h))^{-1} $	$ 1 - (\hat{R}_h^{\text{rot}}(R_g, a_h))^{-3} $
12 blobs	0.936	0.291	4.94×10^{-9}	1.96×10^{-8}
42 blobs	0.959	0.311	9.23×10^{-9}	1.22×10^{-8}
162 blobs	0.974	0.345	1.97×10^{-8}	1.96×10^{-8}
642 blobs	0.984	0.388	5.19×10^{-9}	9.15×10^{-9}

translational motion. However, if rotational velocities of spheres are to be considered, we have large errors in \hat{R}_h^{rot} as presented in Table 3, resulting in large relative errors in the rotational mobility coefficient $1/(8\pi(\hat{R}_h^{\text{rot}})^3)$. We will show the consequence of such an error in multi-particle simulations in Section 4.1. The blob radius ratio $a_h/s = 0.5$ has previously been the choice in e.g. [20,29], motivated by a view where blobs are seen as tightly packed spheres covering the surface of a particle, but not overlapping.

Table 3

Relative errors in the rotational radius and rotational mobility coefficient using $a_h/s = 0.5$ and R_g chosen so that $\hat{R}_h^{\text{trans}}(R_g, a_h) = 1$. This is the relation chosen between the geometric and hydrodynamic radius of the sphere in [12]. Compare the error levels in the rightmost column to those presented in Table 2, where we have optimised for R_g and a_h/s to minimise the errors in both mobility coefficients (matching both \hat{R}_h^{trans} and \hat{R}_h^{rot}). The error in the rotational mobility coefficient is much larger if the blob radius ratio $a_h/s = 0.5$ is fixed.

Resolution	R_g	$ 1 - (\hat{R}_h^{\text{rot}}(R_g, a_h)) $	$ 1 - (\hat{R}_h^{\text{rot}}(R_g, a_h))^{-3} $
12 blobs	0.792	2.46×10^{-2}	7.77×10^{-2}
42 blobs	0.891	1.79×10^{-2}	5.58×10^{-2}
162 blobs	0.950	5.48×10^{-3}	1.66×10^{-2}
642 blobs	0.977	1.25×10^{-3}	3.77×10^{-3}

Remark 1. We could also compute an effective stress radius for a sphere in shear flow, similarly as we do for the rotational and translational radii, and use also this radius as a parameter that we choose to match for in the optimisation problem. Computing the stress radius is done in [12]. Along a similar note, there might be other effective quantities of the multiblob particle that affects e.g. the rheological properties of a multi-particle suspension. In an application of multiblobs to model physical particles with specific properties, the objective function in (14) can potentially be modified, or possible additional constraints added to account for such properties.

2.2. Rods: the rt-grid

For any axisymmetric particle, there are four mobility coefficients that need to be matched: two for translation, $1/\xi_t^{\parallel}$ and $1/\xi_t^{\perp}$, and two for rotation, $1/\xi_r^{\parallel}$ and $1/\xi_r^{\perp}$. For a rod-like particle with length L and radius R , there are approximate expressions from slender body theory for the translational resistance coefficients, where [48,49]

$$\xi_t^{\parallel} = \frac{2\pi\eta L}{\ln(L/2R)}, \quad \xi_t^{\perp} = 2\xi_t^{\parallel}, \quad (15)$$

however only valid in the limit $L \rightarrow \infty$. For the rotational resistance coefficients, theoretical results mainly focus on rotation perpendicular to the axis of symmetry, as the rotation around the axis of symmetry is assumed negligible for an axisymmetric particle. From slender body theory, the coefficient ξ_r^{\perp} for an infinitely thin rod can be approximated by, [48,49]

$$\xi_r^{\perp} = \frac{\pi\eta L^3}{3\ln(L/2R)}. \quad (16)$$

As these resistance coefficients are inaccurate for any rod of finite aspect ratio, we will use the resistance coefficients $\{\xi_t^{\parallel}, \xi_t^{\perp}, \xi_r^{\parallel}, \xi_r^{\perp}\}$ computed with the BIE-method to determine how the blobs should be placed on the surface of a rod of a given length and radius with smooth caps. In [12], mobility coefficients are extrapolated for a cylinder, using that the convergence to the true mobility coefficients is linear in a_h . In this paper, we employ two different matching techniques, one of which is presented in this section.

Rods of two different aspect ratios are studied: a fat rod with $L/R = 4$ and a slender rod with $L/R = 20$, with the accurate mobility coefficients reported in Table A.12 in Appendix A. Given the length and radius determining the shape of the ideal particle, a geometric radius, R_g , and a geometric length, L_g , are determined for each rod size and discretisation so that the hydrodynamic response of the multiblob particle closely matches that of the BIE-rod. More specifically, the mobility matrix for a single multiblob rod is matched as closely as possible to the mobility matrix for the corresponding BIE-rod. We seek L_g and R_g , along with the hydrodynamic radius of the blobs, a_h , that minimise the maximum relative error in the mobility coefficients. We can phrase this as an optimisation problem on the form

$$\min_{L_g, R_g, a_h} \max \left\{ \xi_r^{\perp} \left| \frac{1}{\xi_r^{\perp}} - \frac{1}{\hat{\xi}_r^{\perp}(L_g, R_g, a_h)} \right|, \xi_r^{\parallel} \left| \frac{1}{\xi_r^{\parallel}} - \frac{1}{\hat{\xi}_r^{\parallel}(L_g, R_g, a_h)} \right|, \right. \\ \left. \xi_t^{\perp} \left| \frac{1}{\xi_t^{\perp}} - \frac{1}{\hat{\xi}_t^{\perp}(L_g, R_g, a_h)} \right|, \xi_t^{\parallel} \left| \frac{1}{\xi_t^{\parallel}} - \frac{1}{\hat{\xi}_t^{\parallel}(L_g, R_g, a_h)} \right| \right\}. \quad (17)$$

To find a reasonable local minimum to this problem, the three variables R_g , L_g and a_h are related to the ideal particle and constrained so that

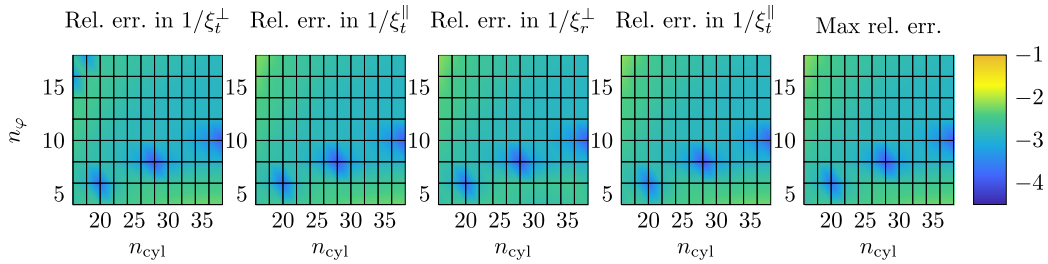


Fig. 7. Relative error (\log_{10}) in the four mobility coefficients, $1/\xi_t^\perp$, $1/\xi_t^\parallel$, $1/\xi_r^\perp$ and $1/\xi_r^\parallel$, for a rod of aspect ratio $L/R = 20$ discretised with $n_{\text{cap}} = 3$ and varying n_{cyl} and n_ϕ at the optimal value of the problem in (17). Here, the cap is sampled with Gauss Legendre nodes in the axial direction and on the cylindrical part of the rod, layers of blobs in the axial direction are aligned.

$$\begin{aligned} L/2 &\leq L_g \leq 1.5L, \\ R/2 &\leq R_g \leq 2R, \\ 0.2 &\leq a_h/s \leq 1. \end{aligned} \tag{18}$$

The problem is solved with `fminimax` in MATLAB with a solver tolerance set to 10^{-11} . The minimisation of the maximum error in (17) ensures that the error level is kept small and approximately equal in all four mobility coefficients. We term the optimised grid the *rt-grid*, as both rotational and translational mobility coefficients are matched.

The geometry and parameterisation of the surface of rods with smooth caps, designed for the BIE-method, are described in detail by Bagge & Tornberg in the appendix of [33]. For the multiblob rods, we use the same parameterisation of the rod geometry. For a multiblob rod of length L_g and radius R_g , three parameters determine its discretisation: The top and bottom cap each occupy a length corresponding to $1.5R_g$ (a choice made for smooth caps for the BIE-rods – the caps are hence not half-spheres) and are discretised in the axial direction by n_{cap} nodes. The middle cylindrical part of the rod is discretised with n_{cyl} equally spaced nodes (in the BIE-method, these are chosen as Gauss-Legendre nodes). Both the cap and the cylindrical middle part of the rod is discretised with n_ϕ equally spaced points in the cross-section of the rod. Thus, a total of $n_b = (2n_{\text{cap}} + n_{\text{cyl}})n_\phi$ blobs discretise the rod. We have experimented with different ways of sampling these parameters: *aligning* the different layers along the axial direction of the rod or *shifting* the layers so that every second layer is aligned and sampling the cap in the axial direction with either equally spaced nodes or Gauss-Legendre nodes in the arc length parameter. Two example rods are displayed in Fig. 4, where the grid is shifted on one rod and aligned on the other. Different strategies of placing the blobs do not result in very large differences in neither the computed geometries (L_g , R_g , a_h) nor the resulting error levels upon solving (17). Hence, only a subset of the results is presented. We have also tried to approximate the rod with a cylinder without caps, to quantify the importance of the caps in the hydrodynamic response (results not reported, but the errors in the mobility coefficients are larger).

The optimisation problem in (17) is solved for a large set of different discretisation triplets $\{n_{\text{cap}}, n_{\text{cyl}}, n_\phi\}$ for each of the two aspect ratios $L/R = 20$ and $L/R = 4$. The parameter n_{cap} is kept moderate to avoid excessive clustering of blobs near the endpoints of the rods, which might cause ill-conditioning of the matrix \mathbf{N} . The results of the optimisation for slender rods of aspect ratio $L/R = 20$ are displayed in Figs. 7–10 and the results for fat rods of aspect ratio $L/R = 4$ are displayed in Figs. 11–13. The relative errors in the four different mobility coefficients are visualised for a specific choice of n_{cap} for the aspect ratio $L/R = 20$ in Fig. 7. Note that the relative error is equal in all four mobility components. The computed length, radius, aspect ratio, blob radius ratio a_h/s , blob radius and total number of blobs for varying n_{cyl} and n_ϕ is visualised in Figs. 8 and 11 for two specific choices of the parameter n_{cap} for the slender and fat rod. These figures illustrate that $L_g \rightarrow L$ and $R_g \rightarrow R$ for well-resolved models of the rod. Figs. 9–10 and 12–13 display the maximum relative error in the mobility coefficients depending on the number of blobs used to discretise the particle, presented in terms of $\{n_{\text{cap}}, n_{\text{cyl}}, n_\phi\}$, for shifted and aligned grids. Note that some certain choices of $\{n_{\text{cap}}, n_{\text{cyl}}, n_\phi\}$ will be considerably more favourable than others, and it is not only the total number of blobs that is important, nor an increased refinement in a certain direction. The distribution of blobs with shifted layers generally constitutes the best choice for both aspect ratios, but for the aligned grid (Figs. 9 and 12), some choices of $\{n_{\text{cap}}, n_{\text{cyl}}, n_\phi\}$ are better than the shifted grid (Figs. 10 and 13).

The error in the solution of a mobility problem with multiple particles will be correlated with the error levels of the mobility coefficients for the *rt-grid*, as we will see in the numerical results section for rods, 4.2, at least if the particles are sufficiently separated. If the error is in one of the translational coefficients, then, the translational velocity error would be limited from below on a level correlated with that coefficient error and similarly for the rotational coefficients and rotational velocity errors. We can expect no errors in the multi-particle tests for rods to be smaller than the errors seen for the single particle mobility coefficients. Reasonable error levels for multi-particle mobility problems are expected if a set of discretisation parameters are chosen so that the maximum relative mobility coefficient error is low for the *rt-grid*. Solving a minmax problem ensures that the error is weighted equally for rotation and translation, parallel and perpendicular to the axis of symmetry. To illustrate the importance of optimising for (L_g, R_g, a_h) , the optimisation landscape is visualised in Fig. 14, where one parameter at the time is fixed at its optimum and the other two are varied. From the figure, it is clear that the error levels in the four mobility coefficients are highly sensitive to the choice of (L_g, R_g, a_h) .

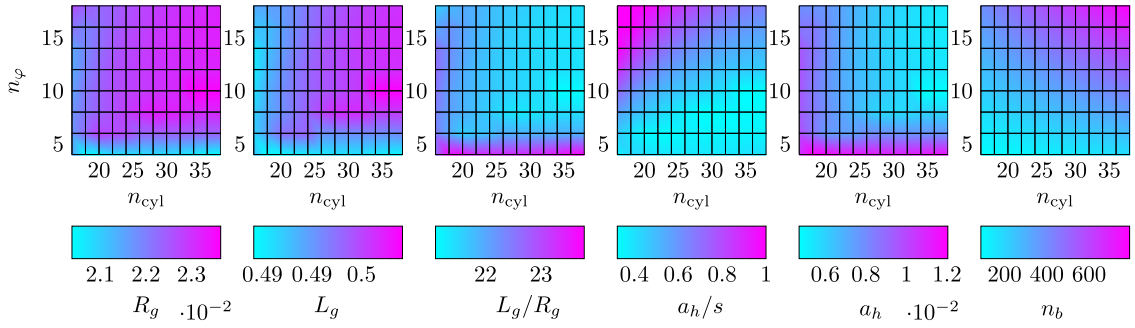


Fig. 8. Same as Fig. 7, but showing computed geometric radii, R_g , geometric lengths, L_g , aspect ratio L_g/R_g , blob scalings, a_h/s , hydrodynamic blob radii, a_h , and number of blobs, n_b , for each grid. The length and radius of this slender rod is $L = 0.5$ and $R = 0.025$. With finer grid discretisation in the axial direction of the particle (larger n_{cyl}), R_g closely matches R and L_g closely matches L .

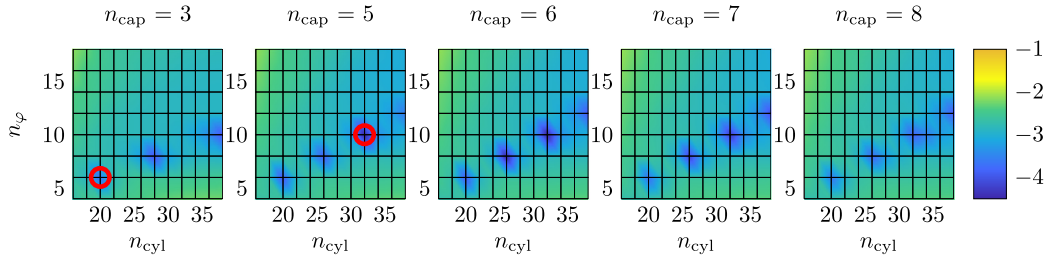


Fig. 9. Maximum relative error (\log_{10}) in the four mobility coefficients for rods of aspect ratio $L/R = 20$ with increasing n_{cap} from left to right, corresponding to the rightmost panel in Fig. 7. The grid is aligned along the axial direction of the rod and Gauss-Legendre nodes are used for the caps. For some relation between the grid spacings in the vertical and horizontal direction along the rod, errors are significantly smaller. Two sets of grid parameters $\{n_{cap}, n_{cyl}, n_\varphi\}$ are marked in red, to be considered for further numerical investigations in Section 4.2.

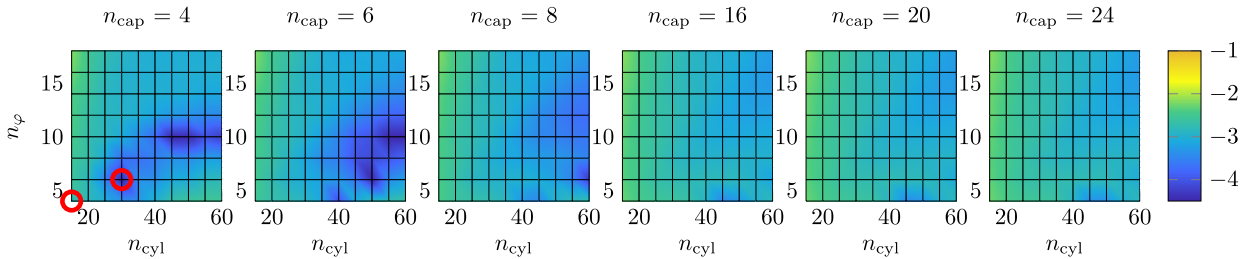


Fig. 10. The same type of investigation as in Fig. 9, but with the grid rotated between each layer so that not all the blobs of all layers are aligned, but only the blobs of every second layer. Errors are generally slightly lower than if the grid is not rotated (compare to Fig. 9). Gauss-Legendre nodes are used for the caps. This choice was found superior to using equally spaced points along the axial direction of the caps or chopping the caps of the rod to avoid clustered blobs near the rod endpoints (results not displayed).

For illustration purposes, we pick a number of discretisation triplets $\{n_{cap}, n_{cyl}, n_\varphi\}$ with varying number of total blobs and varying error levels for each aspect ratio to be used in numerical experiments. These sets are marked in red in Figs. 9 and 10 for the slender rod and in Figs. 12 and 13 for the fat rod. For these particular choices, the single particle error levels are presented in Table 4. Note that the error level is the same in all four mobility coefficients. The corresponding optimised parameters L_g , R_g and a_h/s are displayed in Table 5.

To stress the importance of optimising also for a_h/s , and not only for the geometric surface in terms of (L_g, R_g) , we solve the related problem

$$\min_{L_g, R_g} \max \left\{ \xi_r^\perp \left| \frac{1}{\xi_r^\perp} - \frac{1}{\hat{\xi}_r^\perp(L_g, R_g)} \right|, \xi_r^\parallel \left| \frac{1}{\xi_r^\parallel} - \frac{1}{\hat{\xi}_r^\parallel(L_g, R_g)} \right|, \xi_t^\perp \left| \frac{1}{\xi_t^\perp} - \frac{1}{\hat{\xi}_t^\perp(L_g, R_g)} \right|, \xi_t^\parallel \left| \frac{1}{\xi_t^\parallel} - \frac{1}{\hat{\xi}_t^\parallel(L_g, R_g)} \right| \right\}, \quad (19)$$

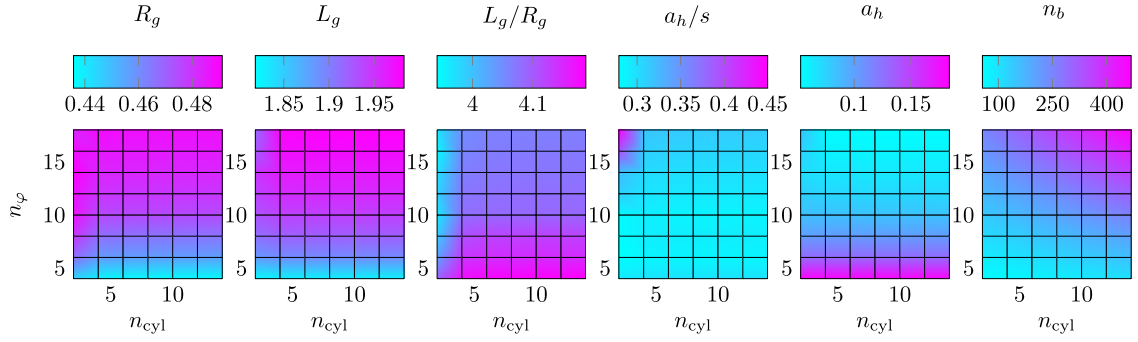


Fig. 11. Solving the optimisation problem (17) for a rod of aspect ratio $L/R = 4$, for which $L = 2$ and $R = 1/2$. The particle is discretised with $n_{\text{cap}} = 6$ and varying n_{cyl} and n_{ϕ} . Displayed are the computed geometric radii, R_g , geometric lengths, L_g , computed aspect ratio L_g/R_g , blob scalings, a_h/s , hydrodynamic blob radii, a_h , and the number of blobs for each grid. With finer grid discretisation in the cross-section of the particle (larger n_{ϕ}), R_g closely matches R and L_g closely matches L .

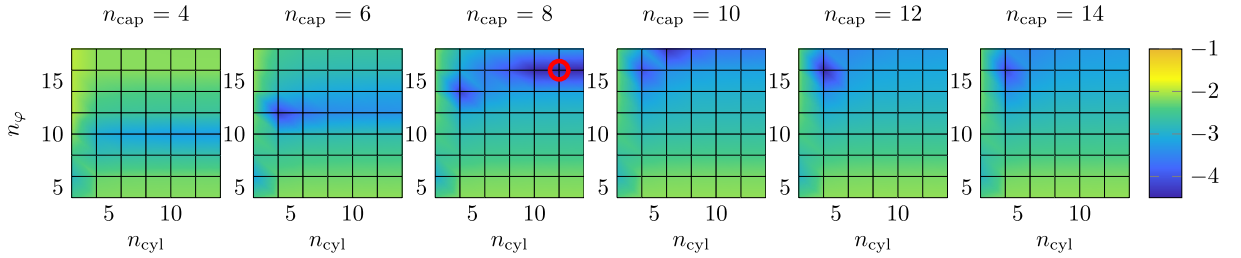


Fig. 12. Maximum relative error in the four mobility coefficients (\log_{10}) for rods of aspect ratio $L/R = 4$ with increasing n_{cap} from left to right, computed as optimal values in (17). Here, all layers of blobs along the axial direction of the rod are aligned and caps are discretised with Gauss-Legendre nodes in the axial direction. The choice $\{n_{\text{cap}}, n_{\text{cyl}}, n_{\phi}\}$ marked in red is considered in numerical experiments in Section 4. The error and corresponding optimised parameters are presented in Tables 4 and 5.

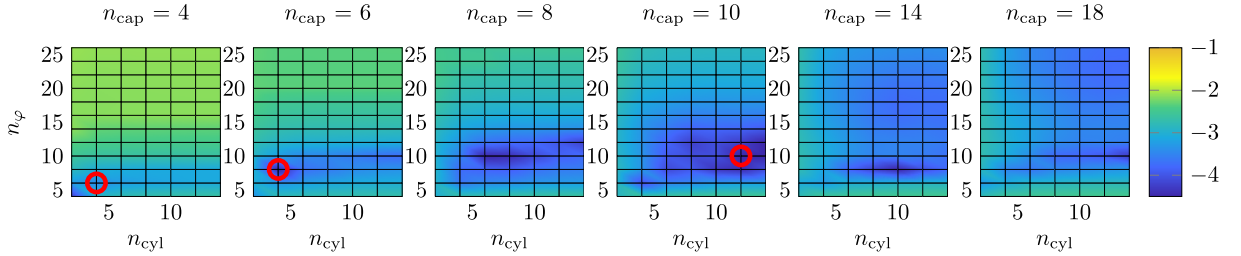


Fig. 13. The same type of investigation as in Fig. 12, but with the grid shifted between each layer of the rod. Results are improved with larger n_{cap} . Three sets of discretisation triplets $\{n_{\text{cap}}, n_{\text{cyl}}, n_{\phi}\}$ are marked in red and are considered in numerical experiments in Section 4.




Fig. 14. Optimisation landscape for the problem in (17) for a rod of aspect ratio $L/R = 4$ and $n_{\text{cap}} = 6$, $n_{\text{cyl}} = 4$, $n_{\phi} = 8$, with one parameter at the time fixed at its optimal value. The optimal value in all three variables ($R_g, L_g, a_h/s$) is marked with a red circle in each of the three plots. The minimum of the maximum relative mobility coefficient error is highly dependent on the precise choice of (L_g, R_g, a_h) . This error level will set a baseline for the error also in a multi-particle simulation.

Table 4

Relative error levels in the mobility coefficients for a single particle for four chosen discretisations, marked with red circles in Figs. 9, 10, 12 and 13. The presented errors are the computed optimal values in (17). The “aligned” column indicates if all blobs along the axial direction of the rod are aligned or shifted for every second layer.

n_{cap}	n_{cyl}	n_{φ}	n_b	aligned	$ 1 - \xi_t^{\parallel}/\hat{\xi}_t^{\parallel} $	$ 1 - \xi_t^{\perp}/\hat{\xi}_t^{\perp} $	$ 1 - \xi_r^{\parallel}/\hat{\xi}_r^{\parallel} $	$ 1 - \xi_r^{\perp}/\hat{\xi}_r^{\perp} $
$L/R = 20$								
4	15	4	92	No	1.73×10^{-3}	1.73×10^{-3}	1.73×10^{-3}	1.73×10^{-3}
3	20	6	156	Yes	1.92×10^{-4}	1.92×10^{-4}	1.92×10^{-4}	1.92×10^{-4}
4	30	6	228	No	2.61×10^{-5}	2.61×10^{-5}	2.61×10^{-5}	2.61×10^{-5}
5	32	10	420	Yes	2.21×10^{-6}	2.21×10^{-6}	2.21×10^{-6}	2.21×10^{-6}
$L/R = 4$								
4	4	6	72	No	3.66×10^{-4}	3.66×10^{-4}	3.66×10^{-4}	3.66×10^{-4}
6	4	8	128	No	2.17×10^{-5}	2.17×10^{-5}	2.17×10^{-5}	2.17×10^{-5}
10	12	10	320	No	3.10×10^{-5}	3.10×10^{-5}	3.10×10^{-5}	3.10×10^{-5}
8	12	16	448	Yes	3.46×10^{-6}	3.47×10^{-6}	3.47×10^{-6}	3.46×10^{-6}

Table 5

Optimised geometry ($L_g, R_g, a_h/s$) upon solving the minmax problem (17), compared to the resulting geometry (L_g^*, R_g^*) from solving (19), with fixed blob radius ratio $a_h^*/s = 0.5$. The geometries correspond to the error levels presented in Table 6. For reference, the slender ideal particle has length $L = 0.5$ and radius $R = 0.025$ and the fat particle has length $L = 2$ and radius $R = 0.5$.

n_{cap}	n_{cyl}	n_{φ}	n_b	aligned	L_g	L_g^*	R_g	R_g^*	a_h/s
$L/R = 20$									
4	4	15	92	No	0.485	0.478	0.0213	0.0191	0.359
3	20	6	156	Yes	0.493	0.484	0.0228	0.0209	0.334
4	30	6	228	No	0.495	0.483	0.0232	0.0208	0.267
5	32	10	420	Yes	0.496	0.491	0.0235	0.0225	0.341
$L/R = 4$									
4	4	6	72	No	1.933	1.762	0.476	0.412	0.222
6	4	8	128	No	1.953	1.820	0.479	0.433	0.221
10	12	10	320	No	1.966	1.858	0.483	0.446	0.201
8	12	16	448	Yes	1.966	1.921	0.483	0.467	0.304

with fixed $a_h/s = 0.5$, using shifted grids along the particle axis. The corresponding maximum relative error levels for different grids are presented in Fig. 15 and Table 6 for the two rod sizes. The errors in the self-interaction are much larger if we fix a_h/s instead of optimising for this ratio. The corresponding (L_g, R_g) is reported in Table 5.

We have experimented with rotations of the multiblob grid about the particle axis when determining the parameters in the rt-grids in this section and found that even if the multiblob rods are not truly axisymmetric, the magnitudes of the mobility coefficients for a single particle are equal to 14 digits with different rotations. In an unbounded fluid with no other particles present, this is not surprising. We can however conclude from these experiments that the known block diagonal form of the mobility matrix for axisymmetric particles holds also for our multiblobs. To ensure that our solutions of the minmax problem (17) are not local minimas, we have also perturbed the initial guesses in all three variables (L_g, R_g, a_h). These perturbations have not been done for all grids, but we are confident to conclude that the convergence to a solution in (17) is not dependent on the initial guesses.

2.3. Rods: the r - and t -grids

Even if we solve the minimisation problem (17) (for what we hereafter refer to as the self-interaction) to determine R_g, L_g and a_h in the rt-grid, the relative errors in the mobility coefficients may not be as small as desired. For a given number of blobs on the particle, the question is if there is any way to obtain smaller errors in the mobility coefficients. In this section, we explore the idea of combining two different discretisations to reduce the errors further.

The idea is to match the grid of the multiblob particle twice: *either* so that translational coefficients are matched, with

$$\xi_t^{\perp} \left| \frac{1}{\xi_t^{\perp}} - \frac{1}{\hat{\xi}_t^{\perp} (L_g^t, R_g^t, a_h^t)} \right| \leq \epsilon \text{ and } \xi_t^{\parallel} \left| \frac{1}{\xi_t^{\parallel}} - \frac{1}{\hat{\xi}_t^{\parallel} (L_g^t, R_g^t, a_h^t)} \right| \leq \epsilon \quad (20)$$

or so that rotational coefficients are matched, with

$$\xi_r^{\perp} \left| \frac{1}{\xi_r^{\perp}} - \frac{1}{\hat{\xi}_r^{\perp} (L_g^r, R_g^r, a_h^r)} \right| \leq \epsilon \text{ and } \xi_r^{\parallel} \left| \frac{1}{\xi_r^{\parallel}} - \frac{1}{\hat{\xi}_r^{\parallel} (L_g^r, R_g^r, a_h^r)} \right| \leq \epsilon, \quad (21)$$

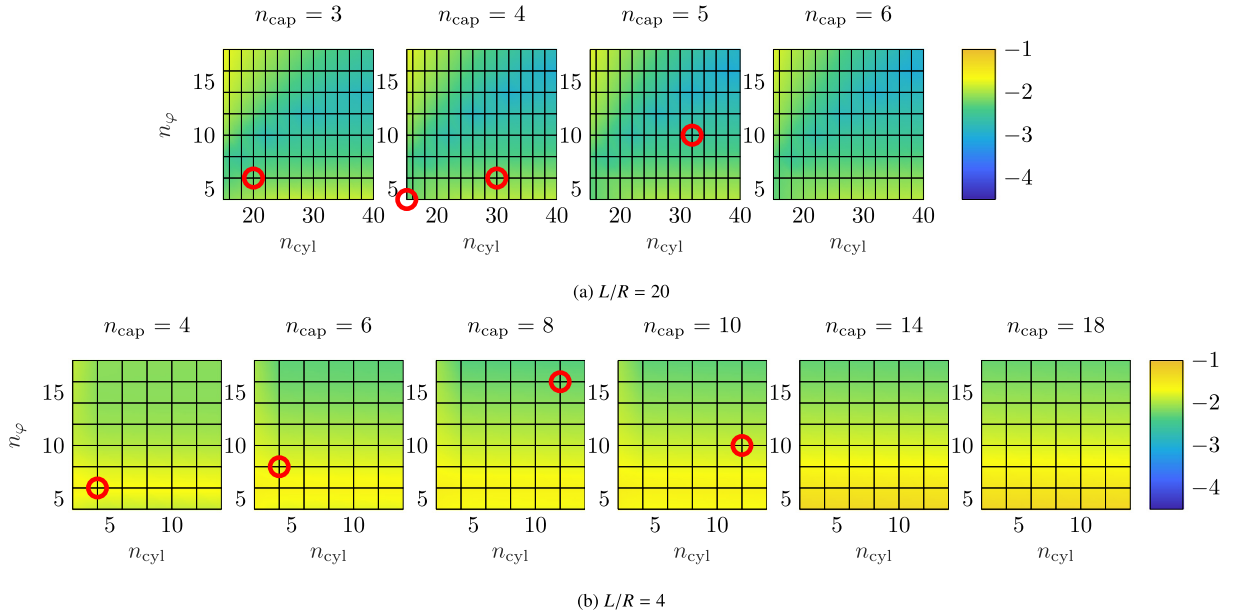


Fig. 15. Maximum relative error in the four mobility coefficients (\log_{10}) while keeping the ratio $a_h/s = 0.5$ fixed and solving the problem (19) for L_g and R_g . The same four sets of grid parameters $\{n_{\text{cap}}, n_{\text{cyl}}, n_\varphi\}$ are marked in red as in Figs. 9 and 10 for the slender rod and in Figs. 12 and 13 for the fat rod, where we optimise also for a_h/s . In Table 6, the error levels are compared, and it can be concluded that the errors are considerably larger if we do not optimise for a_h/s . Only results for grids shifted along the axis of the rod are displayed.

Table 6

Comparing relative errors in the mobility coefficients when optimising for L_g and R_g , using either an optimal blob radius ratio that is solved for in the optimisation problem, or choosing $a_h/s = 0.5$. The relative error levels in the mobility coefficients for a single particle indicate that errors are considerably larger if $a_h/s = 0.5$ is chosen. Data is presented for four chosen discretisations marked with red circles in Fig. 15, with the optimal parameters presented in Table 5. Note that the relative errors are equal for all of the four mobility coefficients (see Table 4 for the case with optimal a_h/s) and therefore only one number is reported per discretisation.

n_{cap}	n_{cyl}	n_φ	n_b	Coeff. err., optimal a_h/s	Coeff. err., $a_h/s = 0.5$
<i>L/R = 20</i>					
4	15	4	92	1.73×10^{-3}	5.23×10^{-3}
3	20	6	156	1.92×10^{-4}	5.65×10^{-3}
4	30	6	228	2.61×10^{-5}	6.67×10^{-3}
5	32	10	420	2.21×10^{-6}	2.78×10^{-3}
<i>L/R = 4</i>					
4	4	6	72	3.66×10^{-4}	2.11×10^{-2}
6	4	8	128	2.17×10^{-5}	1.71×10^{-2}
10	12	10	320	3.10×10^{-5}	1.32×10^{-2}
8	12	16	448	3.46×10^{-6}	7.51×10^{-3}

for some small ϵ (note that in general, it is not possible to satisfy (20) and (21) with $\epsilon = 0$). In other words, we match the self-interaction for the particle for translation and rotation separately. We denote the grid where the translational components are matched, fulfilling (20), as the *t-grid* and the grid where rotational components are matched as the *r-grid*, (21). If the t-grid is used to discretise the particle, we obtain errors of size ϵ in the translational velocity for a single particle. On the other hand, if the r-grid is used, errors of size ϵ are instead obtained in the rotational velocity. The errors in the rotational coefficients for the t-grid, and vice versa, can however be expected to be larger. We would like to utilise the good properties of the r- and t-grids also in a multi-particle setting. For increasing particle-particle distances, we have already mentioned that the particles behave more and more like isolated entities and the self-interaction becomes increasingly dominant. We can then expect to capture the translational velocities accurately with the t-grid and the rotational velocities accurately with the r-grid, also in the multi-particle case. The idea is therefore to solve a multi-particle mobility problem twice, once for each of the grids, keeping the translational velocities from the solution stemming from the t-grid and the rotational velocities computed with the r-grid. We term this solution strategy a *combined solve*.

Let us for a moment consider the structure of the mobility matrix \mathbf{M} . If we were to explicitly compute \mathbf{M} , the errors in the diagonal blocks would be small by doing a combined solve. These blocks represent the force-translation and torque-rotation coupling in the self-interaction for all particles in the suspension. A consequence is that we will have small relative velocity errors with the combined solve if the force and torque magnitudes are approximately equal – then, the diagonal blocks of the mobility matrix are dominating the matrix vector product $\mathbf{M}\mathbf{F}$. However, for a very large torque in relation to the force on a particle, or vice versa, off-diagonal blocks of \mathbf{M} , representing force-rotation and torque-translation couplings through the fluid, will have a larger impact. The four mobility coefficients ξ_t^\perp , ξ_t^\parallel , ξ_r^\perp and ξ_r^\parallel will be important in the representation of these blocks too. The matrix \mathbf{M} is symmetric and the r-grid will affect not only rotational components of the velocity, but also translational, and similarly for the t-grid. Hence, despite the fact that we seek two different grids, we would like the relative error in the translational mobility coefficients to be small when using the r-grid and vice versa. We refer to these errors as the *cross-errors*.

If we do not match the r- and t-grids carefully, there is a risk of large cross-errors. The idea is to try to minimise the cross-errors by varying the blob radius a_h such that the errors in the translational mobility coefficients are small when matching for the r-grid, and similarly, that the errors in the rotational mobility coefficients are small when matching for the t-grid. We solve two different optimisation problems: For the t-grid, we find a blob radius a_h^t and geometry pair (L_g^t, R_g^t) such that the error in the rotational mobility coefficients is minimised. For the r-grid, we solve the opposite problem: we find a blob radius a_h^r and geometry pair (L_g^r, R_g^r) such that the error in the translational coefficients is minimised. Mathematically, we can write the two problems as

$$\begin{aligned} \min_{L_g^r, R_g^r, a_h^r} \quad & \max \left\{ \left| \xi_t^\perp \left| \frac{1}{\xi_t^\perp} - \frac{1}{\hat{\xi}_t^\perp(L_g^r, R_g^r, a_h^r)} \right| \right|, \left| \xi_t^\parallel \left| \frac{1}{\xi_t^\parallel} - \frac{1}{\hat{\xi}_t^\parallel(L_g^r, R_g^r, a_h^r)} \right| \right| \right\}, \\ \text{s.t.} \quad & \left| \xi_r^\perp \left| \frac{1}{\xi_r^\perp} - \frac{1}{\hat{\xi}_r^\perp(L_g^r, R_g^r, a_h^r)} \right| \right| \leq \epsilon, \\ & \left| \xi_r^\parallel \left| \frac{1}{\xi_r^\parallel} - \frac{1}{\hat{\xi}_r^\parallel(L_g^r, R_g^r, a_h^r)} \right| \right| \leq \epsilon. \end{aligned} \quad (22)$$

and

$$\begin{aligned} \min_{L_g^t, R_g^t, a_h^t} \quad & \max \left\{ \left| \xi_r^\perp \left| \frac{1}{\xi_r^\perp} - \frac{1}{\hat{\xi}_r^\perp(L_g^t, R_g^t, a_h^t)} \right| \right|, \left| \xi_r^\parallel \left| \frac{1}{\xi_r^\parallel} - \frac{1}{\hat{\xi}_r^\parallel(L_g^t, R_g^t, a_h^t)} \right| \right| \right\}, \\ \text{s.t.} \quad & \left| \xi_t^\perp \left| \frac{1}{\xi_t^\perp} - \frac{1}{\hat{\xi}_t^\perp(L_g^t, R_g^t, a_h^t)} \right| \right| \leq \epsilon, \\ & \left| \xi_t^\parallel \left| \frac{1}{\xi_t^\parallel} - \frac{1}{\hat{\xi}_t^\parallel(L_g^t, R_g^t, a_h^t)} \right| \right| \leq \epsilon, \end{aligned} \quad (23)$$

with ϵ a small tolerance to be chosen. With a larger ϵ , we allow for some slackness in the constraints and a larger feasibility region, potentially leading to smaller cross-errors, while with smaller ϵ , smaller errors are obtained in the translational coefficients for the t-grid and in the rotational coefficients for the r-grid. These errors appear in the 3×3 diagonal blocks of the mobility matrix and since these diagonal blocks are dominating the matrix vector product $\mathbf{M}\mathbf{F}$, we would like to pick a small ϵ . From empirical studies of the optimal values in (22) and (23), we pick $\epsilon = 10^{-7}$ for rods with $L/R = 4$, while for $L/R = 10$, we pick $\epsilon = 10^{-6}$. As a rule of thumb, we allow for cross-errors a few orders larger in magnitude than the self-interaction error. The problems in (22) and (23) are solved with `fminimax` in MATLAB.

Note that for spheres, we can find an optimal blob radius, such that all mobility coefficients are matched simultaneously, meaning that it suffices to use a *single* optimised grid. For rods, the relative error in all four mobility coefficients cannot be forced to a prescribed tolerance ϵ simultaneously, i.e. we cannot determine (R_g, L_g, a_h) fulfilling both (20) and (21), and we therefore consider this idea where the mobility problem is determined with a combined solve from the r- and t-grids. The sizes of the cross-errors are displayed for r- and t-grids minimising (22) and (23) for a few chosen discretisation triplets $\{n_{\text{cap}}, n_{\text{cyl}}, n_\varphi\}$ in Table 7. The corresponding blob geometries (R_g^t, L_g^t, a_h^t) and (R_g^r, L_g^r, a_h^r) are reported in Table B.13 in Appendix B. Cross-errors are visualised for a larger number of discretisation triplets in Fig. 16. Note that the presented cross-errors are larger than the errors in the self-interaction obtained when all mobility coefficients are matched simultaneously for the rt-grid. The cross-errors however do not affect the largest contribution to the hydrodynamical interaction of the particles, related to diagonal blocks of the mobility matrix, and are hence less severe. In a multi-particle simulation where force and torque magnitudes are not equal, we expect the relative velocity errors to plateau at a level correlated with the cross-errors for well-separated particles, if the combined solve is applied. We investigate the accuracy using a combined solve in numerical experiments in Section 4.2.

Table 7

Relative coefficient errors in matching the r-grid and t-grid for both types of rods, upon solving (22)-(23). The cross-errors "r-trans", $|1 - \xi_t/\xi_r(R_g^t, L_g^t, a_h^t)|$, and "t-rot", $|1 - \xi_r/\xi_t(R_g^t, L_g^t, a_h^t)|$, are displayed, i.e. the relative error in the translational coefficients using the r-grid and vice versa. Note that errors in parallel and perpendicular coefficients are of the same magnitude and therefore not distinguished. For rods with $L/R = 4$, $\epsilon = 10^{-7}$ in the problems (22)-(23), while for $L/R = 20$, $\epsilon = 10^{-6}$. These levels agree well with the obtained error levels "t-trans" and "r-rot", corresponding to $|1 - \xi_t/\xi_r(R_g^t, L_g^t, a_h^t)|$ and $|1 - \xi_r/\xi_t(R_g^t, L_g^t, a_h^t)|$, meaning that the constraints in (22)-(23) are met. The resulting optimal geometry for the two grids is reported in Table B.13 in Appendix B.

n_{cap}	n_{cyl}	n_φ	aligned	r-trans	t-rot	t-trans	r-rot
$L/R = 20$							
4	15	4	No	2.31×10^{-3}	6.96×10^{-3}	1.00×10^{-6}	1.00×10^{-6}
3	20	6	Yes	2.51×10^{-4}	8.09×10^{-4}	1.00×10^{-6}	8.45×10^{-7}
4	30	6	No	3.45×10^{-5}	1.01×10^{-4}	1.00×10^{-6}	1.00×10^{-6}
5	32	10	Yes	2.20×10^{-6}	8.09×10^{-6}	1.00×10^{-6}	4.97×10^{-7}
$L/R = 4$							
4	4	6	No	4.20×10^{-4}	2.84×10^{-3}	1.00×10^{-7}	1.00×10^{-7}
6	4	8	No	2.46×10^{-5}	1.82×10^{-4}	1.00×10^{-7}	1.00×10^{-7}
10	12	10	No	3.53×10^{-5}	2.54×10^{-4}	1.00×10^{-7}	1.00×10^{-7}
8	12	16	Yes	3.93×10^{-6}	2.77×10^{-5}	1.00×10^{-7}	1.00×10^{-7}

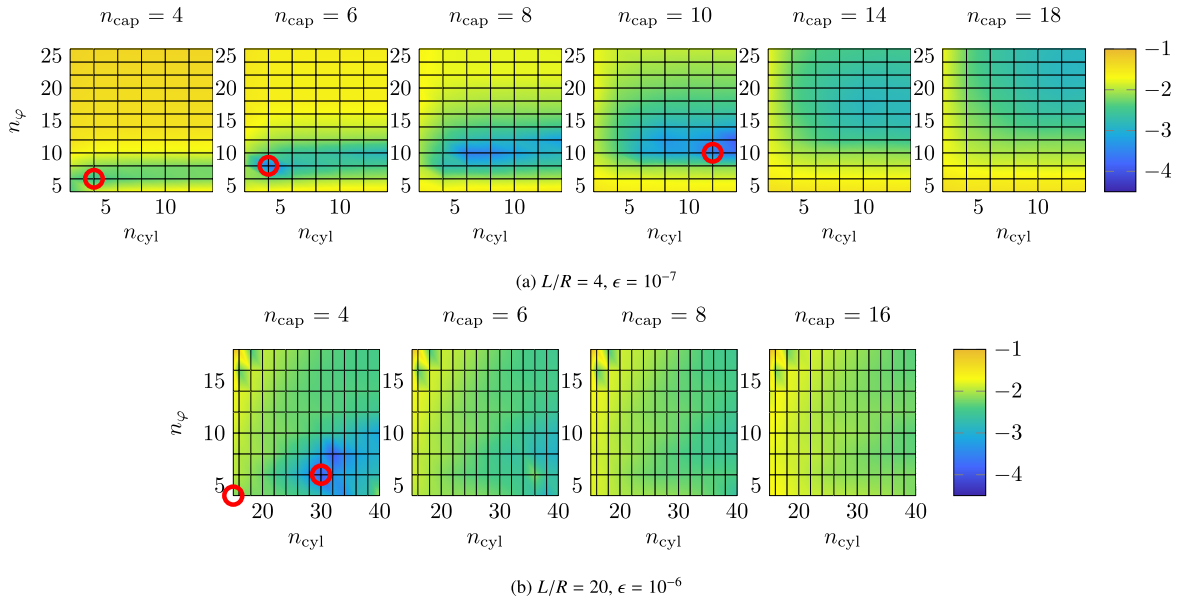


Fig. 16. Maximum cross-errors for the r- and t-grids used in the combined solve, with blob geometry computed as solution to (22)-(23). Cross-errors affect the off-diagonal blocks of the mobility matrix. In the diagonal blocks, the relative error is limited by the set tolerance ϵ in (22) and (23) for sufficiently separated particles. With a combined solve, the error level in the diagonal blocks is smaller than when using the rt-grid. Marked in red are discretisation parameters $\{n_{\text{cap}}, n_{\text{cyl}}, n_\varphi\}$ for each rod size that are considered for further numerical investigations in Section 4.2. Results for grids with shifted layers of blobs in the axial direction of the rod are displayed.

Remark 2. To be physically sensible, the mobility matrix \mathbf{M} needs to be symmetric and positive definite. First, we introduce the notation

$$\mathbf{B}_u = \begin{bmatrix} \mathbf{I} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{bmatrix} \in \mathbb{R}^6 \times \mathbb{R}^6 \text{ and } \mathbf{B}_\omega = \begin{bmatrix} \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{I} \end{bmatrix} \in \mathbb{R}^6 \times \mathbb{R}^6. \quad (24)$$

Then, we can write the mobility matrix where the t-grid and r-grid are combined as

$$\mathbf{M}_{\text{combined solve}} = (\mathbf{I} \otimes \mathbf{B}_\omega) \mathbf{M}_r + (\mathbf{I} \otimes \mathbf{B}_u) \mathbf{M}_t, \quad \mathbf{I} \in \mathbb{R}^p \times \mathbb{R}^p, \quad (25)$$

with \mathbf{M}_r the mobility matrix computed with the r-grid and \mathbf{M}_t the mobility matrix computed with the t-grid. Theoretically, we want to ensure that the mobility matrix is symmetric, which could be done by computing

$$\bar{\mathbf{M}} = (\mathbf{M}_{\text{combined solve}} + \mathbf{M}_{\text{combined solve}}^T) / 2. \quad (26)$$

Numerically however, the symmetry error is small and it is therefore sufficient to solve the problem directly as in (25), or equivalently, extracting translational velocities from the t-grid and rotational velocities from the r-grid.

Remark 3. A further option to consider with the combined solve is the application a self-correction, with $\mathbf{R}_{\text{corr}}^l = \mathbf{R}_{\text{exact}}^l - \mathbf{R}_{\text{blob}}^l$ introduced for particle l , such that the new resistance matrix is given by

$$\mathbf{R} = \mathbf{M}_{\text{blob}}^{-1} + \mathbf{R}_{\text{corr}}^{\text{self}}, \quad (27)$$

with $\mathbf{R}_{\text{corr}}^{\text{self}}$ a block diagonal matrix with

$$\mathbf{R}_{\text{corr}}^{\text{self}} = \begin{bmatrix} \mathbf{R}_{\text{corr}}^1 & & & \\ & \mathbf{R}_{\text{corr}}^2 & & \\ & & \ddots & \\ & & & \mathbf{R}_{\text{corr}}^p \end{bmatrix}. \quad (28)$$

The correction matrix $\mathbf{R}_{\text{corr}}^l$ only depends on the geometry of particle l and could be constructed from an accurate representation of the resistance matrix for a single particle, ignoring all other particles, and subtracting the corresponding contribution from the multiblob method, not to count the self-interaction twice. The technique allows for an accurate self-interaction in the diagonal blocks of \mathbf{M} in dilute suspensions, however at high cost: namely, large pollution in off-diagonal blocks of the mobility matrix as a result of the global operation of inverting the resistance matrix in (27). Such pollution is not found with the combined solve, and therefore that technique is favoured.

3. Pair-corrections

With optimised blob grids as presented in Section 2, we aim at good accuracy for well-separated particles. In this section, we consider a pair-correction strongly inspired by Stokesian dynamics, aiming to improve the accuracy also for closely interacting and moderately separated particles. Denoting the correction of the resistance matrix by \mathbf{R}_{corr} , the corrected system can be written on a saddle-point form as

$$\begin{bmatrix} \mathbf{N} & -\mathbf{K} \\ -\mathbf{K}^T & -\mathbf{R}_{\text{corr}} \end{bmatrix} \begin{bmatrix} \boldsymbol{\lambda} \\ \mathbf{U} \end{bmatrix} = \begin{bmatrix} \mathbf{0} \\ -\mathbf{F} \end{bmatrix}, \quad (29)$$

as identified by Fiore & Swan [31] (note that this is only a modification of the system in (10) to be solved in the non-corrected case). Solution methods for saddle-point problems are well-studied and reviewed in [50] by Benzi, Golub & Liesen. For this particular system, fast solution techniques can be applied as summarised in [12], utilising a fast implementation of an approximation of the matrix vector product $\mathbf{N}\boldsymbol{\lambda}$. Fiore & Swan also stress that the resistance matrix \mathbf{R} (such that $\mathbf{R}\mathbf{U} = \mathbf{F}$) can be identified as the negative Schur complement to the matrix in (29) and that the system in (29) is block diagonalisable using the Schur complement. This means that the system can be solved using GMRES (with efficient preconditioning outlined in [12,31]) to compute the matrix vector product $\mathbf{M}\mathbf{F}$ at low cost. Moreover, the correction matrix \mathbf{R}_{corr} only induces an extra cost in setting up the system, but not in solving it. It is of course possible to track the mobility matrix explicitly, solving (29) for \mathbf{U} straightforwardly to obtain $\mathbf{U} = (\mathbf{K}^T \mathbf{N}^{-1} \mathbf{K} + \mathbf{R}_{\text{corr}})^{-1} \mathbf{F}$ and identifying the corrected mobility matrix as

$$\mathbf{M} = (\mathbf{K}^T \mathbf{N}^{-1} \mathbf{K} + \mathbf{R}_{\text{corr}})^{-1}. \quad (30)$$

Accounting for lubrication effects for particles in close proximity is expensive in any grid-based method in which the fluid domain or the particle boundaries are discretised; a fine grid of the particle surfaces must be used to properly resolve the physics of the fluid. A pair-correction, as introduced in this section, can be included to avoid resolving the particle surface, but still obtain reasonable accuracy.

In Stokesian dynamics [39,51,40,29,31], a correction is added to the resistance matrix accounting for lubrication forces, so that

$$\mathbf{R} = \mathbf{M}_{\text{coarse}}^{-1} + \mathbf{R}_{\text{corr}}^{\text{pair}}, \quad (31)$$

with $\mathbf{R}_{\text{corr}}^{\text{pair}}$ built pair-wise by blocks on the form $\mathbf{R}_{\text{corr}}^{ij} = \mathbf{R}_{\text{exact}}^{ij} - \mathbf{R}_{\text{coarse}}^{ij}$, for all particles (i, j) closer to each other than some set cut-off, ignoring all other particles in the system. The correction is constructed from accurate analytical lubrication expressions for the particle pair, subtracting off the corresponding approximate resistance matrix constructed from a multipole expansion, making sure not to count the same contribution twice in the coarse representation. The small 12×12 correction matrix for the pair takes the form

$$\mathbf{R}_{\text{corr}}^{ij} = \begin{bmatrix} \mathbf{R}_{\text{corr}11}^{ij} & \mathbf{R}_{\text{corr}12}^{ij} \\ \mathbf{R}_{\text{corr}12}^{ij} & \mathbf{R}_{\text{corr}11}^{ij} \end{bmatrix}, \quad (32)$$

where the blocks on the diagonal represent self-interaction within the pair and the blocks off the diagonal represent interaction with the other particle in the pair. The contribution from every pair is added, such that the correction to the resistance matrix takes the form

$$\mathbf{R}_{\text{corr}}^{\text{pair}} = \begin{bmatrix} \sum_{j \neq 1} \mathbf{R}_{\text{corr}11}^{1j} & \mathbf{R}_{\text{corr}12}^{12} & \cdots & \mathbf{R}_{\text{corr}11}^{1p} \\ \mathbf{R}_{\text{corr}12}^{21} & \sum_{j \neq 2} \mathbf{R}_{\text{corr}11}^{2j} & & \mathbf{R}_{\text{corr}12}^{2p} \\ \vdots & & \ddots & \vdots \\ \mathbf{R}_{\text{corr}12}^{p1} & \mathbf{R}_{\text{corr}12}^{p2} & \cdots & \sum_{j \neq p} \mathbf{R}_{\text{corr}11}^{pj} \end{bmatrix}. \quad (33)$$

The submatrix $\mathbf{R}_{\text{corr}}^{ij}$ as defined in (32) is set to zero for well-separated particles. Corrections on the particle level is a standard idea in Stokesian dynamics. In contrast to treating the interaction between particles directly as in Stokesian dynamics with an RPY or higher order multipole expansion (treating spherical particles as a single blob), the coarse mobility and resistance matrices are in this paper computed using the multiblob method. The strategy to introduce corrections for multiblob particles was presented and favoured in [31]. Here, we do not think about this as a correction to account for lubrication forces only, but as a means of encoding accurate information for pairs of particles to improve on the total mobility description of the particle system. We therefore term this correction the *pair-correction*. It is however still natural to introduce a cut-off distance so that the correction is set to zero for well-separated particles.

Remark 4. Note that we only correct the resistance matrix and not the so called *grand* resistance matrix that relates the rigid body motion of the particles relative to the background flow velocity *together* with the local rate of strain of the flow to the force-torque-stresslet triplet [31]. Our frame-work could be extended to handle the grand resistance matrix, but this is outside the scope of this paper.

Analytical results for the close interaction of a particle pair are however known only for spherical and spheroidal particle geometries. In this paper, we investigate how the technique can be generalised to other particle shapes, but instead of using known analytical results for the pair interaction, the pair-wise resistance matrix is precomputed with the BIE-solver equipped with QBX (one could use any accurate method at hand). Hence, we are able to generalise also to distances outside of where any lubrication approximation would be accurate. For spheres, the resistance matrix for a particle pair depends on the distance only and we can compute an accurate interpolant for the pair resistance, which allows for rapid evaluations of the corrections. For rods, pre-computations can still be done, but creating a multi-variate interpolant is more cumbersome and is outside the scope of this work. In Section 4, pair-corrections are employed both for rods and spheres. To the knowledge of the authors, a pair-correction of this type has previously not been adopted for other particle shapes than spheres and spheroids. For spheroidal geometries, a new advantage is that the pair-corrections are applicable also for larger particle separations. This is in contrast to Stokesian dynamics, where it numerically has been shown that the technique is sensitive to the approximation introduced in an additive correction of this type, when the correction is not dominating the behaviour of the resistance matrix (which is only the case when treating lubrication for close to touching bodies) [44]. See also the discussion in Appendix D.

When pair-corrections are applied to the rt-grid, we solve the system in (29). With pair-corrections instead applied to the combined solve introduced in Section 2.3, we solve two corrected mobility problems on the form (29), one for the t-grid and one for the r-grid. This means that pair-corrections infer no additional solving cost, compared to using the combined solve as is. The small pair-correction matrix for the isolated particle pair is formed explicitly for all pairs sufficiently close to each other using the combined grid. Forming this matrix is however a cheap operation compared to working with the mobility matrix for a large system of particles. Note that we, with a pair-correction (on the rt-grid or to the combined solve), cannot guarantee that the correction is symmetric positive definite nor prove that the corrected mobility matrix has the same property. We will return to this question in the numerical results section for rods, 4.2, where we check for positive definiteness.

An alternative discussed in [31] is to correct for lubrication on the *blob* level, applicable mainly for particle shapes with an unknown analytical pair-wise resistance matrix. Interpreting each blob as a rigid sphere individually affected by lubrication forces from any close blobs belonging to *other* particles, the lubrication effects on the blob level could be incorporated by constructing the resistance matrix with corrections included for any *blobs* sufficiently close to each other [31,45,52]. As multiblob particles have a rough surface by construction, they are not allowed to come too close to each other, i.e. blobs on different particles should physically not be allowed to overlap (even if this is not mathematically hindered). Therefore, a pair-correction would only be of interest for particles of some minimum separation. On the scale of the blobs, this would correspond to large blob-blob separations relative to the blob radius between pairs of blobs for which lubrication effects are to be extracted. Therefore, each contribution to such a correction is expected to be too small in magnitude to be able to correct the error inherent in the multiblob method. Numerical tests have confirmed this hypothesis. As a consequence, we will choose to only present results for corrections on the particle level.

Table 8
Overview of numerical tests presented for spheres.

Geometry	Studied properties	Subsection
Two spheres	<ul style="list-style-type: none"> Comparing optimised grids to the choice $\hat{R}_h^{\text{trans}} = 1$ with $a_h/s = 0.5$. Grid orientation and resolution dependence for the relative error in M. 	4.1.1
A tetrahedron of spheres	<ul style="list-style-type: none"> Accuracy in \mathbf{u}^i and ω^i with the optimised grid with and without pair-correction for varying resolutions, depending on the gap δ. Accuracy for different α, with $\ \mathbf{t}^i\ = \alpha \ \mathbf{f}^i\$ and different gaps δ. 	4.1.2
Five random spheres	<ul style="list-style-type: none"> Accuracy in \mathbf{u}^i and ω^i for asymmetric particle configurations with $\alpha = 1$ and the optimised grid with and without pair-correction for varying n_b. 	4.1.3
Twisted chain of spheres.	<ul style="list-style-type: none"> Generalisation of the tetrahedron test, but for a larger number of particles and a fixed α. Pair-corrections applied only for neighbouring particles. 	Appendix C.1

4. Numerical results

We would like to quantify the error in the mobility matrix for each particle configuration, aiming for a general result for the worst possible error in solving mobility problems with particles of a certain type. One option could then be to compute the relative error in the mobility matrix, $\|\mathbf{M}_{\text{BIE}} - \mathbf{M}\|_2 / \|\mathbf{M}_{\text{BIE}}\|_2$, as this metric sheds light on the appearance of the error not only for a specific right hand side \mathbf{F} , but also for a general force/torque vector. The mobility matrix however contains elements of largely varying magnitude and is dominated by its diagonal blocks representing the force-translation and torque-rotation couplings in the self-interaction for all particles. Quantifying the relative error in the norm of the mobility matrix would hence mainly capture errors in these diagonal blocks. The interaction between particles is nevertheless important, and so is different force-rotation and torque-translation couplings through the fluid. We therefore instead quantify the error for a large number of force/torque vectors \mathbf{F} in each test, with the components of the particle forces and torques drawn independently from some distribution. In many of the simulations, we let

$$\|\mathbf{t}^i\| = \alpha \|\mathbf{f}^i\|, \quad (34)$$

with the direction of \mathbf{t}^i and \mathbf{f}^i drawn from the unit sphere. The parameter α is varied to account for three important scenarios: a dominating torque, a dominating force, or approximately equal magnitudes of the force and torque. For these cases, different blocks of the mobility matrix will be important and we can in this way better understand the error level in different parts of the mobility matrix and the worst error level in the matrix as a whole.

One strategy for quantifying the error in the matrix vector product $\mathbf{U} = \mathbf{M}\mathbf{F}$ is to extract the relative errors in the translational and rotational velocities, \mathbf{u}^i and ω^i , separately. This is the choice made for spheres. For the rods, remember that we are ultimately interested in particle dynamics. A large relative error in rotational velocities does not necessarily have a large impact in cases where the translational velocities are large, and vice versa. Hence, we choose to instead compute the relative error in the velocity at the tip of the particle, given by $\mathbf{v}_{\text{tip}}^i = \mathbf{u}^i + \omega^i \times ((L/2)\mathbf{s}^i)$, with \mathbf{s}^i the unit direction of the symmetry axis of the particle.

The error in the solutions to the mobility problems will mainly depend on the shortest distance between particles, hereafter referred to as δ or the “gap”. For rods, there is also a dependence on the relative orientation of the particles. For large gaps, we will see that the error in the self-interaction in the rt-grid will set the lowest error level attainable for dilute suspensions, in accordance with our previous hypothesis.

Parameters and settings used for the computations of the accurate reference solutions using the BIE-solver with QBX (see [34,33]) are reported in Appendix A.

4.1. Spheres

The accuracy of the multiblob method is presented with and without pair-corrections for a few spheres in different configurations with different degrees of symmetry, as summarised in Table 8. We observe that the mobility errors decrease with particle separations and that error levels obtained with the optimised multiblob grids in R_g and a_h/s can be further improved with pair-corrections. We also study the dependence on the particle resolution for the accuracy and describe the error in the rotational and translational velocities for different relations between the force and torque magnitudes (different α in (34)). The optimised grid where we have solved for R_g and a_h/s is compared to choosing R_g such that $\hat{R}_h^{\text{trans}} = 1$, with $a_h/s = 0.5$, as in [12]. We illustrate that it is important to match all mobility coefficients, as long as we do not only care about the translational velocity. As we have at hand an accurate interpolant for the mobility matrix for a pair of particles at any separation distance, the pair-correction can easily be computed for varying separation distances.

4.1.1. Two spheres

Consider a setting with two spheres at the set separation distances $\delta \in \{0.5, 1, 2, 5, 10\}$. Each particle is assigned a force and torque with $\|\mathbf{t}^i\| = \alpha \|\mathbf{f}^i\|$ and varying α . For each α , 20 different force/torque sets are assigned to the particles, with

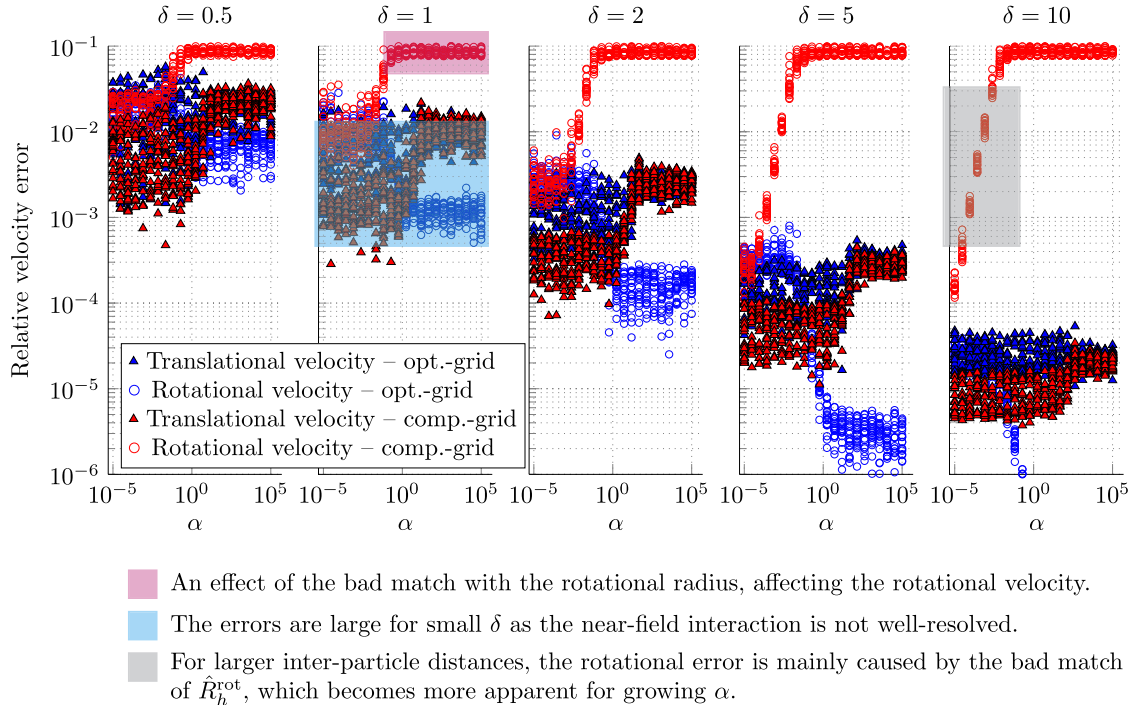


Fig. 17. Velocity errors for translation and rotation for two multiblob spheres with two different grids of 12 blobs: the comparative grid chosen so that $\hat{R}_h^{\text{trans}} = 1$ (to match a unit sphere) and $a_h/s = 0.5$, and the optimised grid given from the matching problem in (14). Both spheres are affected by a force and torque with $\|\mathbf{t}^i\| = \alpha \|\mathbf{f}^i\|$ and directions uniformly sampled from the unit sphere. For each α , 20 force/torque pairs are assigned to the particles. The depicted error for each α is the relative error for each realisation of forces and torques, displaying the maximum taken over the two spheres. The error in the rotational velocity plateaus at a high level for the comparative grid, due to the large error in the rotational hydrodynamic radius, see Table 3. The particle-particle distance δ is varied between the different panels.

directions uniformly sampled from the unit sphere. We solve the mobility problem with the optimised grid and with the comparative grid, with $\hat{R}_h^{\text{trans}} = 1$ and $a_h/s = 0.5$, as in [12].⁵ In Fig. 17, 12 blobs discretise each particle. If we only care about the translational motion of spheres, the error is low at large δ if the comparative grid is used. This is reasonable to do in a dynamic simulation, as the rotational velocity has no effect on the position of the spheres. On the other hand, if the optimised grid is used, we can obtain good accuracy at large δ also for the rotational velocity. For the comparative grid, errors in the rotational velocity plateau at the level of the relative error in the rotational mobility coefficient, as presented in Table 3. For the optimised grid, on the other hand, only the inter-particle error has an effect.

The relative error in the mobility matrix is visualised for varying δ in Fig. 18. Spheres of three different resolutions are investigated and it is clear that the error in the rotational mobility coefficient sets the error level for the comparative grid, whereas the velocity errors decrease with increasing δ for the optimised grid. For each δ , 10 different orientations of the grid are considered, and it can be concluded that the grid orientation has a very small impact on the error level. For the smallest gaps, the comparative grid captures the lubrication effects between close to touching spheres better than with the optimised grid. Note however that this is for distances closer than where blobs on adjacent particles start to overlap (marked with vertical lines in the figure). Due to the dominance of the self-interaction error using the comparative grid, we will only use the optimised grid in the remaining tests.

4.1.2. A tetrahedron of spheres

Four multiblob spheres are placed at the vertices of a tetrahedron as in Fig. 2a and the distance δ between the particles (equal among all pairs) is varied. We also vary the forces and torques by setting $\|\mathbf{t}^i\| = \alpha \|\mathbf{f}^i\|$ and sampling the direction of the force and torque uniformly from the unit sphere. The constant α is equal for all particles and for each α , 200 different force/torque sets are considered. In Fig. 19, the mean and maximum error is displayed for the translational and rotational velocity over all particles resulting from all such sets of forces and torques, using 42 blobs to discretise each sphere. Note that the behaviour of the error is the same regardless of resolution, but that the number of blobs used in the discretisation sets the error level (not displayed). The mirrored S-shape in the error for the rotational velocity indicates small errors for large α , that is, in simulations where the torque is dominant relative to the force. The smallest total error taken over all velocity components is obtained for $\alpha \approx 1$.

⁵ In that work, the effective hydrodynamic radius is chosen as $R = \hat{R}_h^{\text{trans}}$ for each resolution, given $R_g = 1$.

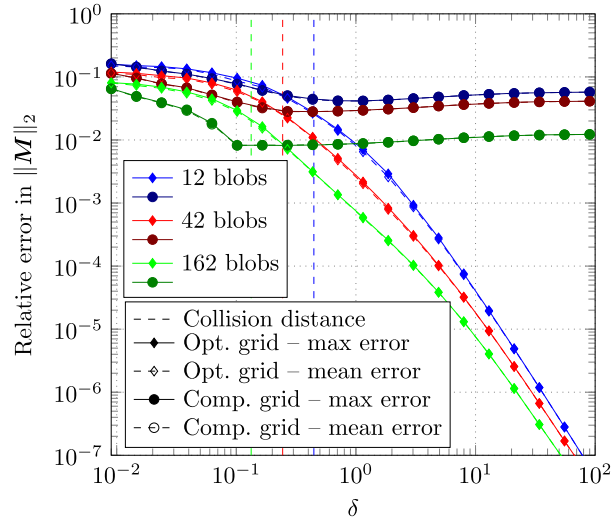


Fig. 18. Relative error in the mobility matrix for two spheres at varying gaps δ . Close to coinciding maximum and mean errors using 10 different rotations of the multiblob particles are displayed for spheres of three different resolutions. The errors with the comparative grids plateau at the level of the self-interaction error (for the rotational mobility coefficient), while the accuracy for the optimised grid is improved with δ .




Fig. 19. Velocity errors for a tetrahedron of four equally spaced spheres (see Fig. 2a) discretised with 42 blobs each, affected by a force and torque with $\|\mathbf{r}^i\| = \alpha \|\mathbf{f}^i\|$ and directions uniformly sampled from the unit sphere. Note that the same α is used for all four spheres. For each α , 200 different force/torque sets are assigned to the particles. The total error level counting both translation and rotation varies with varying α and the smallest total error is obtained for $\alpha \approx 1$. Here and in Figs. 21–23, the optimised grid is used.

We now try to improve on the accuracy further by applying pair-corrections. The test is first done with $\alpha \in \{10^{-4}, 1, 10^4\}$ and $n_b = 42$ in Fig. 20. Again, for each α , 200 sets of randomly directed forces and torques are considered and the mean and max velocity error is visualised for each α and velocity type. The pair-correction decreases the error with approximately one order of magnitude for sufficiently close particles. The improvement is notable for $\delta < 1$ and affects the dominating components of the velocity the most, i.e. translation or rotation, depending on if α is small or large. Note however that the maximum error for the smallest gaps is relatively large. In Fig. 21, the number of blobs used to discretise the particles is varied and only the error in the translational velocity is considered for $\alpha = 1$. A small improvement in accuracy can be noted with larger n_b .

4.1.3. Five random spheres

Random configurations of five unit spheres are considered of different particle densities. A minimum allowed particle-particle distance δ_{\min} is set, with δ_{\min} in the ordered list $\{0.5, 3, 6, 10\}$, meaning that no particle-pair is closer to each other than δ_{\min} for each setting. Particles are positioned at random in a cube of side length l such that the distance for each sphere to any other sphere is at least δ_{\min} , with $l = \{10, 10, 15, 20\}$. Each pair of δ_{\min} and l is tested with 20 different

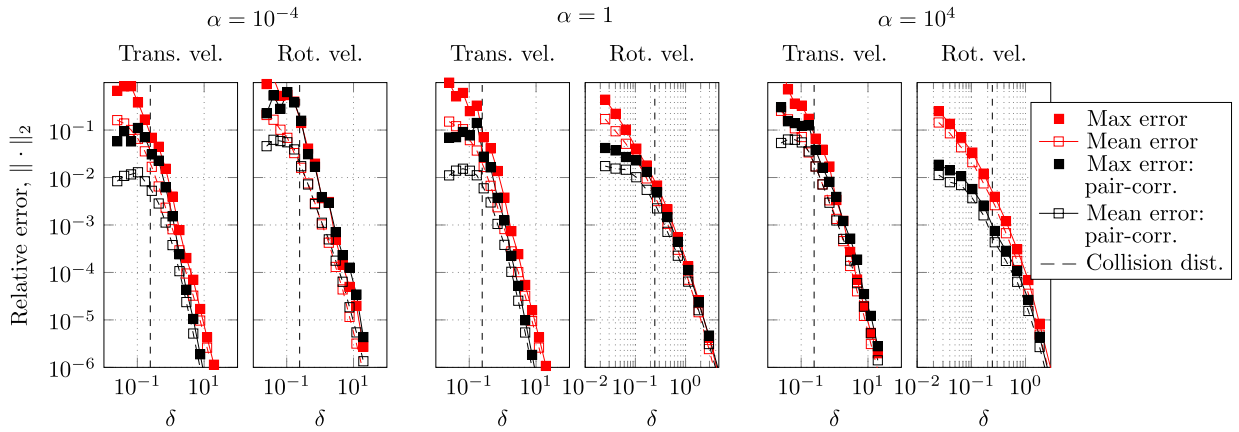


Fig. 20. Applying pair-corrections to the tetrahedron test, with spheres consisting of 42 blobs. See caption of Fig. 19 for further details. The error in the rotational velocities decays faster than the error for translational velocities when the torque is non-negligible (seen for $\alpha \geq 1$). Pair-corrections have the largest impact on the dominating component of the velocity – translational velocities for small α and rotational velocities for large α .

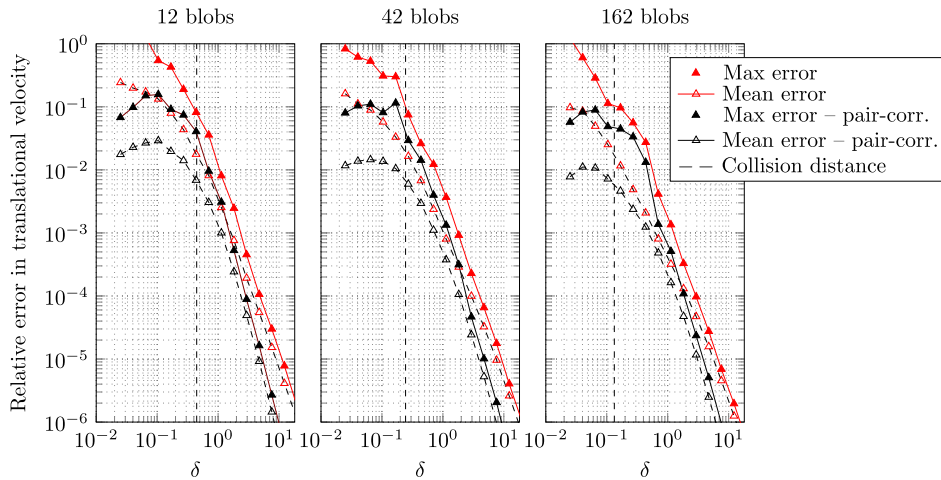


Fig. 21. Applying pair-corrections to the tetrahedron test, with spheres of varying resolutions. Only the error in translational velocity for the case $\alpha = 1$ is displayed. A slight improvement can be noted with resolution. The distance for which collision with blobs on adjacent particles might occur is marked for each resolution. See caption of Fig. 19 for further details.

random configurations, with each sphere in each configuration assigned a random force and torque, with every component independently drawn from $\mathcal{U}(-1, 1)$. This means that also here, $\|\mathbf{f}^i\| \approx \|\mathbf{t}^i\|$. The relative error in the translational velocity, \mathbf{u}^i , and angular velocity, $\boldsymbol{\omega}^i$, is computed for each sphere i for the original multiblob method with the optimised grid. Results are depicted in Fig. 22. Note the difference in accuracy using a coarse or a fine grid.

In Fig. 23, we apply pair-corrections to the resistance matrix obtained with the optimised grid for the coarsest blob resolution of the sphere. We can conclude that pair-corrections improve on the accuracy for all studied ranges of particle separations and the improvement is approximately one order of magnitude. By comparing Figs. 22 and 23, we can conclude that a similar improvement of the accuracy can be obtained with pair-corrections as by increasing the blob-resolution of the spheres.

4.2. Rods

An overview of numerical experiments involving rods is presented in Table 9. For illustration purposes, rods of two different aspect ratios are studied, a fat rod with $L/R = 4$ and a slender rod with $L/R = 20$. We remind of the relative error in the single particle mobility coefficients for the rt-grid, referred to as the self-interaction error, found in Table 4, and of the cross-errors in the r- and t-grids (which together gives the combined solve) in Table 7, for four chosen discretisations of the rod for each rod size. These errors predict the errors for sufficiently separated particles in a multi-particle setting. For some separation distance for each discretisation and grid, the error in the self-interaction will be dominating the pair-interaction error. The cross-errors affect the accuracy of the combined solve, but only enter in off-diagonal blocks of the mobility matrix (not in the self-interaction 3×3 blocks on the main diagonal).




Fig. 22. The example in Section 4.1.3: Random configurations of five unit spheres of different resolutions, affected by forces and torques with each component sampled from $\mathcal{U}(-1, 1)$. The spheres are contained in a box with minimum allowed separation distance δ_{\min} (see details in main text). The relative errors in the translational velocity, \mathbf{u}^i , and angular velocity, $\boldsymbol{\omega}^i$, are displayed for each particle in 20 different randomised configurations, yielding five marks for each configuration for each blob resolution.




Fig. 23. As in Fig. 22, but comparing the results with and without pair-corrections for the discretisation with 12 blobs, a very coarse discretisation.

In tests where the pair-correction is applied, the correction is constructed from BIE-results involving two isolated rods. To the best of our knowledge, pair-wise corrections have not previously been implemented for the rod geometry, as no exact analytical lubrication expressions exist for other shapes than spheroidal particles. For each new relative configuration of two rods (among a larger set), a new accurate resistance matrix has to be computed for the pair. In practice, the column j in the resistance matrix can be computed by solving the resistance problem with a unit translational/rotational velocity vector with elements $\mathbf{U}_i^j = \delta_{ij}$. For two interacting particles, it then takes 12 Stokes solves to determine the form of the resistance matrix. To avoid a large number of such (costly) computations, we first stick to geometries where the pair can be obtained from translation and/or rotation of some basic configurations. We will however consider a random configuration of rods in the final numerical example.

4.2.1. Two parallel rods

In this test, we explore the relation between the error in the mobility coefficients with a certain grid and the multi-particle errors at large particle-particle distances. The setting is two parallel rods of aspect ratio $L/R = 4$ with increasing separation distance. We consider five different strategies for solving the mobility problem: (i) using the optimised rt-grid where the error in all four mobility coefficients are minimised simultaneously, (ii) using the t-grid separately, where only translational mobility coefficients are matched, (iii) using the r-grid separately (matching rotational mobility coefficients), (iv) using the combined solve, where the t-grid is used to compute the translational velocities and the r-grid to compute

Table 9

Overview of numerical tests presented for rods with a progression of complexity between the consecutive tests.

Geometry	Studied properties	Section
Two parallel rods	<ul style="list-style-type: none"> The need of optimising for the blob grid geometry: A comparison of the rt-grid and a comparative grid where we only optimise for (R_g, L_g), but set $a_h/s = 0.5$. Effect of the self-interaction error in the r- t- and rt-grids and combined solve at large separations. Error levels with the rt-grid vs. a combined solve for varying n_b. 	4.2.1
Sweep test for two rods	<ul style="list-style-type: none"> As in 4.2.1, but for general relative particle orientations, visualising the gain from the smaller self-interaction error with a combined solve compared to the larger self-interaction error using the rt-grid. The interplay between the self-interaction error and the pair-interaction error. 	4.2.2
Twisted rod chain – three rods	<ul style="list-style-type: none"> As in 4.2.2, with pair-corrections applied to the rt-grid or to the combined solve for neighbouring particles, with coarse and fine meshes. The dependence of the pair-correction on the self-interaction error of the underlying grid. 	4.2.3
Twisted rod chain – eight rods	<ul style="list-style-type: none"> As in 4.2.3, but applied to a larger set of particles, displaying the error in \mathbf{v}^i, \mathbf{u}^i and ω^i. 	Appendix C.2
Random rods	<ul style="list-style-type: none"> As in Appendix C.2, but for general and asymmetric particle configurations, displaying the accuracy in \mathbf{v}^i with and without corrections to the rt-grid or to the combined solve for slender rods of a coarse and a fine resolution. Comparing a pair-corrected combined solve for a coarsely resolved fat rod to a pair-corrected solution with a fine rt-grid. 	4.2.4

the rotational velocities and, finally, (\mathbf{v}) using a comparative grid, where we optimise for (R_g, L_g) , but set $a_h/s = 0.5$. This last solution strategy is included for comparison to show the importance of optimisation of all parameters. For each inter-particle distance, the quotient between the force and torque magnitude is varied, with $\|\mathbf{t}^i\| = \alpha \|\mathbf{f}^i\|$. For each α , 100 sets of randomly oriented forces and torques are assigned to the particles, drawn from a uniform distribution on the unit sphere. In Fig. 24, the maximum relative error in the translational and rotational velocity, as compared to the corresponding BIE-solution, is visualised for each α and each choice of $\{n_{\text{cap}}, n_{\text{cyl}}, n_{\varphi}\}$. For the rightmost panels in Fig. 24, compare with Tables 4 and 7 and note that

- For the t-grid, the rotational velocity plateaus at the single particle rotational error level as represented by the t-grid cross-error, $|1 - \xi_r/\hat{\xi}_r(R_g^t, L_g^t, a_h^t)|$. For the r-grid, the error for the translational velocity plateaus at the translational error level represented by the r-grid cross-error, $|1 - \xi_t/\hat{\xi}_t(R_g^r, L_g^r, a_h^r)|$.
- The rotational velocity is resolved to an error level of 10^{-7} for the r-grid and the translational velocity is resolved to the same level for the t-grid. This is due to the tolerance ϵ chosen for the r- and t-grid optimisation problems in (22)-(23). For large α , such that the force is dominant, and for small α , where the torque is dominant, the cross-errors will have an impact on the solution.
- For the original rt-grid, both the relative error for the translational and for the rotational velocity plateau at the error level presented by the relative error in the mobility coefficients (the self-interaction error for the rt-grid).
- For the comparative grid, considerably larger error levels are obtained than with the rt-grid for all choices of discretisation parameters. The error level corresponds to the self-interaction error for the comparative grid, as presented in Table 6.
- For the combined solve, the results are better than those obtained with the rt-grid and the error plateaus at a lower level for discretisation sets $\{n_{\text{cap}}, n_{\text{cyl}}, n_{\varphi}\}$ where the error level with the rt-grid is large. This however comes at the cost of solving two mobility problems instead of one.

As the tolerance for solving the optimisation problems for the r- and t-grids in (22)-(23) is chosen to $\epsilon = 10^{-7}$, the error in mobility problems with the combined solve will not decay below this level; This is the self-interaction error for the combined solve. Similarly as for the spheres, translational relative errors are the smallest when the force magnitude is dominating and rotational relative errors are the smallest when the torque is dominating. Effects from particle interactions are important for closely interacting particles, which is a reason to why the difference between different discretisations and solution strategies is small for small δ in Fig. 24.

Similar results are obtained with the larger aspect ratio, $L/R = 20$, not included here.

4.2.2. Sweep test

In this numerical test, we investigate the worst velocity error for two neighbouring particles with varying relative orientations and compare results from using an rt-grid and a combined solve. This is a more general test than the one for

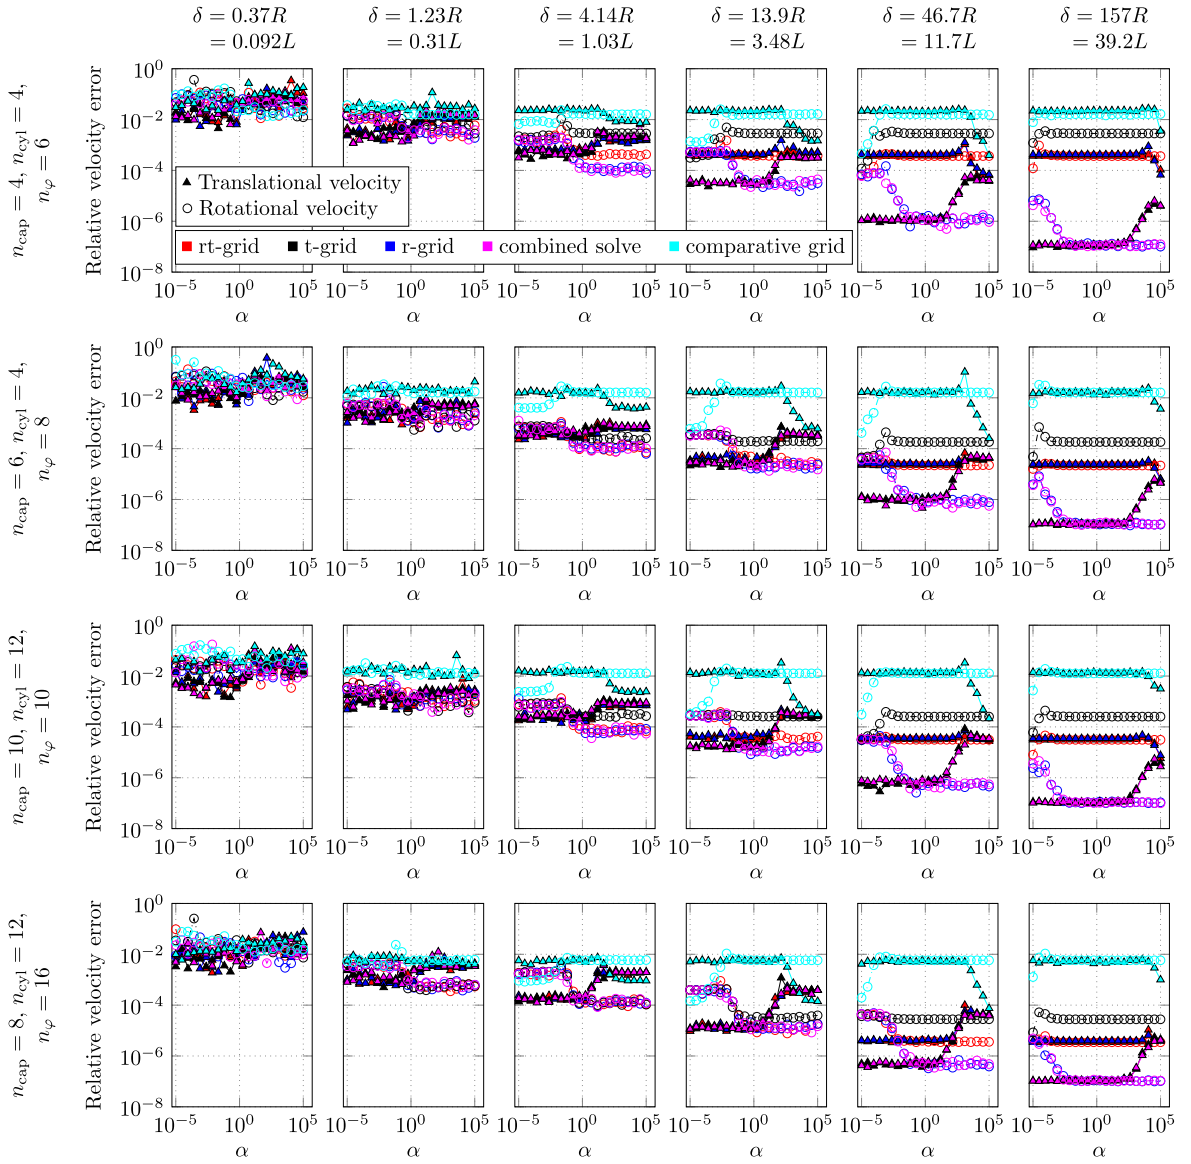


Fig. 24. Two parallel rods of aspect ratio $L/R = 4$ with different blob grids $\{n_{\text{cap}}, n_{\text{cyl}}, n_{\phi}\}$ in the four rows and different gaps δ from left to right. For large gaps (see the rightmost panels), the error levels for the r-, t- and rt-grids are all predicted from Tables 4 and 7 and the error level for the comparative grid is predicted from Table 6 (see bullet points in main text). For small gaps, the error in the pair-interaction is dominating, independently of the choice of grid.

two parallel rods in that we vary both distances and orientations, bearing in mind that parallel rods constitute a particular geometry with a high degree of symmetry. We perform a sweep test, with one rod fixed at the origin while the other rod is placed in a Cartesian 2D-grid (corresponding to the x-z-plane in Fig. 25) relative to the first rod. For each new position, the second rod is oriented randomly in the first orthant with 10 different orientations, such that the minimum distance to the first rod is kept. See Fig. 25 for an illustration with the second rod in four different positions.

The mobility matrix for the pair of rods is determined for each position and orientation. For each such configuration, 200 sets of randomly sampled forces and torques with $\|\mathbf{t}^i\| = \alpha \|\mathbf{f}^i\|$ are assigned to the pair and the resulting velocity is computed at the tip of each rod. The maximum error over the orientations in the tip velocity, as compared to the corresponding BIE-solution for the same particle configuration, is taken as a measure of the accuracy in that node of the Cartesian grid. Results are depicted in Fig. 26 for rods of aspect ratio $L/R = 20$ discretised with two different grids, $\{n_{\text{cap}} = 4, n_{\text{cyl}} = 15, n_{\phi} = 4\}$ and $\{n_{\text{cap}} = 3, n_{\text{cyl}} = 20, n_{\phi} = 6\}$. The test is conducted both with the rt-grid and the combined solve.

There is an interplay between the self-interaction error and the error due to the interaction between rods. Similarly as for the two parallel aligned rods, the error flattens out to the error level of the self-interaction for particles not too




Fig. 25. Test for two rods, where the second rod is placed in every node of a 2D-grid and randomly oriented to ten orientations for each new position. Grey indicates a fixed rod, whereas red indicates positions in a 2D grid for which the second rod is rotated in the first orthant (blue).

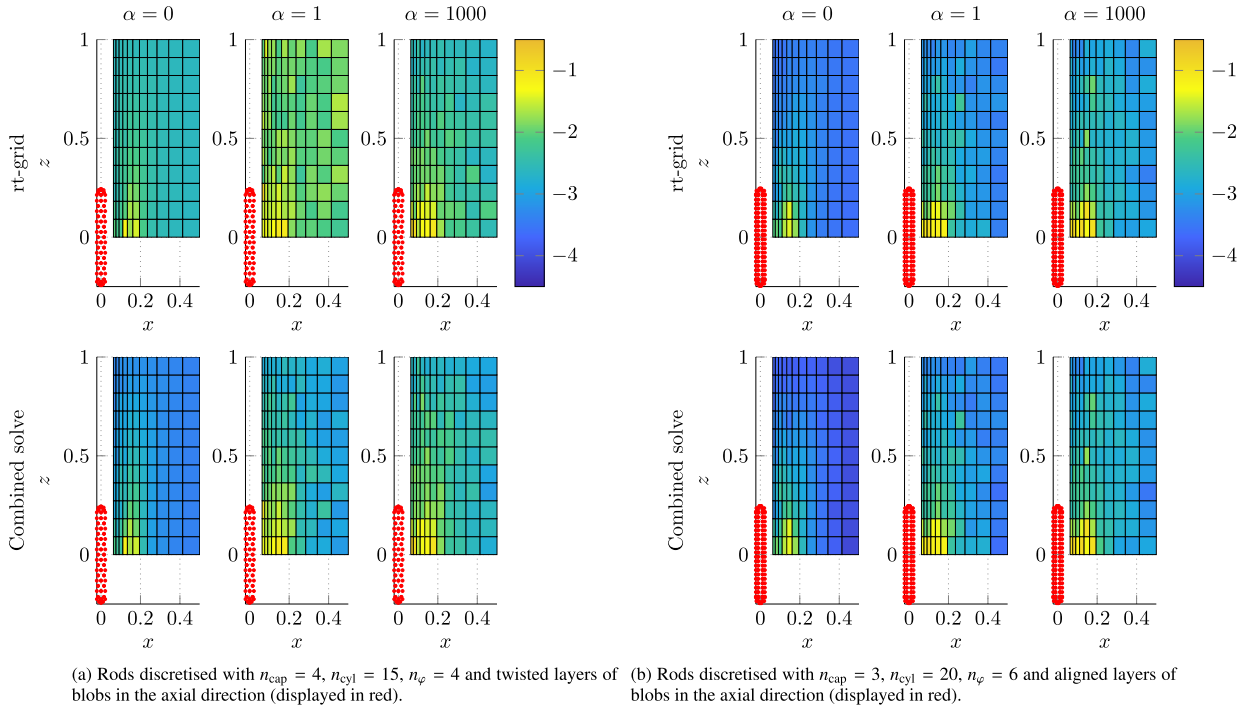


Fig. 26. Sweep test for two rods of aspect ratio $L/R = 20$. The relative error in the velocity at the tip of the two particles is measured and the maximum over all sweeps and 200 different force/torque sets is displayed in each Cartesian coordinate (\log_{10}). Three different settings are considered, with $\|\mathbf{t}^i\| = \alpha \|\mathbf{f}^i\|$, where $\alpha \in \{0, 1, 1000\}$. In the second and third column of plots, the angular velocity is more dominant, due to the choice of α . For the coarsest choice of $\{n_{\text{cap}}, n_{\text{cyl}}, n_{\varphi}\}$ in (a), a clear improvement is noted if the combined solve is used instead of the rt-grid, as the self-interaction error is reduced to a level smaller than the pair-interaction error. With the finer grid in (b), the self-interaction error is already smaller than the pair-interaction error. Hence, we cannot further reduce this error by doing a combined solve.

close to contact when using the coarse rt-grid, $\{n_{\text{cap}} = 4, n_{\text{cyl}} = 15, n_{\varphi} = 4\}$, see Fig. 26a. With the combined solve, which has a smaller self-interaction error, the error level decreases with separation, and it is instead the pair-interaction error that is dominant. For the finer grid in Fig. 26b, the pair-interaction error is already dominant and nothing is gained by doing a combined solve for the studied range of relative positions. A combined solve is then only inferring an increased computational work as two mobility problems are solved instead of one to determine each velocity vector, given a vector of forces and torques.

4.2.3. Twisted rod chain - three rods

This numerical example illustrates the relation between the mobility coefficient errors (the self-interaction error) with the rt-grid or the combined solve and the multi-particle error at large distances, similarly to the test for parallel rods and the sweep test. Here, we investigate the effect of also adding pair-corrections, which can be done with two different approaches: directly to the rt-grid or to the combined solve. In this test, these techniques are compared and we show numerically that for a pair-correction to improve on the errors due to particle interactions, the error in the self-interaction error has to be sufficiently small. If small self-interaction errors cannot be obtained with the rt-grid, a combined solve might be needed.

A chain of rods is considered, where for each chain, a unit direction vector $\mathbf{d} \in \mathbb{R}^3$ and rotation (θ, ϕ) are drawn at random from the first orthant. The chain is constructed with the first rod placed in the origin with orientation coinciding with the z-axis. The consecutive rods are obtained from the previous by translating the centre coordinate by $\beta\mathbf{d}$ in the coordinate frame of the previous particle, with the constant $\beta(\delta_{\min})$ determining the magnitude of the translation such that the smallest distance between a pair of particles is δ_{\min} , and then rotating by (θ, ϕ) . See Fig. C.37 in Appendix C.2 for chains of rods of aspect ratio $L/R = 20$.

We first do the test for a small system of three rods and generalise the test to eight rods in Section Appendix C.2. For the three rods, ten different directions \mathbf{d} and rotations (θ, ϕ) are used to generate different configurations. For each chain and inter-particle distance, 50 different sets of randomly sampled forces and torques are applied to the particles (to make the test as general as possible), with the directions drawn from the unit sphere. The error is quantified as the maximum relative error in the tip velocity of any of the three particles for each chain. If the self-interaction error in the rt-grid is large, a lot can be gained by applying the combined solve instead, to decrease the self-interaction error to a level below the error due to the interaction with other particles. Even more is gained by applying a pair-correction to the combined solve so that this interaction error is reduced. If the self-interaction error in the rt-grid is small, however, not much is gained by using the combined solve, and it suffices to do a pair-correction using the rt-grid. We see that pair-corrections to the rt-grid is beneficial for small to moderate distances, for which the error due to particle interactions is larger than the self-interaction error. The exact distance for when this happens depends on the choice of discretisation. For larger inter-particle distances, the self-interaction error starts to become dominant, which means that the error in the pair-corrected solution plateaus at the same level as without corrections.

The same conclusion can in general be drawn for all relations α between the magnitudes of the forces and torques, with $\|\mathbf{t}^i\| = \alpha\|\mathbf{f}^i\|$, but some differences can be noted. In Fig. 27, solutions for $L/R = 4$ with the larger error coarse grid $\{n_{\text{cap}} = 4, n_{\text{cyl}} = 4, n_{\varphi} = 6\}$ and the coarse grid $\{n_{\text{cap}} = 6, n_{\text{cyl}} = 4, n_{\varphi} = 8\}$ that has one order of magnitude smaller self-interaction error are compared with the finer grids $\{n_{\text{cap}} = 10, n_{\text{cyl}} = 12, n_{\varphi} = 10\}$ and $\{n_{\text{cap}} = 8, n_{\text{cyl}} = 12, n_{\varphi} = 16\}$, which has the best accuracy (two orders of magnitude smaller self-interaction error than the coarsest grid). Note that the direction of the torque relative to the orientation of the particles matters a lot here. For certain chains and torque directions, the error from the translation and rotation add up for $\alpha = 1$. For the rods of aspect ratio $L/R = 20$, with results displayed in Fig. 28, errors plateau at approximately the same constant level independently of α , corresponding to the self-interaction error. For the coarsest grid, the large cross-errors will affect the accuracy of the solution for the combined solve as a result of the very small n_{φ} used for this grid. The slight difference in the error curves for $\alpha = 1$ between rod types could be explained by the relative magnitudes of the resistance coefficients, $\xi_t^{\parallel}, \xi_t^{\perp}, \xi_r^{\parallel}, \xi_r^{\perp}$, where for instance ξ_r^{\parallel} is very large for the slender rod, see Table A.12 in Appendix A.

Similar conclusions can be drawn for longer chains of rods, see the discussion in Appendix C.2 and compare the results of Figs. 28 and C.36.

The smallest eigenvalues of the mobility matrix for the different solution strategies are displayed for the coarsest blob grids in Fig. 29 for the two rod types. It can be concluded that the smallest eigenvalues in \mathbf{M}_{BIE} are correctly captured with a pair-correction (both the one applied to the rt-grid and the correction applied to the combined solve), which indicates that lubrication effects are correctly represented. Note a small difference between the smallest eigenvalues for the pair-corrected and non-corrected solutions. The mobility matrix is positive definite for all separation distances and every tested solution strategy, which is an important property of the mobility matrix.

4.2.4. Random configurations of rods

In the previous tests with chains of twisted rods in Section 4.2.3 (and Appendix C.2), there is a lot of symmetry and pair-corrections are applied only to neighbouring particles. Now, we consider a more general setting, with random configurations of rods, with the smallest allowed distance between particle surfaces, denoted by δ_{\min} , set to a given multiple of the rod radius. Each particle is repeatedly positioned and oriented in a cube of side length l at random until its smallest distance to any other particle is larger than δ_{\min} . The δ_{\min} and l used for the two rod types are reported in Table 10. Two example geometries are visualised in Fig. 2 for the smallest δ_{\min} for each rod type. Note that, in this test, we choose to apply the pair-correction to all pairs of particles and not only to particles within some set cut-off.

Each particle is assigned a net force and torque with each component independently drawn from a uniform distribution: for the slender rods, $\mathbf{f}^i \sim \mathcal{U}(-12\pi, 12\pi)$ and $\mathbf{t}^i \sim \mathcal{U}(-1/8, 1/8)$, while for the fat rods, $\mathbf{f}^i \sim \mathcal{U}(-1, 1)$ and $\mathbf{t}^i \sim \mathcal{U}(-1, 1)$. The different scaling is motivated by the fact that a large torque on slender rods has a large impact on the rotational velocities (see the mobility coefficients presented in Table A.12 in Appendix A). For each δ_{\min} , ten random configurations are created and we study the relative errors in the tip velocities.

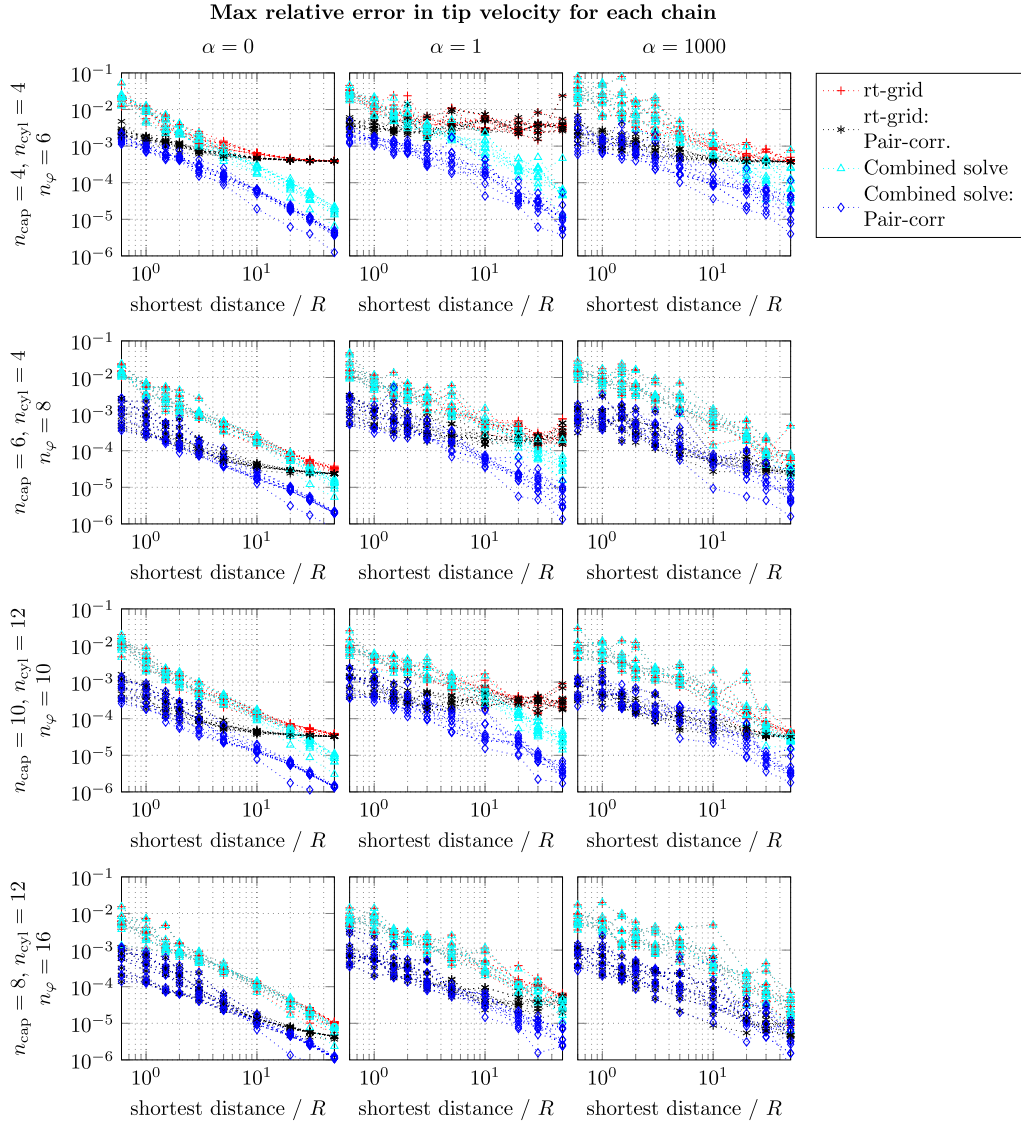


Fig. 27. Twisted chains of three rods of aspect ratio $L/R = 4$. The maximum relative tip velocity error is displayed for each of the ten chains, given 50 force/torque sets for each α . Note that pair-corrections are only applied to neighbouring particles in the chain. The error level is correlated with the error in the self-interaction for the rt-grid for all three choices of α . For overlapping red and cyan curves, the pair-interaction error is dominant - no differences can be seen between the solution with the rt-grid and the combined solve as both approaches imply a self-interaction error below the error level in the interaction between particles.

Table 10

Choices of the parameters l and δ_{\min} used in the example in section 4.2.4: random configurations of seven rods.

$L/R = 4$						
δ_{\min}	$0.5R$	$3R$	$6R$	$10R$		
l	4.5	7.0	10.0	13.0		
$L/R = 20$						
δ_{\min}	R	$6R$	$10R$	$20R$	$50R$	$100R$
l	0.375	0.75	1.0	1.6	3.0	5.5

Consider first configurations of slender rods, where the mobility problems are solved with and without pair-corrections for the rt-grid and the combined solve. Results are displayed in Figs. 30–31 for a coarse and a fine discretisation. The errors decrease with δ_{\min} only up until the distance for which the self-interaction error dominates, which happens first for the coarse discretisations (this is the case already for $\delta_{\min} = R$ in Fig. 30). Remember that δ_{\min} is only the shortest distance

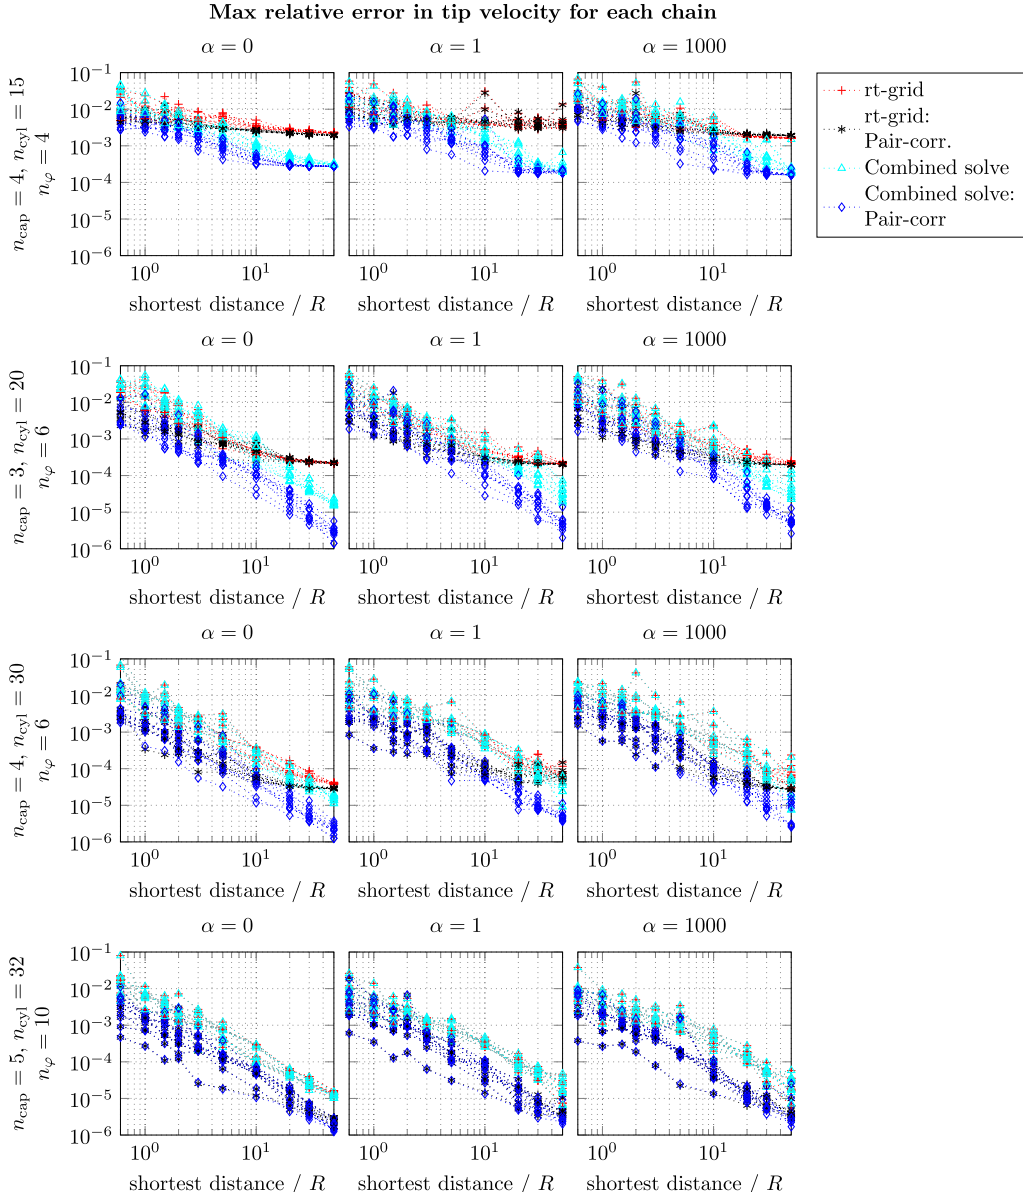


Fig. 28. The same test as in Fig. 27 (twisted chains of three rods), but for rods of aspect ratio $L/R = 20$. Errors plateau at approximately the same constant level independently of α , corresponding to the self-interaction error. For the coarsest grid in the top row of panels, the large cross-errors will affect the accuracy of the solution for the combined solve. Note that these might be due to the very small n_ϕ used for this grid.

between particles and some pairs of particles will be considerably further apart. Hence, the self-interaction error is visible also for small δ_{\min} . Note that the error level for the pair-corrected calculation plateaus at a constant level correlated with the self-interaction error in the underlying solution strategy (the rt-grid or the combined solve) and that self-interaction error is limiting the error level from below for all δ_{\min} . A small self-interaction error is essential for a pair-correction to improve on the error. Note that a pair-correction might result in an error larger than for the underlying solution strategy *without* pair-corrections. This can be seen for the pair-correction applied to the rt-grid in Fig. 30 and for the pair-correction applied to the combined solve in the right-most panels of both Fig. 30 and 31, for which the self-interaction error is at the set level ϵ used in the optimisation problems to find r- and t-grids in (22)–(23). An explanation is that the error in the self-interaction of the underlying solution strategy pollutes other blocks of the mobility matrix upon solving the linear system used to define the pair-correction. A pair-correction can hence never reduce a self-interaction error. Pair-corrections applied to the rt-grid and to the combined solve are comparable in accuracy for the minimum distances δ_{\min} for which the error due to interactions between particles is dominant. To conclude, we do not wish to apply pair-corrections for well-separated particles if the self-interaction error is large.

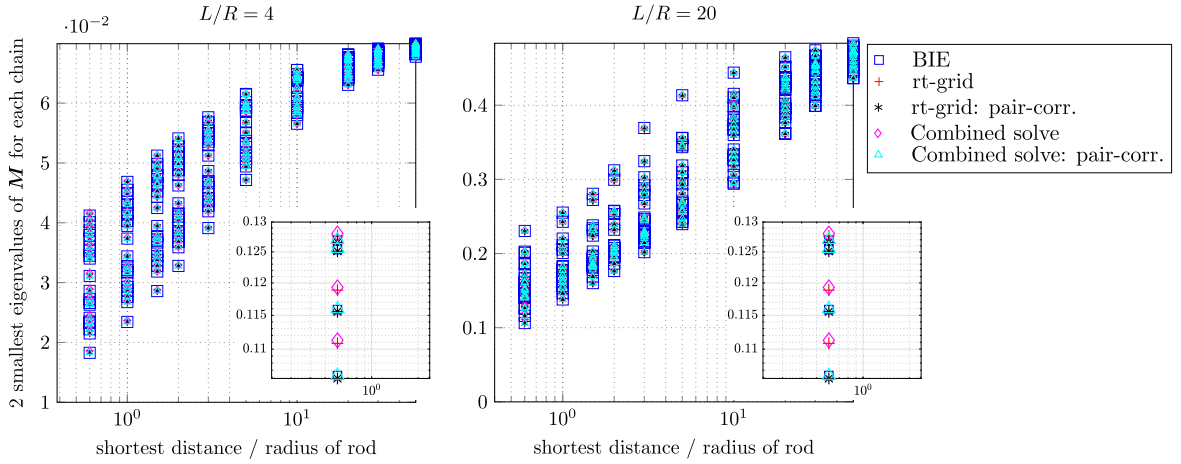


Fig. 29. For each distance, the two smallest eigenvalues are displayed for each of the ten chains, so that a total of 20 eigenvalues are shown for each distance and solution strategy. The very smallest eigenvalues are displayed in the insets. The coarsest blob discretisation is used for each aspect ratio. The smallest eigenvalues from a pair-corrected mobility matrix coincide with the smallest eigenvalues of the reference BIE-matrix, in eyeball norm. No eigenvalue from any of the multiblob strategies of solving the mobility problem is smaller than the corresponding eigenvalue with the BIE-solution and all the mobility matrices are positive definite, as required.

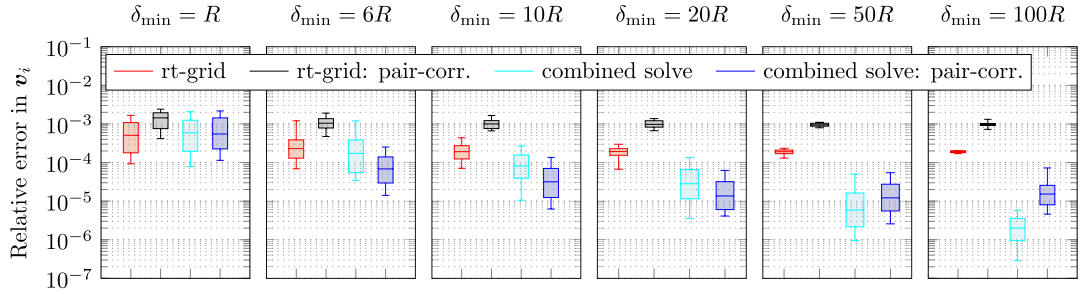


Fig. 30. Random configurations of seven rods of aspect ratio $L/R = 20$ positioned in a cube of length l (see Table 10). The rods are discretised with the coarser grid $n_{\text{cap}} = 3$, $n_{\text{cyl}} = 20$, $n_{\varphi} = 6$. Box plots show the relative error in the particle tip velocities \mathbf{v}^i for all rods in ten different configurations, with whiskers displaying the minimum and maximum error, box edges displaying the 10th and 90th percentile of the error and the box centre lines displaying the median. The combined solve with or without pair-corrections is to be preferred over solutions with the rt-grid, as the self-interaction error with the rt-grid is large.

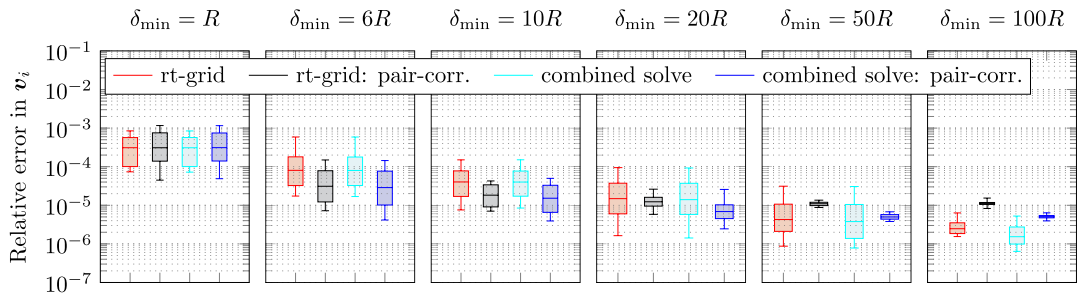


Fig. 31. As in Fig. 30, but for the finer discretisation $n_{\text{cap}} = 5$, $n_{\text{cyl}} = 32$, $n_{\varphi} = 10$. The self-interaction error with the rt-grid is almost as small as for the combined solve. For large δ_{\min} , pair-corrections plateau at a level correlated with the self-interaction error in the underlying solution strategy, but at a level larger than the self-interaction error observed for the non-corrected rt-grid and combined solve, due to error pollution (see discussion in main text).

Next, consider configurations of fat rods of aspect ratio $L/R = 4$. We use a combined solve with and without pair-corrections for a coarse grid. Results are compared with a solution computed with the rt-grid with and without corrections using a fine grid. The fine grid has more than twice the number of blobs as the coarse grid. From Fig. 32, we can conclude that by solving a mobility problem twice with a coarse grid, we can obtain the same accuracy as by using a fine grid. The pair-corrected combined solve with a coarse blob resolution outperforms the pair-corrected rt-grid with a fine resolution. The key is the reduction of the self-interaction error obtained with the combined solve. An interpretation of the result is that even if the mobility problem is solved with linear complexity, we gain from using a coarse grid and a combined solve.

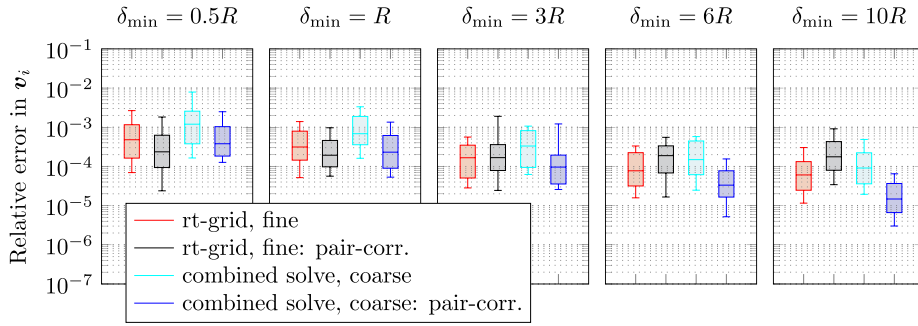


Fig. 32. As in Fig. 30, but for fat rods of aspect ratio $L/R = 4$. We compare the combined solve for a coarse grid with $n_{\text{cap}} = 4$, $n_{\text{cyl}} = 4$, $n_{\varphi} = 6$ and a solution with the rt-grid for the fine discretisation $n_{\text{cap}} = 10$, $n_{\text{cyl}} = 12$, $n_{\varphi} = 10$, i.e. with 4.44 times the number of blobs. The accuracy is approximately of the same order for the non-corrected rt-grid and combined solve. With a pair-correction applied to the combined solve, smaller errors can be obtained than by pair-correcting the rt-grid with the fine resolution.




Fig. 33. The ten smallest eigenvalues of the mobility matrix for each of the configurations of rods of aspect ratio $L/R = 4$ in Fig. 32 corresponding to $\delta_{\min} = 6R$ and the coarse grid $n_{\text{cap}} = 4$, $n_{\text{cyl}} = 4$, $n_{\varphi} = 6$. For all solution strategies with the multiblob method, the eigenvalues are very close to the eigenvalues of the reference BIE-matrix (see inset to the right for the very smallest eigenvalue in the larger left plot). This means that the corrected matrix is positive definite in all cases.

Finally, we consider the eigenvalues of the corrected mobility matrix in one specific case: rods of aspect ratio $L/R = 4$ and all ten particle configurations with $\delta_{\min} = 6$, to visualise that the corrected matrix is positive definite. The smallest eigenvalues of the mobility matrices corresponding to all four multiblob solution strategies are displayed in Fig. 33. Upon comparison with the eigenvalues of the reference BIE-matrix, it can be concluded that the ten smallest eigenvalues agree if a pair-correction is applied and that the corrected matrices are positive definite.

5. Recommended solution strategies

In this work, techniques for solving the mobility problem have been demonstrated with numerical examples presenting systems of rods or spherical particles. We would like to complete these numerical examples with a suggested solution procedure for a given general axisymmetric particle type, where we assume that accurate mobility coefficients are known. We also assume that the multiblob geometry is determined by a set of geometric parameters that we can optimise for together with the blob radius ratio a_h/s . Proceed as follows:

1. Solve the optimisation problem in (17) to minimise the maximum relative error in the four mobility coefficients $\{1/\xi_t^{\parallel}, 1/\xi_t^{\perp}, 1/\xi_r^{\parallel}, 1/\xi_r^{\perp}\}$ to find the rt-grid. Choose the number of blobs along different directions of the particle such that this error is small. This error level will bound the error from below in a multi-particle suspension of the given particle type. The total number of blobs, n_b , should be chosen so that the resulting optimised blob radius does not lead to overlapping blobs on distinct particles for inter-particle distances where the model is typically used. For very

closely interacting particles, accuracy can be gained by picking a slightly larger number of blobs, as this would reduce the pair-interaction error. This however comes at an increased cost in solving mobility problems.

2. For some combinations of particle geometry, total number of blobs and discretisation of the particle, it might be impossible to choose a grid with small errors in both translational and rotational mobility coefficients. As it is crucial to push down the self-interaction error, we then propose a different approach. If errors are not sufficiently small using the rt-grid, we can choose to optimise for two different grids, the r- and t-grids, matching for the rotational and translational mobility separately as in (22)–(23), and combine solutions from two separate mobility problems. As two solves then are required to get a solution to a single mobility problem, we need to reduce n_b for this approach.

There is hence an important choice to make in how to reduce the self-interaction error – one could try to obtain this by increasing n_b and compute an optimal rt-grid or by doing a combined solve for a smaller n_b .

3. Given a small self-interaction error using either an rt-grid or a combined solve, pair-corrections can be applied to reduce errors for interacting particles, assuming that we have an accurate method at hand that can be applied quickly for a small system of two particles only. Apply pair-corrections only for pairs of particles sufficiently close to each other by introducing a cut-off, below which the error due to interactions between particles is larger than the self-interaction error.

6. Conclusions

The accuracy of the multiblob method for the Stokes mobility problem has been studied in the deterministic setting for spherical and rod-like particles. Comparisons have been made to a reference solution accurately computed with a numerical method based on a boundary integral equation of the second kind equipped with QBX-quadrature. A particle discretised with a number of spherical blobs determining its surface is seen as a model for some ideal particle of a certain shape. Blob-blob interactions, as described via the RPY-tensor, can be viewed as a truncated expansion for the hydrodynamic interaction of the particle, while the cumulative effect from many blobs constrained to move as a rigid body capture higher-order terms of the expansion. Blobs are placed to match the single particle mobility matrix, referred to as the self-interaction, of the ideal particle, motivated by the fact that in a very dilute suspension, the hydrodynamic self-interaction of a particle is dominating the effects due to the presence of other particles. This dominance is more pronounced the larger the particle separations. Independent of how dilute the suspension is, the error level will never be lower than the error in the self-interaction. We employ an optimisation procedure to match this self-interaction. The resulting optimised grid geometry (parameters needed to determine the geometric surface where blobs are to be placed, together with the blob radius) is referred to as the rt-grid. The error level obtained upon optimising for the rt-grid will set the velocity error level for any multi-particle suspension: For sufficiently large particle-particle distances, the error plateaus at the self-interaction error level from the rt-grid. The self-interaction error sets a baseline also for denser suspensions, even though we cannot isolate the error from the interaction error with other particles in this case. We would like to stress that by solving an optimisation problem, we can obtain the smallest possible self-interaction error, given a number of blobs and a means of discretising the particle. If such an optimisation strategy is not employed and a reasonable blob-grid geometry is simply hand-picked, the self-interaction error might become several orders larger in magnitude. We have demonstrated this with numerical examples both for the sphere and the rod geometry. The level of the self-interaction error is very sensitive to the precise choice of the grid geometry and even a slight change to the optimal parameters can lead to a drastic error increase.

For spheres, the self-interaction error level also for the coarse optimised rt-grids is in the order of 10^{-8} and negligible independently of the blob grid resolution compared to errors stemming from interactions between particles. We showed that in contrast to choosing the blob grid as in [12], where only the translational hydrodynamic radius is matched, small errors can be obtained in both translational and rotational velocities with our optimisation technique. Due to the smaller optimised hydrodynamic radii a_h of the blobs, compared to the specific choice made in [12], lubrication effects are not as accurately captured as the fluid “leaks” in between blobs, distorting the repelling forces between particles at very small distances. This effect is however small and notable only for distances where blobs start to overlap. Hence, we cannot expect a sphere model to be accurate in this regime of particle-particle distances anyway.

For rods, the self-interaction error in the rt-grid is the dominant source of error for coarse discretisations. A different choice is to match for rotation and translation mobility components separately, such that we obtain one optimal blob-grid for translation and one for rotation. We can then solve the mobility problem twice and combine the solutions, extracting translational velocities from the translationally matched grid and rotational velocities from the rotationally matched. This strategy is referred to as a combined solve. The negative aspect of applying the combined solve is that we have to solve each mobility problem twice, hence doubling the cost. On the other hand, the self-interaction error can be reduced significantly, and is for large particle gaps mainly characterised by a parameter ϵ used in the optimisation procedure to adjust the interplay between different contributions to the error in the single particle mobility.

We also showed numerically that a pair-correction inspired by Stokesian dynamics can be applied for clusters of particles to improve the accuracy further. Pair-corrections applied to the rt-grid improve on the error for all gaps small enough that the self-interaction error is not the dominating contribution to the total error. For larger separations, the self-interaction error of the rt-grid limits the velocity errors from below on a certain level. Pair-corrections can hence only be successfully

applied if pair-interaction errors are larger than the self-interaction error level. For coarse discretisations, a pair-correction applied to the combined solve is therefore beneficial, as the self-interaction error of the combined solve is lower. Said differently, a well-matched self-interaction is a necessity for pair-corrections to improve on the accuracy.

The mobility problem has been studied for a small number of particles with varying relative distances and orientations. In conclusion, coarsely resolved particles of the studied type can successfully be used to approximate the solution to mobility problems, with controllable accuracy compared to a set of reference particles, where the mobility problem is solved with a much more accurate but also computationally expensive method. In future work, the multiblob method will allow for the study of heterogeneous suspensions containing particles of different shapes and/or aspect ratios, each with an optimised grid geometry. Another future research direction is to extend the blob grid optimisation strategy to systems with boundaries and/or periodicities.

In contrast to the case for Stokesian dynamics, the additivity assumption inherent in the pair-corrected method does not violate the accuracy. From the construction of the pair-corrections, we however have no technique to guarantee the positive-definiteness of the correction matrix or the resulting mobility matrix. In all numerical tests performed with the pair-correction, the resulting (corrected) mobility matrix is positive definite, which has been checked by explicitly computing its eigenvalues. As pointed out in a recent paper by Orsi et al. [53], recommended by one of the anonymous reviewers, the problem with additivity can be resolved by correcting only for *relative* motion in each particle pair. Such a correction can be constructed by a projection onto the space of relative motion [14]. We note however that differences in the resulting accuracy with such a projection compared to what is done in this paper are small, the cause of which is further discussed in Appendix D. Summarising the findings in Appendix D, most of the corrected velocity is already inherently from relative motion, which we, in turn, observe is another consequence of using a well-matched grid. For the pair-correction to be competitive in large scale simulations, a multivariate interpolant over relative orientations and particle distances is required for the fast evaluation of the pair mobility matrix for non-spherical particles. To be more specific, the pair-corrections for a pair of general axisymmetric particles can be obtained from interpolation in four dimensions as specified in Appendix E. One option to approach the problem could be to apply radial basis functions as has recently been proposed in 2D, where a partition of unity approach is applied to reduce problems of ill-conditioning [54,55]. Benefits are that a tensor product structure could be avoided for the 4D problem, that the resolution can be adaptively set by changing the size of the partitions, and that the interpolation can be done locally. A pre-computed function of the described type is to be developed in future work. If Brownian motion is accounted for in the system, it is desired that the correction built from this interpolant is positive definite – a property that also has to be considered in the sphere case where a 1D interpolant already is available.

In this work, a smaller number of particles have been studied, allowing for comparisons to an accurate method where the surface of the particles is well resolved. In future work, we plan to use the same framework but move from the deterministic setting to study the diffusive behaviour in Brownian systems containing a large particle number. Control of the deterministic error in the multiblob approach, as provided by the techniques presented in this work, is a necessary piece in understanding also other contributions to the total error in a dynamic simulation, such as a time-discretisation error and a statistical error in a Brownian setting.

CRedit authorship contribution statement

Anna Broms: Conceptualization, Data curation, Formal analysis, Investigation, Methodology, Software, Validation, Visualization, Writing – original draft, Writing – review & editing. **Mattias Sandberg:** Methodology, Supervision, Writing – review & editing. **Anna-Karin Tornberg:** Funding acquisition, Methodology, Supervision, Writing – review & editing.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

Data will be made available on request.

Acknowledgements

We acknowledge the support from the Swedish Research Council: grant no. 2019-05206 and the research environment grant INTERFACE (biomaterials), no. 2016-06119, and thank the anonymous reviewers for several useful comments. This work is done in connection with the Linné Flow Centre at KTH.

Appendix A. BIE-parameters and reference mobility coefficients

Parameters used to compute the reference solution with the BIE-method equipped with QBX are presented in Table A.11, with the discretisation described in detail by Bagge & Tornberg in the appendix of [33]. For spheres, 60 Gauss-Legendre

Table A.11

Parameters for computing the reference solution for rods using the BIE-solver with QBX.

Aspect ratio	n_{cap}^{QBX}	n_{cyl}^{QBX}	n_{φ}^{QBX}
$L/R = 4$	40	10	25
$L/R = 20$	35	60	18

Table A.12

Mobility coefficients computed with the BIE-method for a single rod. These will be used as reference when optimising for the blob grid.

	$1/\xi_t^{\parallel}$	$1/\xi_t^{\perp}$	$1/\xi_r^{\parallel}$	$1/\xi_r^{\perp}$
$L/R = 4$	0.0713091832	0.0804688924	0.0836722897	0.153993117
$L/R = 20$	0.504333544	0.701528202	14.4347584	251.837496

Table B.13

Geometry obtained for the r- and t-grids upon solving the optimisation problems in (22) and (23), corresponding to the error levels reported in Table 7 in Section 2.3. For finer discretisations, the difference in the parameters ($L_g, R_g, a_h/s$) is small for the r- and t-grids; they are however not identical, even if it appears so in the table.

n_{cap}	n_{cyl}	n_{φ}	aligned	L_g^t	R_g^t	L_g^r	R_g^r	a_h^t/s	a_h^r/s
$L/R = 20$									
4	15	4	No	0.484	0.0213	0.486	0.0214	0.359	0.359
3	20	6	Yes	0.493	0.0229	0.492	0.0228	0.329	0.336
4	30	6	No	0.495	0.0232	0.495	0.0232	0.267	0.267
5	32	10	Yes	0.496	0.0235	0.496	0.0235	0.342	0.342
$L/R = 4$									
4	4	6	No	1.929	0.477	1.933	0.476	0.222	0.223
6	4	8	No	1.953	0.479	1.952	0.479	0.221	0.221
10	12	10	No	1.967	0.483	1.966	0.484	0.202	0.201
8	12	16	Yes	1.966	0.483	1.966	0.483	0.305	0.305

Table B.14

Grid parameters for each of the five fat rods of aspect ratio $L/R = 4$ forming a circle in the example presented in Fig. 1. Each rod is discretised with $n_{\text{cap}} = 6$, $n_{\text{cyl}} = 4$ and $n_{\varphi} = 8$ and the non-optimised grid is the optimised grid perturbed by 1%.

Grid	L_g	R_g	a_h/s
Optimised	1.953	0.479	0.221
Non-optimised	1.972	0.484	0.223

points are used in the axial direction (from pole to pole) and 60 equidistant points are chosen in the azimuth direction for each spherical particle.

The computed accurate mobility coefficients for rods of aspect ratio $L/R = 4$ and $L/R = 20$ are reported in Table A.12.

Appendix B. Grid parameters

Optimal blob geometries corresponding to the r- and t-grids are presented in Table B.13 for selected discretisation triplets $\{n_{\text{cap}}, n_{\text{cyl}}, n_{\varphi}\}$ for the slender rod with $L/R = 20$ and the fat rod with $L/R = 4$. For the introductory numerical example in Fig. 1, parameters are presented in Table B.14.

Appendix C. Pair-corrections for particle chains

C.1. A twisted chain of spheres

We consider a setting where seven spheres form a chain: The first sphere is placed at the origin, with the next placed at a prescribed distance δ in the randomised direction \mathbf{v} from the first particle, rotated by (θ, φ) . Here, \mathbf{v} is a vector in the first orthant and (θ, φ) are sampled from the first orthant. The next spheres in the chain are placed accordingly, by translating the centre coordinate by $\delta\mathbf{v}$ in the coordinate frame of the previous particle and then rotating. Example

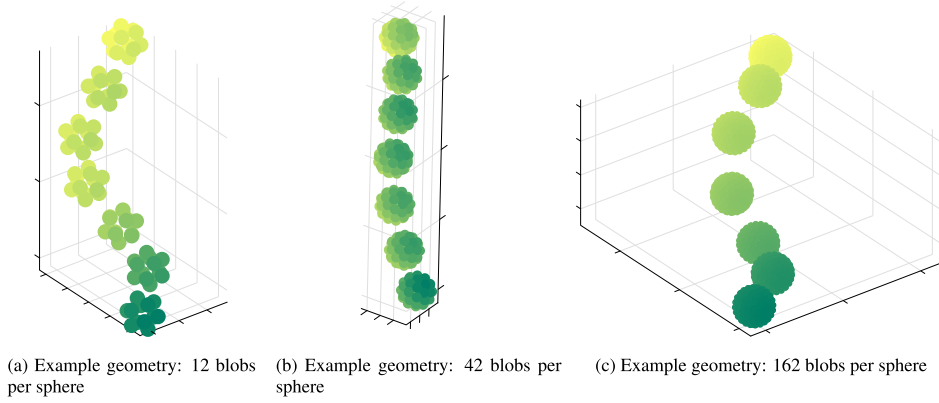


Fig. C.34. The example in Section Appendix C.1: Chains of seven spheres with the same relative transformation used between two consecutive particles. Depth is indicated by colour.

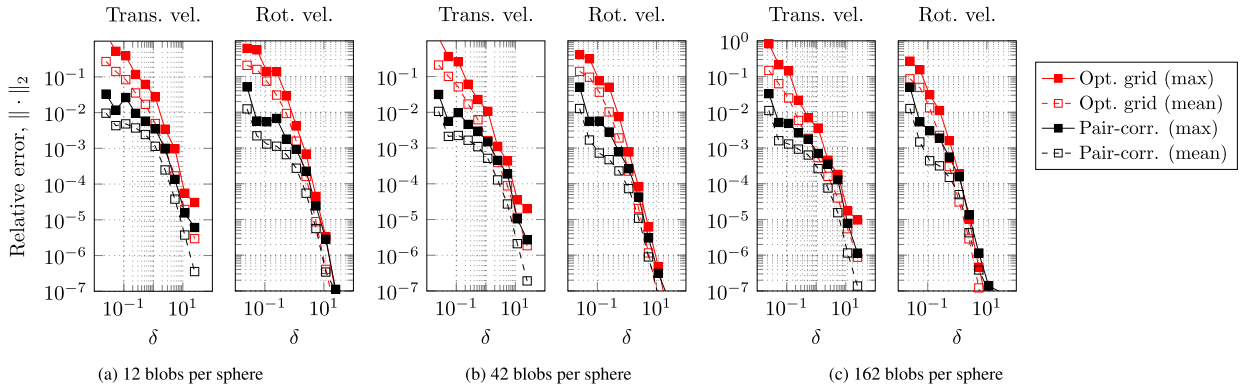


Fig. C.35. Velocity errors for data accumulated from all seven spheres in ten different randomised particle chains, with example chains displayed in Fig. C.34, formed with the same relative transformation between two consecutive particles. The optimised multiblob grids defined by (R_g, a_h) are used for each resolution. The mean and max velocity error is displayed for the translational velocity and rotational velocity, versus the inter-particle distance δ for different number of blobs in the discretisation. Pair-corrections improve on the error if δ is sufficiently small.

geometries are visualised in Fig. C.34. For each chain, a particular force and torque is set for each sphere, chosen with each component independently drawn from $\mathcal{U}(-1, 1)$ resulting in a somewhat arbitrary force/torque vector with $\|\mathbf{f}^i\| \approx \|\mathbf{t}^i\|$. This choice means that we mainly look at the error in the main diagonal blocks of the mobility matrix as these blocks will be dominant. The mean and maximum relative error in the particle velocity for three different resolutions of the multiblob spheres are displayed in Fig. C.35 with and without pair-corrections. An improvement in accuracy can be noted for the finest blob grid compared to the coarsest and pair-corrections improve on the accuracy for the small inter-particle distances.

C.2. Twisted rod chains - eight rods

We repeat the test with rod chains, as described in Section 4.2.3, but now for eight particles and a specific force/torque vector \mathbf{F} , with each component independently drawn from $\mathcal{U}(-1, 1)$. This means that here, $\|\mathbf{f}^i\| \approx \|\mathbf{t}^i\|$. Statistics for the relative error in the translational, rotational and tip velocities for all particles in ten different chains are visualised in Fig. C.36, for rods with $L/R = 20$ and a semi-coarse discretisation. Pair-corrections are only applied to neighbouring particles. If applied to the rt-grid, error levels are decreased for small δ_{\min} , whereas for distances δ_{\min} where the self-interaction start to become dominant in the mobility matrix relative to the errors due to particle interactions, the errors corresponding to pair-corrections plateau at the same level as when using the non-corrected rt-grid. If the combined solve is used instead, the accuracy can be further enhanced with a pair-correction, as self-interaction errors then are small.

Appendix D. Constraints on the correction matrix

As pointed out in [14,53,44], pair-corrections should be done only to the relative motion within a particle pair. Phrased differently, the pair correction should be frame invariant. In the work by Orsi et al., a set of constraints on the individual functions building up the correcting pair resistance matrix are derived to obtain frame invariance for spheres. The authors

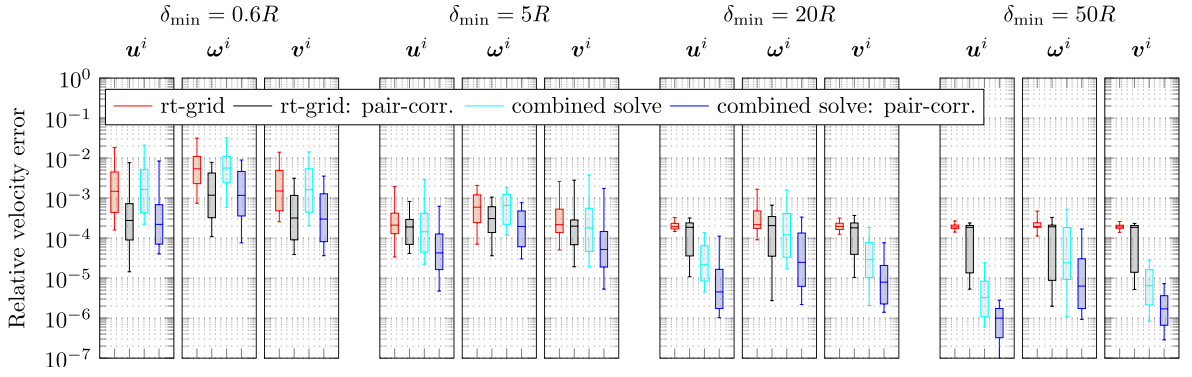


Fig. C.36. Relative error in the translational, rotational and tip velocity for each rod in each of the ten chain configurations displayed in Fig. C.37 for slender rods with $L/R = 20$ discretised with the grid $n_{\text{cap}} = 3$, $n_{\text{cyl}} = 20$, $n_{\varphi} = 6$. Whiskers in box plots display the minimum and maximum error, box edges display the 10th and 90th percentile of the error and the box centre lines display the error median. For larger particle distances (in the rightmost panels), the error using the rt-grid with or without pair-corrections is dominated by the self-interaction error, and plateaus at this level. The self-interaction error can be improved upon with a combined solve and a pair-correction applied to the combined solve can reduce interaction errors. For small separation distances (in the leftmost panels), a pair-correction can be applied directly to the rt-grid to enhance accuracy, as the interaction error is the dominating contribution to the error here. Note that the pair-corrections are only applied to neighbouring particles in this test.

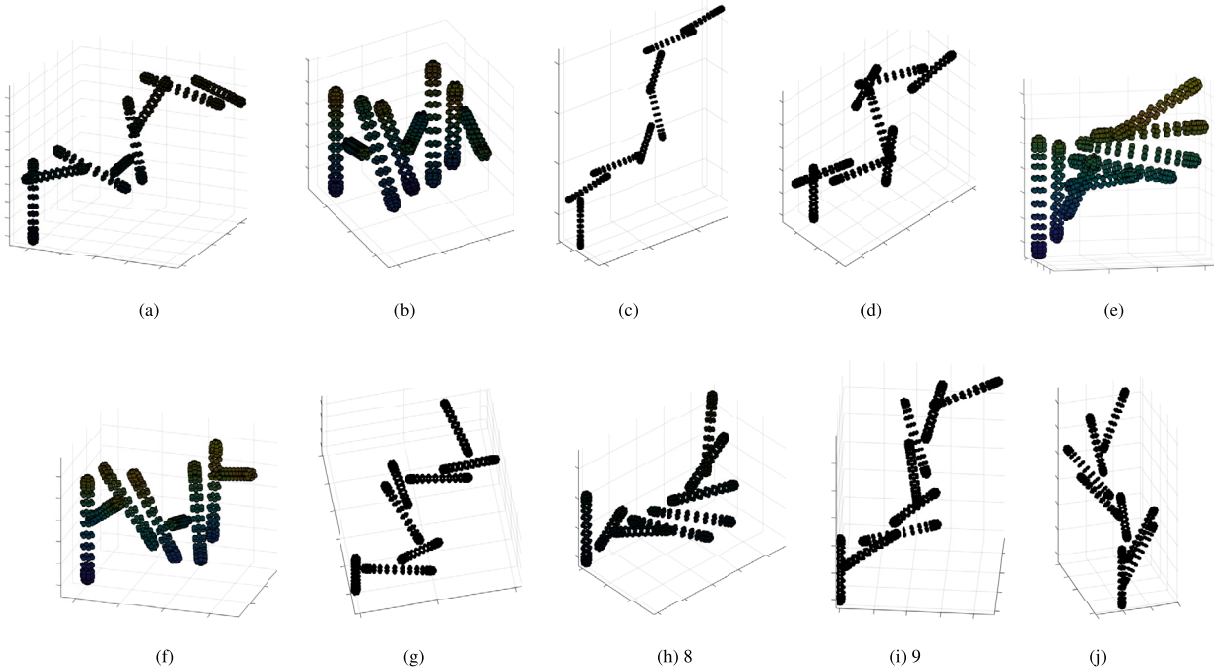


Fig. C.37. Chains of eight rods of aspect ratio $L/R = 20$ with the same relative transformation used between two consecutive rods considered in the example in Section Appendix C.2. Similar chains of three rods are considered in Section 4.2.3.

show that these constraints are naturally obeyed to sufficient accuracy with the numerical method applied in the paper (the fictitious domain method). For completeness, we indicate that the multiblob method of this paper has the same property, by reproducing the example in Section 3.2 and Fig. 2 of [53]: Given two spheres separated by a distance δ , the force on the first particle due to a translation along the line of centres of the pair with a velocity \mathbf{U} should obey (see eq. 26 of [53])

$$\mathbf{f}_1 = \left(X_{11}^A(\delta) + X_{12}^A(\delta) \right) \mathbf{U} = \mathbf{0}, \quad (\text{D.1})$$

where X_{11}^A and X_{12}^A are two of the functions that build up the resistance matrix. Similarly, for a velocity $\hat{\mathbf{U}}$ perpendicular to the line of centres, the relation




Fig. D.38. For a pair-correction only to correct for relative motion of a pair of spheres, two of the constraints that need to be fulfilled are given in (D.3). The relative errors for these constraints are here reported for varying resolutions of multiblob spheres versus the inter-particle distance δ .

$$\mathbf{f}_1 = \left(Y_{11}^A(\delta) + Y_{12}^A(\delta) \right) \hat{\mathbf{U}} = \mathbf{0} \quad (\text{D.2})$$

should hold for two other functions Y_{11}^A and Y_{12}^A determining the resistance matrix [53]. The four functions $X_{11}^A, X_{12}^A, Y_{11}^A$ and Y_{12}^A can be individually identified given a resistance matrix, and the precise structure of the entire resistance matrix is clearly presented in Appendix A1 of [53]. The constraints in (D.1) and (D.2) can be written as

$$\begin{aligned} X_{11}^{A,\text{multiblob}}(\delta) + X_{12}^{A,\text{multiblob}}(\delta) - X_{11}^{A,\text{accurate}}(\delta) - X_{12}^{A,\text{accurate}}(\delta) &= 0, \\ Y_{11}^{A,\text{multiblob}}(\delta) + Y_{12}^{A,\text{multiblob}}(\delta) - Y_{11}^{A,\text{accurate}}(\delta) - Y_{12}^{A,\text{accurate}}(\delta) &= 0, \end{aligned} \quad (\text{D.3})$$

for the resistance functions computed from the multiblob resistance matrix and our accurate interpolant respectively. In Fig. D.38, the differences in the left hand side of (D.3) are visualised for multiblob spheres of varying resolutions. For the multiblob method, the constraints in (D.3) are matched to a relative error level of at most a fraction of a percent, comparable to the results reported in [53]. Note that these results are a consequence of having a well-matched grid for a single particle. If the grid is not carefully matched, the constraints that assure that only relative motion is corrected, e.g. (D.3), are *not* satisfied. Hence, the grid optimisation is important for the pair-corrections to be successful. The full set of constraints that have to be fulfilled to correct only for relative motion can alternatively be represented via a projection onto the space of relative motions [14]. Note that with the projection in place, so that the conditions in (D.3) are fulfilled exactly, we observe numerically that negative eigenvalues that sometimes occur in the small pair-wise correction matrices are removed, which is a welcome property as we would like the global correction matrix to be positive definite.

Appendix E. The four coordinates needed to specify the relative position of two rod-like particles

This section describes the coordinates needed to characterise the relative position of two equal axi-symmetric particles. Interpolation over these four coordinates can be done to facilitate fast evaluation of the pair resistance matrix. The construction of the coordinates is depicted in Fig. E.39. As shown by Fig. E.39a we fix a Euclidean coordinate system (x, y, z) by positioning the origin to the centre of particle 1. The z -axis is chosen parallel to particle 1, and the x - z plane is chosen such that the centre of particle 2 lies in this plane. The remaining y -direction is determined by the fact that we want the (x, y, z) coordinate system to be orthonormal and right-oriented. The distance between the centres of the particles is denoted r , and the angle between the vector from the centre of particle 1 to the centre of particle 2 and the positive x -axis is denoted α . Note that we always can orient the x, y, z -frame so that $-\pi/2 \leq \alpha \leq \pi/2$. The rotational position of particle 2 relative to particle 1 is depicted in Fig. E.39b. We denote by (x', y', z') the coordinate axes parallel to (x, y, z) , and centred at the centre of particle 2. We let ϕ and θ be the spherical angular coordinates of the end point of particle 2 whose z' -coordinate is greater than or equal to zero. In summary it is possible to interpolate the pair resistance matrix using the coordinates $(r, \alpha, \phi, \theta)$ over $0 \leq r \leq r_{\max}$, $-\pi/2 \leq \alpha \leq \pi/2$, $0 \leq \phi \leq \pi/2$, $0 \leq \theta \leq 2\pi$.

Another possible choice of coordinates to interpolate over would be to use the x - and z -coordinates of the centre of particle 2 relative to particle 1, together with the particle 2 angles ϕ and θ . Since, however, the angular coordinates, in contrast to the distance coordinates, have inherent bounds it seems more natural to choose a coordinate system involving three angular coordinates rather than a coordinate system involving only two.




Fig. E.39. Coordinate system introduced to describe relative positions for two axisymmetric particles.

References

- [1] K.M. Håkansson, A.B. Fall, F. Lundell, S. Yu, C. Krywka, S.V. Roth, G. Santoro, M. Kwick, L. Pahl Wittberg, L. Wågberg, L.D. Söderberg, Hydrodynamic alignment and assembly of nanofibrils resulting in strong cellulose filaments, *Nat. Commun.* 5 (2014).
- [2] K.M. Håkansson, F. Lundell, L. Pahl-Wittberg, L.D. Söderberg, Nanofibril alignment in flow focusing: measurements and calculations, *J. Phys. Chem.* 120 (2016) 6674–6686.
- [3] V. Krishne Gowda, C. Brouzet, T. Lefranc, L.D. Söderberg, F. Lundell, Effective interfacial tension in flow-focusing of colloidal dispersions: 3-D numerical simulations and experiments, *J. Fluid Mech.* 876 (2019) 1052–1076.
- [4] W. Yan, H. Zhang, M.J. Shelley, Computing collision stress in assemblies of active spherocylinders: applications of a fast and generic geometric method, *J. Chem. Phys.* 150 (2019).
- [5] V.K. Gowda, T. Rosén, S.V. Roth, L.D. Söderberg, F. Lundell, Nanofibril alignment during assembly revealed by an X-ray scattering-based digital twin, *ACS Nano* 16 (2022) 2120–2132.
- [6] S. Kim, S.J. Karrila, *Microhydrodynamics: Principles and Selected Applications*, Butterworth-Heinemann, 1991.
- [7] D.L. Ermak, J.A. McCammon, Brownian dynamics with hydrodynamic interactions, *J. Chem. Phys.* 69 (1978) 1352–1360.
- [8] M. Graham, *Microhydrodynamics, Brownian Motion, and Complex Fluids*, Cambridge Texts in Applied Mathematics, Cambridge University Press, Cambridge, 2018.
- [9] T. Ando, E. Chow, Y. Saad, J. Skolnick, Krylov subspace methods for computing hydrodynamic interactions in Brownian dynamics simulations, *J. Chem. Phys.* 137 (2012).
- [10] S. Delong, F.B. Usabiaga, R. Delgado-Buscalioni, B.E. Griffith, A. Donev, Brownian dynamics without Green's functions, *J. Chem. Phys.* 140 (2014).
- [11] B. Sprinkle, F. Balboa Usabiaga, N.A. Patankar, A. Donev, Large scale Brownian dynamics of confined suspensions of rigid particles, *J. Chem. Phys.* 147 (2017).
- [12] F.B. Usabiaga, B. Kallemov, B. Delmotte, A.P.S. Bhalla, A. Donev, B.E. Griffith, Hydrodynamics of suspensions of passive and active rigid particles: a rigid multiblob approach, *Commun. Appl. Math. Comput. Sci.* 11 (2016).
- [13] J.G. Kirkwood, The general theory of irreversible processes in solutions of macromolecules, *J. Polym. Sci.* 34 (1954) 597–610.
- [14] B. Cichocki, K. Hinsen, Stokes drag on conglomerates of spheres, *Phys. Fluids* 285 (1994).
- [15] J. García De La Torre, M.L. Huertas, B. Carrasco, Calculation of hydrodynamic properties of globular proteins from their atomic-level structure, *Biophys. J.* 78 (2000) 719–730.
- [16] M.X. Fernandes, J.G. De la Torre, Brownian dynamics simulation of rigid particles of arbitrary shape in external fields, *Biophys. J.* 83 (2002) 3039–3048.
- [17] M. Zurita-Gotor, J. Bławdziewicz, E. Wajnryb, Motion of a rod-like particle between parallel walls with application to suspension rheology, *J. Rheol.* 51 (2007) 71–97.
- [18] R. Kutteh, Rigid body dynamics approach to Stokesian dynamics simulations of nonspherical particles, *J. Chem. Phys.* 132 (2010).
- [19] A. Ortega, D. Amorós, J. García De La Torre, Prediction of hydrodynamic and other solution properties of rigid proteins from atomic- and residue-level models, *Biophys. J.* 101 (2011) 892–898.
- [20] J.J. Molina, R. Yamamoto, Direct numerical simulations of rigid body dispersions. I. Mobility/friction tensors of assemblies of spheres, *J. Chem. Phys.* 139 (2013).
- [21] S. Poblete, A. Wysocki, G. Gompper, R.G. Winkler, Hydrodynamics of discrete-particle models of spherical colloids: a multiparticle collision dynamics simulation study, *Phys. Rev. E* 90 (2014) 1–13.
- [22] M. Dhugosz, J.M. Antosiewicz, Toward an accurate modeling of hydrodynamic effects on the translational and rotational dynamics of biomolecules in many-body systems, *J. Phys. Chem.* 119 (2015) 8425–8439.
- [23] A. Pandey, P.B. Sunil Kumar, R. Adhikari, Flow-induced nonequilibrium self-assembly in suspensions of stiff, apolar, active filaments, *Soft Matter* 12 (2016) 9068–9076.
- [24] S. Delong, F. Balboa Usabiaga, A. Donev, Brownian dynamics of confined rigid bodies, *J. Chem. Phys.* 143 (2015).
- [25] F.B. Usabiaga, B. Delmotte, A numerical method for suspensions of articulated bodies in viscous flows, arXiv pre-print, <https://doi.org/10.48550/ARXIV.2107.10800>, 2021.
- [26] E.B. van der Wee, B.C. Blackwell, F.B. Usabiaga, A.V. Sokolov, I.T. Katz, B. Delmotte, M.M. Driscoll, A Simple Catch: Thermal Fluctuations Enable Hydrodynamic Trapping of Microrollers by Obstacles, 2022.

- [27] Q. Brosseau, F.B. Usabiaga, E. Lushi, Y. Wu, L. Ristroph, J. Zhang, M. Ward, M.J. Shelley, Relating rheotaxis and hydrodynamic actuation using asymmetric gold-platinum phoretic rods, *Phys. Rev. Lett.* 123 (2019) 178004.
- [28] Q. Brosseau, F.B. Usabiaga, E. Lushi, Y. Wu, L. Ristroph, M.D. Ward, M.J. Shelley, J. Zhang, Metallic microswimmers driven up the wall by gravity, *Soft Matter* 17 (2021) 6597–6602.
- [29] J.W. Swan, J.F. Brady, R.S. Moore, Modeling hydrodynamic self-propulsion with Stokesian dynamics. Or teaching Stokesian dynamics to swim, *Phys. Fluids* 23 (2011).
- [30] A.M. Fiore, J.W. Swan, Rapid sampling of stochastic displacements in Brownian dynamics simulations with stresslet constraints, *J. Chem. Phys.* 148 (2018).
- [31] A.M. Fiore, J.W. Swan, Fast Stokesian dynamics, *J. Fluid Mech.* 878 (2019) 544–597.
- [32] A. Vázquez-Quesada, F. Balboa Usabiaga, R. Delgado-Buscalioni, A multiblob approach to colloidal hydrodynamics with inherent lubrication, *J. Chem. Phys.* 141 (2014).
- [33] J. Bagge, A.K. Tornberg, Highly accurate special quadrature methods for Stokesian particle suspensions in confined geometries, *Int. J. Numer. Methods Fluids* (2021) 1–50.
- [34] L. af Klinteberg, A.K. Tornberg, A fast integral equation method for solid particles in viscous flow using quadrature by expansion, *J. Comput. Phys.* 326 (2016) 420–445.
- [35] E. Corona, L. Greengard, M. Rachh, S. Veerapaneni, An integral equation formulation for rigid bodies in Stokes flow in three dimensions, *J. Comput. Phys.* 332 (2017) 504–519.
- [36] M. Maxey, Simulation methods for particulate flows and concentrated suspensions, *Annu. Rev. Fluid Mech.* 49 (2017) 171–193.
- [37] R. Cortez, L. Fauci, A. Medovikov, The method of regularized stokeslets in three dimensions: analysis, validation, and application to helical swimming, *Phys. Fluids* 17 (2005).
- [38] J.F. Brady, L. Durlafsky, G. Bossis, Dynamic simulation of hydrodynamically interacting particles, *J. Fluid Mech.* 180 (1987) 21–49.
- [39] J. Brady, G. Bossis, Stokesian dynamics, *Annu. Rev. Fluid Mech.* 20 (1988) 111–157.
- [40] J.F. Brady, A. Sierou, Accelerated Stokesian dynamics simulations, *J. Fluid Mech.* 448 (2001) 115–146.
- [41] J.W. Swan, G. Wang, Rapid calculation of hydrodynamic and transport properties in concentrated solutions of colloidal particles and macromolecules, *Phys. Fluids* 28 (2016).
- [42] D.J. Jeffrey, Y. Onishi, Calculation of the resistance and mobility functions for two unequal rigid spheres in low-Reynolds-number flow, *J. Fluid Mech.* 139 (1984) 261–290.
- [43] H.J. Wilson, Stokes flow past three spheres, *J. Comput. Phys.* 245 (2013) 302–316.
- [44] A. Lefebvre-Lepot, B. Merlet, T.N. Nguyen, An accurate method to include lubrication forces in numerical simulations of dense Stokesian suspensions, *J. Fluid Mech.* 769 (2015) 369–386.
- [45] E. Wajnryb, K.A. Mizerski, P.J. Zuk, P. Szymczak, Generalization of the Rotne-Prager-Yamakawa mobility and shear disturbance tensors, *J. Fluid Mech.* 731 (2013).
- [46] W.P. Breugem, A second-order accurate immersed boundary method for fully resolved simulations of particle-laden flows, *J. Comput. Phys.* 231 (2012) 4469–4498.
- [47] B. Sprinkle, E.B. Van Der Wee, Y. Luo, M.M. Driscoll, A. Donev, Driven dynamics in dense suspensions of microrollers, *Soft Matter* 16 (2020) 7982–8001.
- [48] Y.G. Tao, W.K. Den Otter, J.T. Padding, J.K. Dhont, W.J. Briels, Brownian dynamics simulations of the self- and collective rotational diffusion coefficients of rigid long thin rods, *J. Chem. Phys.* 122 (2005).
- [49] J.K. Dhont, *An Introduction to Dynamics of Colloids*, vol. 2, Elsevier Science, 1996.
- [50] M. Benzi, G.H. Golub, J. Liesen, Numerical solution of saddle point problems, *Acta Numer.* 14 (2005) 1–137.
- [51] J.F. Brady, I.L. Claey, Suspensions of prolate spheroids in Stokes flow. Part 1. Dynamics of a finite number of particles in an unbounded fluid, *J. Fluid Mech.* 251 (1993) 411–442.
- [52] M. Reichert, *Hydrodynamic interactions: Theoretical concepts and simulation methods*, Ph.D. thesis, University of Konstanz, 2006.
- [53] M. Orsi, L. Lobry, F. Peters, Frame-invariant sub-grid corrections to the fictitious domain method for the simulation of particulate suspensions in nonlinear flows using openfoam, *J. Comput. Phys.* 474 (2023).
- [54] F. Fryklund, E. Lehto, A.K. Tornberg, Partition of unity extension of functions on complex domains, *J. Comput. Phys.* 375 (2018) 57–79.
- [55] R. Cavoretto, Adaptive radial basis function partition of unity interpolation: a bivariate algorithm for unstructured data, *J. Sci. Comput.* 87 (2021).