# A barrier method for contact avoiding particles in Stokes flow 

Anna Broms *, Anna-Karin Tornberg<br>KTH Royal Institute of Technology, Lindstedtsvägen 25, 11428 Stockholm, Sweden

## A R T I C L E I N F O

## Keywords:

Stokes flow
Contact problem
Rigid particles
Barrier method
Constrained minimisation


#### Abstract

Rigid particles in a Stokesian fluid experience an increasingly strong lubrication resistance as particle gaps narrow. Numerically, resolving these lubrication forces comes at an intractably large cost, even for moderate system sizes. Hence, it can typically not be guaranteed that artificial particle collisions and overlaps do not occur in a dynamic simulation, independently of the choice of method to solve the Stokes equations. In this work, the potentially large set of non-overlap constraints, in terms of the Euclidean distance between boundary points on disjoint particles, are efficiently represented via a barrier energy. We solve for the minimum magnitudes of repelling contact forces and torques between any particle pair in contact to correct for overlaps by enforcing a zero barrier energy at the next time level, given a contact-free configuration at a previous instance in time. Robustness for the method is illustrated using a multiblob method to solve the mobility problem in Stokes flow, applied to suspensions of spheres, rods and boomerang shaped particles. Collision free configurations are obtained at all instances in time, and considerably larger time-steps can be taken than without the technique. The effect of the contact forces on the collective order of a set of rods in a background flow that naturally promote particle interactions is also illustrated.


## 1. Introduction

We present an algorithm for contact forces and torques between rigid particles with smooth surfaces immersed in an unbounded viscous fluid in 3D. The forces and torques are introduced only when needed to avoid particle collisions and overlaps. In this paper, we consider the ideal case in the zero Reynolds number limit, where the Stokes equations describe the fluid flow and there are no additional interaction forces between particles. Full hydrodynamic interaction is accounted for, meaning that there is a global coupling between all the particles in the system. In time dependent simulations of such systems, artificial collisions commonly occur due to hard-to-avoid numerical artefacts and a contact algorithm has to be introduced to avoid such collisions.

The hydrodynamic interactions are computed with a cheap and fast technique: the so called rigid multiblob method, as carefully described in $[1,2]$. Note that this is only one of many numerical methods for Stokes flows, and the contact avoiding strategy that we develop can be used also in combination with other techniques. The multiblob method is based on the idea that a rigid body can be modelled by a collection of spheres or "blobs" and has been used in a very large number of works, see e.g. [3-7] and references therein. Each blob interacts hydrodynamically with all other blobs in the system in a pairwise manner and blobs belonging to the same particle are constrained to move as a rigid body via forces applied at each blob centre, such that the blob forces sum to the net force and torque on the particle. Mathematically, one could view this as a regularised single layer boundary integral formulation,

[^0]

Fig. 1. Multiblob particles in a Stokesian fluid. Colours indicate the depth in the figures. (For interpretation of the colours in the figure(s), the reader is referred to the web version of this article.)
with the blob radius the regularisation parameter. Example multiblob geometries are displayed in Fig. 1. For each instance in time, we solve the Stokes mobility problem, that is, given assigned external forces and torques on the particles, such as e.g. gravity or some electrostatic forcing, the resulting particle translational and rotational velocities are computed. The inertia of the particles is typically negligible in the Stokes regime, and hence, the computed rigid body velocities can be used to update the particle positions, applying some suitable time-stepping scheme, and the configuration of particles can in this way be studied dynamically. By choosing the regularisation parameter and the surface where blobs are placed in relation to the true surface of the particle as the solution to a small off-line optimisation problem for each (axisymmetric) particle type, good accuracy in the particle velocities can be obtained for moderately separated particles even with coarse grids of the particle surfaces [2].

As particles get closer, however, their interactions become increasingly difficult to accurately resolve. In the multiblob framework, accuracy is suffering for particles close to being in contact. Note, though, that this is a challenge independently of the choice of Stokes solver. Insufficient treatment of close interactions, caused either by numerical errors due to insufficient spatial resolution of the particle grid or by accumulated errors from the time-stepping scheme of the moving particles, can in the worst case lead to particle collisions and overlaps not dictated by the hydrodynamic equations. Even if close interactions are resolved, and an adaptive time-stepping scheme is used, another unwanted effect is stalling, where the adaptive time-step size becomes increasingly small as the time-step is adjusted to the stiff problem of particles in close proximity affected by increasingly strong repulsive forces.

There is a discussion in the literature concerning the applicability of the Stokes equations in specific scenarios, such as when dealing with dense particle suspensions undergoing prolonged shear. Such scenarios would allow for inter-particle gaps of subatomic scales - scales where the continuum approximation is no longer valid [8-11]. Hence, in such cases, it is questionable if it is meaningful to strive to accurately continue to resolve the Stokes equations as gaps narrow. In physical systems, additional short-range interaction forces would typically be present e.g. due to polymer coats or charges on the particle surfaces [10,11]. Under the same conditions as smooth particles stay apart, surface roughness or corners could introduce contact through a rupture of the lubrication layer, such that frictional forces become important. The details of contact between particles in sheared suspensions remain poorly understood, but it is clear that interactions over microscopic length scales in the gap between particles have a pronounced impact on the rheology of concentrated suspensions, such as the phenomenon of shear thickening [11].

If one desires to perform simulations based on the Stokes equations without inclusion of additional interaction forces, it is inevitable that small inter-particle distances will arise, and with this, large lubrication forces that can only be resolved using high fidelity methods with fine spatial discretisation of the surfaces of the particles [12]. Both accurate temporal and spatial discretisation however come at a large computational cost, while underresolved simulations in general will result in artificial collisions and overlaps. Enforcing a minimum distance between particles within the system alleviates the resolution requirements, making simulations more computationally efficient. To achieve this, the implementation of a contact avoiding strategy is needed.

As many numerical methods used for solving the Stokes equations break when particles overlap, a contact avoiding strategy is also instrumental in enabling the exploration of particle suspensions over extended time periods. This is essential to facilitate the collection of data for the study of a diverse range of both linear and nonlinear phenomena. These phenomena encompass aspects such as transport and mixing, effective viscosity in sheared suspensions of particles of varying shapes and densities [13], phase transitions in liquid chrystals [14], or rotational diffusion over long time-scales in systems with Brownian motion imperative to understand collective order and self-assembly for biological or synthetic particles [15].

Contact forces and torques introduced as part of the contact avoiding strategy are artificial to the system and introduced only to decrease the impact from the equally artificial numerical errors difficult to avoid for particles in close proximity. Hence, it is desired to find the smallest possible contact forces for particles to stay apart. At the same time, we also want to avoid introducing stiffness in the system, and eliminate the risk of having to take very small time-steps. Such stiffness is often the drawback of applying a strongly repelling potential to avoid particle overlaps, as an alternative to contact resolution algorithms. Potentials of this type include Lennard-Jones or e.g. a potential based on Hard Gaussian Overlap (HGO), which is designed for ellipsoids [16] and discussed
for rods in [17]. ${ }^{1}$ Due to the problem with stiffness, it is difficult to guarantee non-overlapping particles with a potential-based method, and there is a risk that particles become "soft" [18].

Contact problems with particles without hydrodynamic interactions have a richer literature [19-23] with more benchmarks available. An overview of techniques for contact dynamic problems is found in [24]. There is also a body of work in computer graphics [25,26], and of most relevance to this work, the recently introduced method of Incremental Potential Contact (IPC), which is a penalty method where colliding and overlapping configurations are penalised by a barrier energy [27-29].

Different contact resolution techniques have been suggested to better resolve particle collisions in Stokes flow, see e.g. the works [30-34], where no-slip boundary conditions are imposed at the point of contact, constraining the colliding particles to move with equal speed at the collision point. A complementarity formulation is favoured in [35,14,36-39], where the idea is to solve a nonlinear complementarity problem for contact force and torque magnitudes, given some definition of particle separation. The difficult-to-solve nonlinear problem is in turn approximated by one or a sequence of linear complementarity problems (LCPs) that can be solved more easily. One method in this class is presented in the works by Yan et al., [14,36,38], and has the advantages that it is easily applicable to any hydrodynamic solver and is based on a geometric formulation that fulfils Newton's third law (forces on every pair of particles are balanced). In addition, the collision stress and mechanical pressure can be computed directly from the contact forces. One drawback, however, is that the method cannot guarantee non-overlapping configurations at the end of each time-step, as the solution of a single LCP does not imply a solution to the original nonlinear problem. A second drawback is that non-convex particles cannot be handled. Another method of the same flavour for Stokesian fluids is based on so called Space-Time Interference Volumes (STIV), as in Lu et al. [35,40] and Bystricky et al. [37]. The STIV technique is inspired by the work of Harmon et al. [26] in computer graphics. The general idea is to consider the Stokes equations in variational form and, after a candidate time-step, compute the volume in space-time swept out by the trajectories of the particles in contact. If the volume is negative, particles overlap. Enforcing this volume to be zero, by determining appropriate repulsion forces, gives rise to a complementarity problem. In an STIV approach, no-collision is obtained even for large time-steps by solving sequences of LCPs and particles that pass through each other during one time-step can be detected and the corresponding time-step corrected. However, each STIV-formulation is strongly linked to a specific Stokes solver, the method is not easily applicable to general geometries, and in 3D, the interference volume has to be efficiently computed in four dimensions [40,35]. Another difficulty is that using the STIV, contact forces are not automatically balanced and balanced forces are difficult to obtain (Newton's third law is not satisfied) [41].

We choose to approach the contact problem by introducing repelling contact forces and torques to particles that come too close to each other in terms of the Euclidean distance. Key features of the method are that all contact forces are balanced via Newton's third law and that the next time-step in a sequence of time-steps is guaranteed to be contact free (in fact, by construction, a minimum separation distance is guaranteed between particles). The method is strongly inspired by the work of Yan et al. [14,36,38] in how contact forces and torques are geometrically motivated and by the work of Zorin and co-authors in the work on IPC, [27,29,28], in how non-overlapping constraints are set up. However, instead of penalising contact by a barrier energy as in [27,29,28], a barrier energy is rather used to represent both a large number of non-overlap constraints and the complementarity condition with the associated contact force/torque magnitude. In the work on IPC, [27,29,28], a tetrahedral or triangular mesh of the surface of each particle is considered, where distances are computed between all pairs of edges and points to triangles respectively. In our work, particles are rigid and smooth, and we can utilise the known parameterisation of the particle surfaces to compute distances robustly; given a set of points defining the grid of a particle surface, and for each surface grid node within some set threshold of the other particle, we flag the closest point of contact on the neighbouring particle to be part of the collision handling. This is possible to do as we are not dependent on second order derivatives of the distances in solving the optimisation problem with our formulation, in contrast to IPC; gradients of the distances with respect to particle coordinates suffice. The handling of the geometry and the constraints is also in contrast to the geometric approach by Yan et al., where only a single point of contact or "the most overlapping point" has to be determined, and more similar to the STIV technique, where multiple segments of the surface can be flagged for the same particle contact pair. This is a beneficial property especially for non-convex particle geometries, or if surfaces are close to parallel, where the computed contact torque otherwise becomes very sensitive to the choice of contact point. Details are outlined in Section 2.1. Other differences in our work compared to [27,29,28] are that particles are immersed in a Stokesian fluid, where inertia is negligible, and we solve for contact force/torque magnitudes. For rigid bodies with IPC, on the other hand, one solves for the particle coordinates in an implicit time-step [28]. The dimension of the optimisation variable in the resulting optimisation problem hence becomes smaller here than in IPC. Keeping the rigidity of the particles is a non-issue (no additional constraints have to be enforced, in contrast to what is discussed in [28]). We utilise the linearity of the Stokes equations and formulate a minimisation problem for contact force magnitudes to be solved in every time-step where contact occurs. Note specifically that we do not minimise a combined, weighted, energy formulation as in IPC [27,29,28]. Hence, there are no parameters to tune in our formulation for the importance of the collision constraints relative to the minimisation of other contributions to the total energy of the system.

The time-stepping schemes used both for STIV and the geometric approach by Yan et al. are explicit (to avoid a large cost at every time-step) and of low order [42,14,36,38]. The repulsion forces and torques are assumed to be constant over the course of one time-step, and the basis for both types of contact algorithms in their vanilla version is explicit Euler. This is the time-stepping method that will be used for demonstration also in this work. An additional motivation to this choice is in settings where the contact resolution algorithm is coupled to Brownian motion, modelled by a stochastic differential equation, where it is difficult to obtain

[^1]anything better than first order accuracy in time. Note however that there is nothing that prevents a higher order time-stepping scheme from being used if contact avoiding is applied in a deterministic setting. The most straight-forward idea is then to use a higher order time-stepping method for the particles not involved in any contacts. For particles to which repulsive forcing is added in a certain time-step, there is no expected gain in using a high order method.

### 1.1. The Stokes mobility problem

Before introducing the optimisation problem, we start with some preliminaries. Each 3D particle in the fluid suspension can be described by its centre coordinates $\boldsymbol{x}_{i} \in \mathbb{R}^{3}$ and rotation quaternion ${ }^{2} \boldsymbol{q}_{i} \in \mathbb{R}^{4}$. We collect these generalised coordinates for all the $N$ particles in the system in the vector $\boldsymbol{Q}$ such that

$$
\begin{equation*}
\boldsymbol{Q}=\left[x_{1}^{T}, \boldsymbol{q}_{1}^{T}, \boldsymbol{x}_{2}^{T}, \boldsymbol{q}_{2}^{T}, \ldots, \boldsymbol{x}_{N}^{T}, \boldsymbol{q}_{N}^{T}\right]^{T} \tag{1}
\end{equation*}
$$

Let $\mathcal{U}$ be a vector of all rigid body velocities of the particles in the system, where $\boldsymbol{u}_{i} \in \mathbb{R}^{3}$ is the translational velocity and $\omega_{i} \in \mathbb{R}^{3}$ the rotational velocity of particle $i$. Similarly, let $\mathcal{F}_{\text {ext }}$ be a vector of all the externally applied forces $\boldsymbol{f}_{i} \in \mathbb{R}^{3}$ and torques $\boldsymbol{t}_{i} \in \mathbb{R}^{3}$ on the particles in the system:

$$
\boldsymbol{\mathcal { U }}=\left[\begin{array}{lllllll}
\boldsymbol{u}_{1}^{T} & \boldsymbol{\omega}_{1}^{T} & \boldsymbol{u}_{2}^{T} & \boldsymbol{\omega}_{2}^{T} & \ldots & \boldsymbol{u}_{N}^{T} & \boldsymbol{\omega}_{N}^{T}
\end{array}\right]^{T}, \quad \mathcal{F}_{\mathrm{ext}}=\left[\begin{array}{lllllll}
\boldsymbol{f}_{1}^{T} & \boldsymbol{t}_{1}^{T} & \boldsymbol{f}_{2}^{T} & \boldsymbol{t}_{2}^{T} & \ldots & \boldsymbol{f}_{N}^{T} & \boldsymbol{t}_{N}^{T} \tag{2}
\end{array}\right]^{T} .
$$

Now, the Stokes mobility problem can be stated as

$$
\begin{equation*}
\boldsymbol{\mathcal { U }}=\boldsymbol{\mathcal { U }}_{\mathrm{bg}}+\boldsymbol{\mathcal { U }}_{\text {Brownian }}+\boldsymbol{\mathcal { M }} \mathcal{F}_{\mathrm{ext}} \tag{3}
\end{equation*}
$$

where $\mathcal{M}$ is the mobility matrix of size $6 N \times 6 N$, with each $6 \times 6$ block corresponding to the interaction between a specific pair of particles. The vector $\boldsymbol{\mathcal { U }}_{\mathrm{bg}}$ is the velocity contribution to the particles from an eventual background flow and $\boldsymbol{\mathcal { U }}_{\text {Brownian }}$ is a vector of stochastic velocities included if thermal fluctuations are considered. Note that implementation-wise, $\boldsymbol{\mathcal { M }}$ is only to be interpreted symbolically and $\boldsymbol{\mathcal { M }} \mathcal{F}_{\text {ext }}$ corresponds to solving Stokes equations along with no-slip boundary conditions using the multiblob method as described in [2], given forces and torques for all particles, stacked in $\mathcal{F}_{\text {ext }}$. Hence, $\boldsymbol{\mathcal { M }} \mathcal{F}_{\text {ext }}$ is the vector of rigid body velocities that results from $\mathcal{F}_{\text {ext }}$.

With no contact forces present, $\boldsymbol{Q}$ is related to the velocities in $\boldsymbol{\mathcal { U }}$ as

$$
\begin{equation*}
\dot{Q}=\Psi(Q) \mathcal{U} \tag{4}
\end{equation*}
$$

with $\Psi(Q) \in \mathbb{R}^{7 N} \times \mathbb{R}^{6 N}$ a geometry-dependent matrix relating velocities to particle positions and quaternions. To be more specific, the role of $\Psi(Q)$ is mainly to relate quaternion updates to angular velocities. The matrix $\Psi(Q)$ only contains non-zero blocks on the diagonal and for particle $i$, this block takes the form ${ }^{2}$

$$
\left[\begin{array}{cc}
\boldsymbol{I} & \mathbf{0}  \tag{5}\\
\mathbf{0} & \frac{1}{2} \boldsymbol{Q}\left(\boldsymbol{q}_{i}\right)
\end{array}\right] \text {, where } \boldsymbol{Q}\left(\boldsymbol{q}_{i}\right)=\left[\begin{array}{ccc}
-p_{i}^{1} & -p_{i}^{2} & -p_{i}^{3} \\
s_{i} & p_{i}^{3} & -p_{i}^{2} \\
-p_{i}^{3} & s_{i} & p_{i}^{1} \\
p_{i}^{2} & -p_{i}^{1} & s_{i}
\end{array}\right]
$$

The role of the matrix $\boldsymbol{Q}\left(\boldsymbol{q}_{\boldsymbol{i}}\right)$ is to locally preserve the norm of the unit quaternion $\boldsymbol{q}_{\boldsymbol{i}}[3,43]$.
When particles are in contact, the velocities are corrected using the action of a vector of contact forces and torques $\mathcal{F}_{c}$, such that

$$
\begin{equation*}
\dot{\boldsymbol{Q}}=\boldsymbol{\Psi} \mathcal{U}+\boldsymbol{\Psi} \boldsymbol{\mathcal { M }} \mathcal{F}_{c}, \tag{6}
\end{equation*}
$$

with

$$
\begin{equation*}
\mathcal{F}_{c}=\mathcal{D} \lambda . \tag{7}
\end{equation*}
$$

The matrix $\mathcal{D} \in \mathbb{R}^{6 N} \times \mathbb{R}^{N_{c}}$ determines the direction of contact forces and torques and depends on the particle configuration and geometries and the vector $\lambda \in \mathbb{R}^{N_{c}}$ determines the contact force/torque magnitudes, with $N_{c}$ the number of particle pairs in contact. From now on, we sometimes for simplicity write only "contact force". This should always be read as a contact force/torque pair. Ideally, we would like to determine contact forces such that if contact forces are needed for particles to stay apart during the timestep, the force should be active and the corresponding force magnitude $\lambda_{i}$ positive. On the other hand, if the particles stay separate without contact forces, the contact force magnitude $\lambda_{i}$ should be zero and the contact force passive. The nature of the contact problem is hence a complementarity problem. It is also a nonlinear problem, as the distances at the next instance of time depend nonlinearly on the contact force magnitudes $\lambda$.

[^2]
## 2. A barrier method for particles in contact

In practice, it is not advisable to let particles come arbitrarily close to each other, due to limitations in all numerical solvers for particles in Stokes flow - the accuracy for almost touching particles is suffering if the resolution of the particles is not excessively high. For this reason, we define a contact to occur if the distance $d$ between a pair of particles is smaller than a set buffer distance $\hat{d}$. The parameter $\hat{d}$ is typically set to ensure a prescribed accepted accuracy in the Stokes solver. ${ }^{3}$ We want to find contact force magnitudes $\lambda$ such that this separation distance is guaranteed at the next time-step, i.e.

$$
\begin{equation*}
d_{i}\left(Q_{t+\Delta t}(\lambda)\right)>\hat{d}, \quad \text { for all } i \text { distances measured between distinct particles, } \tag{8}
\end{equation*}
$$

with $d_{i}$ depending non-linearly on $\boldsymbol{Q}_{t+\Delta t}$, the particle configuration in the next time step. The set of relevant distances for each particle pair is properly defined in Section 2.1. The complementarity conditions for the particle distances and contact force magnitudes take the form

$$
\begin{equation*}
\lambda_{k} \max \left(0, \min _{i}\left\{d_{i}\left(\boldsymbol{Q}_{t+\Delta t}\right)-\hat{d}\right\}\right)=0, \quad i \text { associated with particle contact pair } k \text { and } \lambda_{k} \geq 0 . \tag{9}
\end{equation*}
$$

Instead of solving this (modified) nonlinear complementarity problem exactly, we will set up an optimisation problem to fulfil all constraints in (8) and still approximately solve (9). Note that depending on the number and concentration of particles in the system and their geometries, the number of constraints in (8) might be very large. This section explains how an optimisation problem can be formulated and discusses modelling choices for the objective function, the constraints and the direction of the contact forces and torques. The constraints in (8) can be represented via a minimisation of a sum of indicator functions [44]:

$$
\min _{i \geq 0} \quad \sum_{i} I\left(d_{i}\left(Q_{t+\Delta t}(\lambda)\right)-\hat{d}\right), \quad \text { with } I(s)=\left\{\begin{align*}
0, & s \geq 0  \tag{10}\\
\infty, & s<0
\end{align*}\right.
$$

This representation is however difficult to work with, and we therefore replace (10) with a smoother counterpart. The formulation is based on a so-called barrier function, $b$, that is zero for sufficiently large distances $d$ and increasingly large for distances smaller than the set threshold $\hat{d}$ :

$$
b(d, \hat{d})=\left\{\begin{align*}
-(d-\hat{d})^{2} \ln (d / \hat{d}), & 0<d<\hat{d}  \tag{11}\\
0, & d \geq \hat{d}
\end{align*}\right.
$$

We can collect all the non-overlapping constraints in (8) into what we define as a barrier energy, mimicking the sum of indicators in (10). A non-overlapping configuration is one with non-negative contact force magnitudes $\lambda$ applied for $t \in[t, t+\Delta t]$ that minimises the barrier energy,

$$
\begin{equation*}
\min _{\lambda \geq \mathbf{0}} \sum_{i} b\left(d_{i}\left(\boldsymbol{Q}_{t+\Delta t}(\lambda)\right), \hat{d}\right) \tag{12}
\end{equation*}
$$

with the summation being over all geometries in contact (corresponding to all constraints in (8)). In a Stokesian fluid, the particle coordinates then evolve with time according to

$$
\begin{equation*}
\dot{\boldsymbol{Q}}=\boldsymbol{\Psi}(\boldsymbol{U}+\mathcal{M} \mathcal{D} \lambda) \tag{13}
\end{equation*}
$$

where $\mathcal{D} \lambda$ is a force/torque vector and $\mathcal{M} \mathcal{D} \lambda$ results in a velocity vector, as in (6). If we discretise this equation in time, we can express $Q_{t+\Delta t}$ in terms of previous, known, coordinate vectors. With the forward Euler method, the minimisation problem becomes

$$
\begin{array}{ll}
\min _{\lambda \geq \mathbf{0}} & \sum_{i} b\left(d_{i}\left(\boldsymbol{Q}_{t+\Delta t}(\lambda)\right), \hat{d}\right),  \tag{14}\\
\text { s.t. } & \boldsymbol{Q}_{t+\Delta t}=\boldsymbol{Q}_{t}+\Delta t \boldsymbol{\Psi}(\boldsymbol{U}+\boldsymbol{M} \mathcal{D} \lambda)
\end{array}
$$

This is a constrained minimisation problem to solve for $\lambda \in \mathbb{R}^{N_{c}}$, i.e. with a scalar contact force/torque magnitude per particle pair in contact. Note however that for some particle configurations, any vector $\lambda$ with sufficiently large magnitude would solve (14). We hence need to constrain the solution by penalising large $\lambda$, with a natural choice being

$$
\begin{equation*}
\min _{\lambda \geq \mathbf{0}}\|\lambda\|+\alpha \sum_{i} b\left(d _ { i } \left(\boldsymbol{Q}_{t}+\Delta t \boldsymbol{\Psi}(\boldsymbol{\mathcal { U }}+\boldsymbol{\mathcal { M } \mathcal { D } \lambda ) ) , \hat { d } ) , ~ , ~ . ~}\right.\right. \tag{15}
\end{equation*}
$$

where $\alpha$ is a parameter that has to be set to balance the minimisation of the barrier energy at the next time-step and the minimisation of contact forces. The parameter $\alpha$ also has the role of enforcing the complementarity condition between the contact forces and contact distances (relative to the buffer region) in (9). A different way of saying the same thing is that we would like to pick $\lambda_{k}$ sufficiently small so that non-overlap is obtained, but where the force is active, the minimum contact distance should not be larger than $\hat{d}$ after applying the contact forces.

[^3]For the choice of norm in (15), we have two options: A 1-norm penalty of the contact force magnitudes, $\sum_{i} \lambda_{i}$, penalises the forces at all contacts to an equal extent and not only those where the magnitude is large. One could also choose the norm $\left(\lambda^{T} \mathcal{D}^{T} \mathcal{M} \mathcal{D} \lambda\right)^{1 / 2}$, or its square, corresponding to the dissipated energy induced by the contact forces. As a main goal is to minimise the impact of the artificial repulsive forces on the system, the dissipated energy due to the introduced contact forces could physically be a reasonable choice of objective function. For our Stokes solver, the rigid multiblob method, it is however a known problem that the accuracy in $\boldsymbol{\mathcal { M }}$ is suffering a lot for closely interacting particles [2]. The same holds true for many numerical methods to solve the Stokes equations. Hence, $\lambda \mathcal{D}^{T} \mathcal{M} \mathcal{D} \lambda$ is a very crude approximation to the true dissipated energy and a large error is associated with this quantity. For this reason, we will use $\sum_{i} \lambda_{i}$ as our objective function of choice. See section 3.1.2 for a numerical comparison of the two options that will further motivate our choice.

Note that it is not trivial to find the best choice of $\alpha$ in (15) for general applicability, as the magnitude of both $\lambda$ and the barrier energy depends on the number of pairs of particles in contact and how much overlap a non-corrected time-step would yield. A suitable level of $\alpha$ also depends on the tolerance chosen as stopping criterion when solving the problem in (15). This is a similar problem as the one encountered when minimising the energy formulation in the work on IPC, where a penalty has to be chosen in an adaptive manner to give the barrier energy the proper weight; see the supplementary in [27]. Here, we would like to avoid such hyperparameter tuning. What is known, is that the barrier energy is identically zero for a non-overlapping configuration. An alternative formulation of the minimisation problem in (15) is therefore

$$
\begin{array}{ll}
\min _{i \geq 0} & \sum_{i} \lambda_{i} \\
\text { s.t. } & \sum_{i} b\left(d_{i}\left(\boldsymbol{Q}_{t}+\Delta t \boldsymbol{\Psi}(\boldsymbol{U}+\mathcal{M D} \lambda)\right), \hat{d}\right)=0 \tag{16}
\end{array}
$$

where contact force magnitudes are minimised while enforcing a zero barrier energy. The formulation in (16) is simpler than (15) in the sense that one hyperparameter is reduced, but possibly more challenging as a nonlinear equality constraint has been added. Note that the Lagrangians of (15) and (16) are identical, but in the case of (16), the parameter $\alpha$ is a Lagrange multiplier to be solved for instead of a set parameter. In the remaining of this work, we will consider the formulation in (16) for determining contact forces.

Remark 1. In contrast to in the work on IPC by Zorin et al., the value of the barrier energy itself is not used explicitly in the formulation presented in this paper but necessary for

1. Imposing a large number of constraints on minimum distances as a collapsed single constraint. This allows us to seek minimum force/torque magnitudes $\lambda$ only where the barrier energy is zero. What is gained is that we do not have to deal with the large number of original constraints explicitly, and are hence able to solve smaller optimisation problems in every time-step where contact occurs, with just a single dual variable instead of one per colliding particle pair.
2. Obtaining a direction of the contact force and torque that takes more information into account: by using the barrier energy, a large number of pairs of contact points are weighted by their distance to give a net force and torque direction (more on this in Section 2.1). This is more robust than considering only one pair of points per contact pair.

An alternative formulation for the same set of constraints is $\min _{i} d_{i} \geq \hat{d}$. For a large class of optimisation methods, the interior point methods, such a problem requires a feasible starting guess, $\lambda \geq \mathbf{0}$ with all non-overlap constraints satisfied, which might be very hard to obtain even for systems containing a few particles. Rewriting the non-overlap constraints in (8) as a barrier energy is hence necessary for robustness.

Remark 2. Care has to be taken in the forward Euler update in (14) for rotational degrees of freedom. Quaternions have to be updated by means of rotations to guarantee that their unit norm is preserved. As a linearisation, (14) holds, but practically, the angular update $\Delta t \omega_{i}$ is expressed as a unit quaternion of the form

$$
\begin{equation*}
\left[\cos \frac{\Delta t\left\|\omega_{i}\right\|}{2}, \frac{\omega_{i}}{\left\|\omega_{i}\right\|} \sin \frac{\Delta t\left\|\omega_{i}\right\|}{2}\right]^{T} \tag{17}
\end{equation*}
$$

The updated quaternion $\boldsymbol{q}_{i t+\Delta t}$ is obtained via a rotation of the old quaternion $\boldsymbol{q}_{i t}$ so that

$$
\begin{equation*}
\boldsymbol{q}_{i t+\Delta t}=\operatorname{Rot}\left(\Delta t\left\|\omega_{i}\right\|, \omega_{i}\right) \boldsymbol{q}_{i t} \tag{18}
\end{equation*}
$$

where $\boldsymbol{\operatorname { R o t }}(\boldsymbol{q})$ represents a quaternion rotation from the left [3,43], and, with $\boldsymbol{q}:=(s, \boldsymbol{p})$,

$$
\boldsymbol{\operatorname { R o t }}(\boldsymbol{q})=\left[\begin{array}{cccc}
s & -p^{1} & -p^{2} & -p^{3}  \tag{19}\\
p^{1} & s & -p^{3} & p^{2} \\
p^{2} & p^{3} & s & -p^{1} \\
p^{3} & -p^{2} & p^{1} & s
\end{array}\right]
$$



Fig. 2. Illustration of the two types of contact distances, $d_{i}^{\mathrm{p}}$ and $d_{i k}^{\mathrm{s}}$, for two fat rods close to contact in pair $i$. Surface-node-to-particle distances $d_{i k}^{\mathrm{s}}$ are the distances from each surface grid node on one particle (black dots) to the closest point on the other particle, less than a set threshold (buffer region indicated in yellow). Closest points are determined from the shortest distance to the centre line segment of the other particle (blue line). Pairs of grid nodes and computed points are marked with red dots, with red lines drawn in between displaying the distances. The shortest particle-particle distance $d_{i}^{\mathrm{p}}$, on the other hand, is given by the length of the cyan line drawn between the closest points on the two surfaces, computed from the closest distance between the two centre line segments, and subtracting $2 R_{\text {rod }}$.

### 2.1. Defining the contact distance and contact force

Let us now focus on the Euclidean contact distances $d_{i}$ between points on the surfaces of two particles in close proximity. We make a distinction between $d^{\mathrm{p}}$ and $d^{\mathrm{s}}$, with $d^{\mathrm{p}}$ the shortest distance between particles and $d^{\mathrm{s}}$ the shortest distance between a pair of surface points on two particles. In principle, $d_{i}$ in the constrained optimisation problem (16) can denote either $d_{i}^{\mathrm{p}}$ or $d_{i}^{\mathrm{s}}$. Practically, we compute $d^{\mathrm{p}}$ directly using known parameterisations of the particles, while the surface-node-to-particle distance $d^{s}$ is based on discretisations of the particle surfaces, described by a distribution of points. To be more precise, for each node on the surface of one particle in a pair sufficiently close to contact, we then use the known parameterisation of the centre line of the other particle to determine the closest point on this line. The particle surfaces considered in this work are all one radius away from the particle centre line, and the surface-node-to-particle distance can hence easily be computed. Fig. 2 visualises $d^{p}$ and $d^{s}$ for a configuration of two rods. Geometric considerations determine if we choose to use $d^{\mathrm{p}}$ or $d^{\mathrm{s}}$ :

1. The particle-particle distance $d^{\mathrm{p}}$ may be a good choice if the particle shapes are simple enough, such as e.g. for spheres or axisymmetric rods with semi-spherical caps. In the latter case, we solve an equation to determine the closest distance between two line segments, following [45,46], and subtract $2 R_{\text {rod }}$ to determine the closest distance between particles.
2. The surface-point-to-surface-point distance, $d^{s}$, is especially beneficial in two different cases: i.) For close to parallel surfaces, such that it is hard to define a single point of contact and ii.) For non-convex particles. We will focus most of our attention on this case. Note that a requirement for robustness is that the surface nodes well resolve the curvature of the particles.

We think of the contact forces as a (hopefully small) correction to the trial time-step $\boldsymbol{Q}_{t+\Delta t}^{*}$, i.e. the time-step without contact forces, given by

$$
\begin{equation*}
Q_{t+\Delta t}^{*}=Q_{t}+\Delta t \Psi \mathcal{U} \tag{20}
\end{equation*}
$$

Let $B_{i}\left(\boldsymbol{Q}_{t+\Delta t}^{*}, \hat{d}_{\text {try }}\right)$ be the modified barrier energy associated with the particle contact pair $i$ aggregated over all relevant particleparticle distances at the trial time-step, with the threshold $\hat{d}_{\text {try }}$ chosen so that $\hat{d}_{\text {try }}>\hat{d}$. This means that $B_{i}\left(\boldsymbol{Q}_{t+\Delta t}^{*}, \hat{d}_{\text {try }}\right)$ is the sum in (12), but for a single particle pair and with a modified threshold. Then, we define the contact force potential, $\sum_{i} \lambda_{i} B_{i}\left(Q_{t+\Delta t}^{*}, \hat{d}_{\mathrm{try}}\right)$, a weighted sum of the barrier energy in the trial time-step, and let the contact forces and torques be given by the gradient of this contact force potential. Some care has to be taken when differentiating with respect to quaternions and we therefore express $\mathcal{F}_{c}$ as

$$
\begin{equation*}
\mathcal{F}_{c}=-\boldsymbol{\Psi}^{T} \nabla_{Q_{t+\Delta t}^{*}}\left(\sum_{i} \lambda_{i} B_{i}\left(\boldsymbol{Q}_{t+\Delta t}^{*}, \hat{d}_{\mathrm{try}}\right)\right) \tag{21}
\end{equation*}
$$

The matrix $\boldsymbol{\Psi}^{T}$ here is needed to map derivatives with respect to quaternions to torques. As stated in (6)-(7), we may write the contact forces on the form $\mathcal{F}_{c}=\mathcal{D} \lambda$. Let $\boldsymbol{B} \in \mathbb{R}^{N_{c}}$ be a vector of all the modified barrier energies for the $N_{c}$ particle pairs in contact. The sparse matrix $\mathcal{D} \in \mathbb{R}^{6 N \times N_{c}}$, representing contact force and torque directions, depends on the geometry of the particles and their locations at the trial time-step. From (21) we identify that $\mathcal{D}$ is given by

$$
\begin{equation*}
\mathcal{D}=-\boldsymbol{\Psi}^{T} \nabla_{\boldsymbol{Q}_{t+\Delta t}^{*}} \boldsymbol{B}\left(\boldsymbol{Q}_{t+\Delta t}^{*}, \hat{d}_{\mathrm{try}}\right) \tag{22}
\end{equation*}
$$

The matrix has the structure

$$
\mathcal{D}=\left[\begin{array}{llll}
\boldsymbol{D}_{1} & \boldsymbol{D}_{2} & \ldots & \boldsymbol{D}_{N_{c}} \tag{23}
\end{array}\right]
$$

where $\boldsymbol{D}_{i}$ represents contact $i$, between particle $j$ and $k$, and is given by

(a) Using the surface-node-to-particle distances $d_{i k}^{s}$ : The terms in (28a) that con- (b) Using the shortest particle-particle distance, $d_{i}^{\mathrm{p}}$ : The resulting contact force tribute to the direction of the contact force are visualised with red arrows, each direction on each particle is indicated with black thick arrows and have the same with a scaling correlated with their contribution to the total force. There is one red direction as the normal drawn between the points of nearest approach (indicated arrow for each red distance line $d_{i k}^{s}$ marked in Figure 2. The resulting contact force in red).
direction on each particle in (28a) is indicated with black thick arrows.

Fig. 3. Illustration of the contact force directions corresponding to using the contact distance $d_{i k}^{s}$ or $d_{i}^{p}$ for two fat rods close to contact in pair $i$. The corresponding distances are visualised in Fig. 2. The direction of the contact torques induced by the contact forces is indicated in magenta. The directions of the contact torques are such that the particles rotate away from the point of closest approach; the rightmost particle in the pair will rotate clock-wise, while the left-most particle will rotate counter-clockwise. Shadowed arrows indicate a scaling of the contact force and torque determined by $\lambda_{i}$.

$$
\boldsymbol{D}_{i}=\left[\begin{array}{llllllllll}
0 & \ldots & 0 & \hat{\boldsymbol{f}}_{i}^{c T} & \hat{\boldsymbol{t}}_{i j}^{c T} & 0 & \ldots & 0 & -\hat{\boldsymbol{f}}_{i}^{c T} & \hat{\boldsymbol{t}}_{i k}^{c T}  \tag{24}\\
0 & \ldots & 0
\end{array}\right]^{T}
$$

Here, $\hat{\boldsymbol{f}}_{i}^{c}, \hat{\boldsymbol{t}}_{i j}^{c}$ and $\hat{\boldsymbol{t}}_{i k}^{c}$ represent force and torque directions; see [22] for a detailed motivation. For robustness, we rescale and redefine $D_{i}$ as

$$
\boldsymbol{D}_{i}=\left[\begin{array}{llllllllll}
0 & \ldots & 0 & \frac{\hat{\boldsymbol{f}}_{i}^{c T}}{\left\|\hat{\boldsymbol{f}}_{i}^{c}\right\|} & \frac{\hat{\boldsymbol{t}}_{i j}^{c T}}{\left\|\hat{\boldsymbol{f}}_{i}^{c}\right\|} & 0 & \ldots & 0 & -\frac{\hat{\boldsymbol{f}}_{i}^{c T}}{\left\|\hat{\boldsymbol{f}}_{i}^{c}\right\|} & \frac{\hat{\boldsymbol{t}}_{i k}^{c T}}{\left\|\hat{\boldsymbol{f}}_{i}^{c}\right\|}  \tag{25}\\
0 & \ldots & 0
\end{array}\right]^{T}
$$

From (25) and the relation $\mathcal{F}_{c}=\mathcal{D} \lambda$, it is clear that the forces applied to the two particles in pair $i$ are equal in magnitude (with the magnitude given by $\lambda_{i}$ ), and differ only in sign, due to Newton's third law. In the contact force potential, all particles or grid nodes closer to each other than a tolerance $\hat{d}_{\text {try }}$ are flagged from the trial configuration that might come into contact in the corrected time step. All these flagged points are considered when setting up $\mathcal{D}$ and determine the direction of contact forces and torques. The threshold $\hat{d}_{\text {try }}$ is chosen to set up a buffer, not to miss any colliding particles and sets at the same time the dimension of the optimisation problem, i.e. the number of colliding particle pairs $N_{c}$. The contact force potential and the force definition in (21) allows for contact force complementarity with respect to $\hat{d}_{\text {try }}$ at the trial time-step, meaning that we allow for a non-zero contact force for pair $i$ only if there is a contact distance at the trial time-step such that $d_{i}\left(\boldsymbol{Q}_{t+\Delta t}^{*}\right)<\hat{d}_{\text {try }}$.

Remark 3. We could also choose to construct the contact force potential at the previous, contact free time-step, from $\boldsymbol{Q}_{t}$. Note that this would require a larger $\hat{d}_{\text {try }}$, as particles are expected to move more in a full time-step than with only the correction from contact forces.

Remark 4. Even if we have written the forces and torques as the gradient of a barrier energy, it is not the gradient of a conservative potential as also $\lambda$ depends on $Q_{t+\Delta t}^{*}$, implicitly.

Next, we will compare using $d^{\mathrm{P}}\left(\boldsymbol{Q}_{t+\Delta t}^{*}\right)$ or $d^{s}\left(\boldsymbol{Q}_{t+\Delta t}^{*}\right)$ in the contact force potential and hence when assembling $\mathcal{D}$. The differences are also outlined in Fig. 3. In the case of a contact force direction determined for the particle pair in contact $i$, corresponding to using $d_{i}^{\mathrm{p}}$ as the definition of contact distance, we identify that $B_{i}\left(\boldsymbol{Q}_{t+\Delta t}^{*}, \hat{d}_{\mathrm{try}}\right)=b\left(d_{i}^{p}\left(\boldsymbol{Q}_{t+\Delta t}^{*}\right), \hat{d}_{\mathrm{try}}\right)$. Let $\boldsymbol{n}_{i}$ be the outward unit normal from the particle in the contact pair with the lowest index (pointing away from the contact) in the trial time-step. We then let

$$
\begin{align*}
& \hat{\boldsymbol{f}}_{i}^{c}=-\left.\frac{\partial b\left(x, \hat{d}_{\mathrm{try}}\right)}{\partial d_{i}}\right|_{d_{i}^{\mathrm{p}}\left(Q_{t+\Delta t}^{*}\right)} \boldsymbol{n}_{i} \\
& \hat{\boldsymbol{t}}_{i j}^{c}=-\left.\frac{\partial b\left(x, \hat{d}_{\mathrm{try}}\right)}{\partial d_{i}}\right|_{d_{i}^{\mathrm{p}}\left(\boldsymbol{Q}_{t+\Delta t}^{*}\right)}\left(\boldsymbol{s}_{i j} \times \boldsymbol{n}_{i}\right)  \tag{26}\\
& \hat{\boldsymbol{t}}_{i k}^{c}=\left.\frac{\partial b\left(x, \hat{d}_{\mathrm{try}}\right)}{\partial d_{i}}\right|_{d_{i}^{\mathrm{p}}\left(\boldsymbol{Q}_{t+\Delta t}^{*}\right)}\left(\boldsymbol{s}_{i k} \times \boldsymbol{n}_{i}\right),
\end{align*}
$$

where $s_{i j}$ and $s_{i k}$ are the vectors from the centre of each particle to the point of contact, in the global reference frame. Here, we remark that similarly to what is done by Yan et al. in [14], the contact stress and mechanical pressure can be computed from the contact forces.

We can also choose to build $\mathcal{D}$ for all pairs of surface points in contact, with $d^{s}$ the definition of separation used in the contact force potential. Let $l_{j}$ be the index of a surface point on particle $j, l_{k}$ an index of a surface point on particle $k$ and $r_{l}$ denote a surface point. Then, let $C_{i}$ denote the set of surface point pairs on different particles that are within a distance $\hat{d}_{\text {try }}$ from each other. In this
setting, the modified barrier energy $B_{i}\left(Q_{t+\Delta t}^{*}, \hat{d}_{\text {try }}\right)$ denotes the barrier energy for all surface points involved in the collision for the contact pair $i$, so that

$$
\begin{equation*}
B_{i}\left(\boldsymbol{Q}_{t+\Delta t}^{*}\right)=\sum_{l \in C_{i}} b\left(d_{l}^{\mathrm{s}}\left(\boldsymbol{Q}_{t+\Delta t}^{*}\right), \hat{d}_{\mathrm{try}}\right) \tag{27}
\end{equation*}
$$

Following (22), this means that the direction of the force will be given by the sum of the "normal directions" for each pair of grid nodes close to contact, $\left(\boldsymbol{r}_{l_{j}}-\boldsymbol{r}_{l_{k}}\right) /\left\|\boldsymbol{r}_{l_{j}}-\boldsymbol{r}_{l_{k}}\right\|$, weighted by the corresponding derivative of the barrier function. All in all, the components of $\boldsymbol{D}_{i}$ take the form

$$
\begin{align*}
& \hat{\boldsymbol{f}}_{i}^{c}=-\sum_{\left(l_{j}, l_{k}\right) \in C_{i}}\left(\left.\frac{\partial b\left(x, \hat{d}_{\mathrm{try}}\right)}{\partial x}\right|_{x=\left\|\boldsymbol{r}_{l_{j}}-\boldsymbol{r}_{l_{k}}\right\|}\right) \frac{\boldsymbol{r}_{l_{j}}-\boldsymbol{r}_{l_{k}}}{\left\|\boldsymbol{r}_{l_{j}}-\boldsymbol{r}_{l_{k}}\right\|},  \tag{28a}\\
& \left.\hat{\boldsymbol{t}}_{i j}^{c}=-\sum_{\left(l_{j}, l_{k}\right) \in C_{i}}\left(\left.\frac{\partial b\left(x, \hat{t}_{\text {try }}\right)}{\partial x}\right|_{x=\left\|\boldsymbol{r}_{l_{j}}-\boldsymbol{r}_{l_{k}}\right\|}\right)\left(\boldsymbol{r}_{l_{j}}-\boldsymbol{x}_{j}\right) \times \frac{\left(\boldsymbol{r}_{l_{j}}-\boldsymbol{r}_{l_{k}}\right)}{\left\|\boldsymbol{r}_{l_{j}}-\boldsymbol{r}_{l_{k}}\right\|}\right),  \tag{28b}\\
& \left.\hat{\boldsymbol{t}}_{i k}^{c}=\sum_{\left(l_{j}, l_{k}\right) \in C_{i}}\left(\left.\frac{\partial b\left(x, \hat{d}_{\mathrm{try}}\right)}{\partial x}\right|_{x=\left\|\boldsymbol{r}_{l_{j}}-\boldsymbol{r}_{l_{k}}\right\|}\right)\left(\boldsymbol{r}_{l_{k}}-\boldsymbol{x}_{k}\right) \times \frac{\left(\boldsymbol{r}_{l_{j}}-\boldsymbol{r}_{l_{k}}\right)}{\left\|\boldsymbol{r}_{l_{j}}-\boldsymbol{r}_{l_{k}}\right\|}\right) . \tag{28c}
\end{align*}
$$

This means that a certain force direction $\left(\boldsymbol{r}_{l_{j}}-\boldsymbol{r}_{l_{k}}\right) /\left\|\boldsymbol{r}_{l_{j}}-\boldsymbol{r}_{l_{k}}\right\|$ has a larger weight if the corresponding distance $\left\|\boldsymbol{r}_{l_{j}}-\boldsymbol{r}_{l_{k}}\right\|$ is smaller and the "overlap" relative to the minimum allowed distance $\hat{d}$ larger. Distances very close to $\hat{d}$ give only very small contributions to the sums representing force and torque directions in (28).

Note that alternatives could be considered for the barrier function, such as

$$
b_{0}(d, \hat{d})=\left\{\begin{align*}
-\ln (d / \hat{d}), & 0<d<\hat{d},  \tag{29}\\
0, & d \geq \hat{d},
\end{aligned} \quad \text { or } \quad b_{1}(d, \hat{d})=\left\{\begin{aligned}
(d-\hat{d}) \ln (d / \hat{d}), & 0<d<\hat{d} \\
0, & d \geq \hat{d}
\end{align*}\right.\right.
$$

with regularity $C_{0}$ and $C_{1}$ respectively and discussed also in [27]. We choose the barrier function $b$ in (11) due to its smooth transition at $d=\hat{d}$ ( $b$ has regularity $C_{2}$ ). This makes $b$ well-suited for the optimisation problem in (16), especially for computing derivatives of the barrier function, which has to be done in any iterative method to solve (16). For determining the contact force direction at the trial time-step via (28), we have no such regularity demand and choose the $C^{0}$ function $b^{0}$, for which the weight to the force and torque from the derivative of the barrier function is the reciprocal of the distance $\left\|\boldsymbol{r}_{l_{j}}-\boldsymbol{r}_{l_{k}}\right\|$.

Remark 5. Note that particles might not only violate the threshold $\hat{d}$ but also overlap, with $d_{i}<0$, before a contact force of the right magnitude is applied. Even if this is not the case at the trial time-step, it can happen in the iterative optimisation procedure to solve (16). We can, however, not evaluate the barrier function $b$ in (11) with a negative argument. One solution is to map $d_{i} \rightarrow d_{i}+\epsilon_{\text {reg }}$ and $\hat{d} \rightarrow \hat{d}+\epsilon_{\text {reg }}$, with $\epsilon_{\text {reg }}$ a parameter to be chosen. Note that if $\epsilon_{\text {reg }}$ is chosen large, the gradient of the barrier function is small also for small $d_{i}$ relative to $\hat{d}$, which might have an impact on the performance of solving the optimisation problem. Note also that in a Brownian setting, we may need to choose a large $\epsilon_{\text {reg }}$ if a large time-step size is used (depending on how repelling the conservative potential is that gives rise to external forces and torques). A different idea to avoid a large $\epsilon_{\text {reg }}$ is to map all $d_{i}<\epsilon_{\text {cap }}$ linearly such that the barrier function at $d_{i}=\epsilon_{\text {cap }}$ is $C^{1}$ and extended to a linear function, with a typical choice of $\epsilon_{\text {cap }}$ being a small fraction of $\hat{d}$, e.g. $\epsilon_{\text {cap }}=10^{-2} \hat{d}$. The three alternative barrier functions in (11) and (29) and their regularisations are visualised in Fig. 4, with the remapping of distances via $\epsilon_{\text {reg }}$ in Fig. 4a and the linear extension of $b$ in Fig. 4b. In numerical experiments in Section 3, we employ the strategy with $\epsilon_{\text {reg }}$, as we numerically have observed a reduced number of iterations for solving the optimisation problem with this choice, as gradients of $b$ are smaller and hence easier to handle.

### 2.2. Connection to a complementarity formulation

For comparison to complementarity techniques in the literature, as in [35,14,36-39], the accumulated non-overlap constraints in $\sum_{i} B_{i}(\lambda)=0$ can equivalently be expressed as a complementarity problem for the barrier energy of each contact pair and the corresponding contact force magnitude:

$$
\begin{equation*}
\mathbf{0} \leq-\boldsymbol{B}\left(\boldsymbol{Q}_{t+\Delta t}^{*}+\Delta t \boldsymbol{\Psi} \mathcal{M} \mathcal{D} \lambda, \hat{d}\right) \perp \lambda \geq \mathbf{0} . \tag{30}
\end{equation*}
$$

Note that the formulation in (30) holds despite the fact that the barrier function $b$ never takes negative values. The problem in (30) can be solved by first linearising about the trial time-step so that

$$
\begin{equation*}
\mathbf{0} \leq-\boldsymbol{B}\left(\boldsymbol{Q}_{t+\Delta t}^{*}, \hat{d}\right)+\left(\nabla_{Q_{t+\Delta t}^{*}}^{T} \boldsymbol{B}\left(\boldsymbol{Q}_{t+\Delta t}^{*}, \hat{d}\right) \Delta t \boldsymbol{\Psi} \boldsymbol{\mathcal { M }} \boldsymbol{\Psi}^{T} \nabla_{Q_{t+\Delta t}^{*}} \boldsymbol{B}\left(\boldsymbol{Q}_{t+\Delta t}^{*}, \hat{d}_{\mathrm{try}}\right)\right) \lambda \perp \lambda \geq \mathbf{0} \tag{31}
\end{equation*}
$$



Fig. 4. Barrier functions as defined in (11) and (29) are smooth approximations of the indicator function. Here, the barrier function is modified to be defined for negative inter-particle distances, as particles might overlap in the trial time-step and in iterations of the optimisation method before contact forces of the right magnitudes are found.
and techniques for LCPs may be used as in [47]. To guarantee a solution to the original nonlinear problem in (30), a sequence of linear problems of the form in (31) have to be solved, each with an update of the trial time-step, $\boldsymbol{Q}_{t+\Delta t}^{*}$, containing already computed position updates by an accumulated contact force.

The equation in (30) is reduced to the formulation by Yan et al. in $[14,36,38]$ if only one point is involved in the contact per particle per contact pair so that

$$
\begin{equation*}
\mathbf{0} \leq \boldsymbol{d}^{p}\left(\boldsymbol{Q}_{t+\Delta t}^{*}(\lambda)\right)-\hat{d} \perp \lambda \geq \mathbf{0}, \boldsymbol{d}^{p} \in \mathbb{R}^{N_{c}} . \tag{32}
\end{equation*}
$$

The only difference is then that we view contact forces as a correction to a trial time-step at $t+\Delta t$ and linearise about $Q_{t+\Delta t}^{*}$ instead of linearising about the previous coordinate vector $\boldsymbol{Q}_{t}$, which is done in $[14,36,38]$. At $\boldsymbol{Q}_{t}$, the configuration is contact free. That is however not guaranteed at $\mathcal{Q}_{t+\Delta t}^{*}$, which likely affect the speed of convergence of the optimisation problem, as we are likely to start with an infeasible initial guess. In this work, in contrast to the work by Yan et al., we find a solution to the nonlinear complementarity problem in (30) by computing a solution vector $\lambda$ satisfying $\sum_{i} B_{i}(\lambda)=0$, instead of considering a linearised problem. Note that it is not possible to strictly solve (30), as it is possible to have a configuration with $\boldsymbol{B}\left(\boldsymbol{Q}_{t+\Delta t}^{*}+\Delta t \boldsymbol{\Psi} \mathcal{M} \mathcal{D} \lambda, \hat{d}\right)=\mathbf{0}$ and $\lambda=\mathbf{0}$ simultaneously and hence we allow for some relaxation in the complementarity condition. We leave it to later work for a more thorough comparison between the nonlinear barrier method presented in this paper and versions of solving linear complementarity problems of the form in (31) and as discussed in [35,14,36-39].

Remark 6. With the choice $\hat{d}_{\text {try }}=\hat{d}$, methods for Quadratic Programming problems (QPs) [48,49] are eligible to solve (31), as the matrix in the LCP in (31) is symmetric positive semi-definite [14]. However, with such a small choice of the buffer region at the trial time-step (implied by choosing $\hat{d}_{\text {try }}$ to be as small as the natural choice of $\hat{d}$ ), there is a risk of missing potential contacts in the corrected time-step, leading to a long sequence of linear complementarity problems to be solved to eventually converge to a solution to the nonlinear complementarity problem. On the other hand, if $\hat{d}$ is instead increased to be as large as the natural choice of $\hat{d}_{\text {try }}$, to get equality between the two parameters, we would have a very large buffer region around each particle at the corrected time-step, with resulting contact forces heavily affecting the physics of the suspension.

If the dissipative energy induced by contact forces is chosen as objective function, the Lagrangian in our nonlinear optimisation problem (16) takes on a similar form as in the works by Yan et al., in which a QP on the form

$$
\begin{equation*}
\min _{\lambda \geq \mathbf{0}} \quad \lambda^{T} \mathcal{D}^{T} \mathcal{M} \mathcal{D} \lambda+\boldsymbol{G}^{T} \lambda, \tag{33}
\end{equation*}
$$

with $\boldsymbol{G}$ a specified vector is considered. For this QP, the first order optimality conditions constitute an LCP for the contact distances at the known time-step $Q_{t}$ and the force magnitudes $\lambda$. The difference in this paper is that we consider a problem with nonlinear constraints in $\lambda$ and the second term in the Lagrangian of (16) is therefore nonlinear, in contrast to in (33). Solving the problem becomes harder in our setting, but in contrast to the work by Yan et al., non-overlaps can be guaranteed.

### 2.3. Solution strategy and basic examples

The problem (16) is solved with fmincon, an inbuilt solver for constrained minimisation in Matlab 2021a, employing an interior-point method. In this and any other iterative method that could be considered for solving the problem, a full Stokes solve is required for each evaluation of the barrier energy, requiring as much work as one time-step without contact handling. A large
number of such iterations hence become intractably expensive, and it is important to keep the number of iterations as low as possible. This number is determined by a set of parameters that have to be carefully set. The time-step size, $\Delta t$, is set a priori depending on the typical magnitudes of the computed rigid body velocities in the problem, which depend on the type of background flow and the magnitude of applied forces and torques, so that the spatial updates of the particles per time-step are reasonable in magnitude. To demonstrate the robustness of the method, we will in this work vary the threshold $\hat{d}$, determining the minimum allowed distance between particles and the non-zero contribution to the barrier energy at the new time-step. In a general application, we however recommend picking $\hat{d}$ as a fraction of the particle radius, e.g. $\hat{d}=10^{-2} R$. As a rule of thumb for the remaining parameters, we pick:

- The stopping criterion for minimising contact force magnitudes,

$$
\begin{equation*}
\max _{i}\left|\lambda_{i}^{k+1}-\lambda_{i}^{k}\right|<10^{-2}\|\lambda\|_{\infty} \tag{34}
\end{equation*}
$$

- The threshold $\hat{d}_{\text {try }}$ determining both the $N_{c}$ particle pairs to potentially be assigned contact forces and the contact force and torque directions for these pairs assembled in the matrix $\mathcal{D}$. For rods and boomerangs with a given $\Delta t$, we set $\hat{d}_{\text {try }}$ adaptively in each time-step as

$$
\begin{equation*}
\hat{d}_{\mathrm{try}}=\hat{d}+\Delta t \max _{i=1, \ldots, N}\left(u_{i}+(L / 2) s_{i} \times \omega_{i}\right) \tag{35}
\end{equation*}
$$

where the expression in the parenthesis is the tip velocity of a particle, with $s_{i}$ the unit direction along the axis of particle $i$ and $L$ the particle length. For spheres, we pick

$$
\begin{equation*}
\hat{d}_{\mathrm{try}}=\hat{d}+\Delta t \max _{i=1, \ldots, N} \boldsymbol{u}_{i} \tag{36}
\end{equation*}
$$

- The regularisation parameter $\epsilon_{\text {reg }}$ for the barrier function as $\epsilon_{\text {reg }}=\hat{d}_{\text {try }}$.

Numerical examples are reported in Section 3 and the influence of the parameter choices will specifically be considered in Section 3.1.2. We however first consider two introductory examples where contact forces and torques are computed for two rods of the type illustrated in Fig. 3. The aim is to highlight the effect of different choices of $\hat{d}$ in a simplified setting. In example 1 , a single time-step is considered, and in example 2, we look at how a trajectory over time is affected by the contact handling.

## Set-up for two rods in example 1:

- Two rodlike particles of length $L_{\text {rod }}=2$ and radius $R_{\text {rod }}=L_{\text {rod }} / 4$ with particle orientations given by $q_{1}=[0,0,0,1]^{T}$ and $\boldsymbol{q}_{2}=[\sqrt{2} / 2,1 / 4,1 / 4,0]^{T}$, such that the first particle points in the positive $z$-direction and the second has spherical coordinates $\theta=\pi / 4, \phi=\pi / 4$.
- Fixed position of the first particle at the origin, $\boldsymbol{x}_{1}=\mathbf{0}$. The centre coordinate of the second particle is varied, with $\boldsymbol{x}_{2}=\beta \boldsymbol{d}$, where $\boldsymbol{d}$ is a direction vector given by $\boldsymbol{d}=[\cos (\alpha), 0, \sin (\alpha)]^{T}$. The angle $\alpha$ is varied in the range $[0, \pi / 2]$ and the parameter $\beta$ is adjusted such that the particle-particle distance is $\hat{d}$.
- Each particle is assigned a unit translational velocity in the direction of the other particle such that they are pushed together and the contact algorithm is triggered. See the particle configuration and velocities in Fig. 5a.
- Contact forces and torques are determined only for a single time-step, with $\Delta t=0.01$.

In Figs. $5 \mathrm{~b}-5 \mathrm{~d}$, three things are displayed: the magnitude $\lambda$ of the contact force and torque, the $x$-component of the contact force on particle 1 and the $y$-component of the contact torque on the same particle. All quantities are smooth in $\alpha$ and the largest contact force/torque magnitudes are obtained for the smallest choices of $\hat{d}$. This is natural, as for a smaller inter-particle distance, a larger force and torque are required to balance the assigned translational velocities that make the particles come too close to contact. Furthermore, the curves in each of the Figs. 5b-5d coincide for the smallest $\hat{d}$, meaning that slight changes to $\hat{d}$ do not impact the contact handling.

## Set-up for two rods in example 2:

- Particles of the same type as in example 1 are affected by gravity and studied over time.
- One rod aligned with a horizontal axis is placed below another rod aligned with the vertical axis, but with a small horizontal off-set $\delta$ from the centre-line of the top rod, as illustrated in Fig. 6a.
- With time, the top particle will approach the bottom particle, as more drag is experienced by the bottom particle. A periodic kiss-and-tumble motion then follows, provided that the initial distance $d_{\text {start }}$ and the off-set $\delta$ are chosen not very large. Similar phenomena are discussed e.g. by Guazzelli \& Morris in [50].

The behaviour for the off-set $\delta=0.3$ and a few initial centre-centre distances $d_{\text {start }}$ is illustrated in Fig. 6b, with the centre-centre distance and particle-particle alignment as functions of time. In our example, very coarse multiblob grids are used to represent the particles, with only 128 grid points per particle, and the contact algorithm is needed to prevent particle overlaps. The qualitative


Fig. 5. A demonstration of the robustness for the computed contact forces and torques in a single time-step, as the relative position of two closely interacting rods is varied. With the first particle centred at the origin, the angle $\alpha$ (indicated on the $x$-axes) together with the initial inter-particle distance $\hat{d}$ determines the centre coordinate of the second particle. For small $\hat{d}$, the computed quantities collapse onto the same curve.


Fig. 6. Two rods fall in a gravitational field from the initial configuration in (a). The top particle experiences less drag than the bottom one and therefore initially falls faster and approaches the second particle. The particles then follow a periodic kiss-and-tumble motion. The on-set of the periodic motion depends on the initial distance between the particles, $d_{\text {start }}$ and the off-set between their centre-lines, $\delta$. The contact algorithm aids in preventing overlaps between the particles. Only small differences can be noted for different choices of the parameter $\hat{d}$, given the same combination of ( $\delta, d_{\text {start }}$ ).
behaviour of the tumbling of the particles is not affected by small changes in the parameter $\hat{d}$, despite the fact that the particles are studied over a long time-span, with the time-step size set to $\Delta t=0.01$.

## 3. Numerical results

The two small examples illustrated in Section 2.3 demonstrate robustness in the computed contact forces and torques. In this section, we perform more elaborate numerical experiments with spheres, rods and boomerangs, with larger particle numbers and therefore more complex contact scenarios. Despite the result demonstrated in Fig. 5, it is generally difficult to study the performance of the contact algorithm dynamically, as a small error in the time-discretisation or small differences in the contact forces might result in completely different particle trajectories after a sufficiently long time. No standard benchmarks are available for systems of multiple particles. For this purpose, we choose to focus on two different types of tests that:

1. Investigate the ability of the contact resolution strategy to find contact force/torque magnitudes that maintain the allowed distance $\hat{d}$. This is a test of the complementarity condition, stating that a non-zero force and torque should be applied if and only if the contact is "active". The test is performed for geometries where the particles are pushed to come into contact by an external force, starting from a contact free configuration.
2. Investigate the ability to preserve symmetries in a background flow that naturally enhance particle interactions to quantify the impact of different choices of $\hat{d}$ and $\Delta t$.


(c) Inter-particle distances for the $N_{c}$ contact pairs before and after applying contact forces, with statistics collected from the 200 configurations with packing density $24 \%$. For the size of $\hat{d}_{\text {try }}$ relative to $\hat{d}$, see panel (d).
(d) Minimum distances upon applying contact forces, $d_{\min }(\lambda)$, respect the set distance $\hat{d}$ both for the denser and coarser configurations. For each configuration, the corresponding $\hat{d}_{\mathrm{try}}$ is also displayed, determining the $N_{c}$ particle pairs flagged to be part of the contact force optimisation.
(e) Same as in (c) but with distances scaled relative to $\hat{d}$ instead of $\hat{d}_{\text {try }}$. The distribution of larger particle distances remains after applying the contact forces. Only the smallest distances are affected.

Fig. 7. Contact forces are computed for all particle pairs that violate $\hat{d}_{\text {try }}$ among 500 randomly positioned spheres in a cube, where the packing density is $24 \%$ and $12 \%$ respectively. The spheres are affected by external forces with directions randomly drawn from the unit sphere and $\left\|f_{i}\right\|=100$, and a single time-step is taken with $\Delta t=0.01$. Particles are in contact if they are closer to each other than $\hat{d}$, here picked for each random configuration as the minimum separation distance at the previous, contact-free time-step. Minimum separation distances $\hat{d}$ are respected with the contact forces, and moreover, the forces do not drastically change the inter-particle distances for the closest particles, meaning that the forces are not unnecessarily large.

In some examples, we will exaggerate $\Delta t$ and/or external force magnitudes and magnitudes of the background flow to trigger the contact avoiding algorithm, with the purpose of demonstrating its robustness.

### 3.1. Spheres

### 3.1.1. A random suspension of spheres

We perform an experiment with configurations of 500 non-overlapping unit spheres randomly distributed in a cube of length $L$, suspended in a fluid of unit viscosity. The geometry is exemplified in Figs. 7 a and 7 b for the packing densities $24 \%$ and $12 \%$. In each of 200 configurations for each density, $\hat{d}$ is set to be the minimum separation distance. The spheres are then assigned external forces uniformly sampled from a sphere and scaled so that $\left\|f_{i}\right\|=100$. A trial time-step with $\Delta t=0.01$ is taken and for spheres that have come too close, contact forces are computed with $d_{i}$ determined at the particle level ( $d_{i}^{\mathrm{p}}$ is used). Note that contact torques are not computed for spheres, as normal collision forces induce no torques for spheres. Among all the corrected particle pairs, the smallest distance at the next time-level is reported vs $\hat{d}$ in Fig. 7d. In Figs. 7c and 7e, normalised distances are displayed for the particles flagged to potentially be in contact during the time-step, before and after applying the contact forces and scaled relative to $\hat{d}_{\text {try }}$ and $\hat{d}$ respectively. In conclusion, the forces do not artificially push the network of particles apart, but approximately keep the original inter-particle distances, except for the particle-pairs that have come too close, for which the "overlap" is avoided by applying a contact force.

### 3.1.2. Hyperparameter robustness test

We perform the same test as for the dense suspension in Section 3.1.1 but vary two parameters: $\hat{d}_{\text {try }}$, that determines the buffer region around each particle in the trial time-step and the number of particle pairs flagged to potentially come in contact during

(a) Bars display statistics for the contact force magnitudes with different buffer regions as given by $\hat{d}_{\text {try }}$ and stopping criteria $\max _{i}\left|\lambda_{i}^{k+1}-\lambda_{i}^{k}\right|<\operatorname{TOL}\|\lambda\|_{\infty}$. Whiskers show the minimum and maximum contact force magnitudes relative to $\left\|f_{i}\right\|$, the upper box edge the top $2.5 \%$ of force magnitudes, the lower box edge the bottom $1 \%$ and the line inside the box, the median.

(b) Histogram of the contact force magnitudes for the contact pairs with $\hat{d}_{\text {try }}>\hat{d}_{\text {try }}^{*}$ not flagged for collision with $\hat{d}_{\mathrm{try}}=\hat{d}_{\mathrm{try}}^{*}$. All are small in magnitude.

Fig. 8. Almost all contact forces are small in magnitude relative to the external force, which implies that only the pair or pairs that actually violate the set minimum allowed distance $\hat{d}$ are assigned a significant contact force, even if more particle pairs are flagged to be part of the contact optimisation. For each hyperparameter combination ( $\hat{d}_{\text {try }}$, TOL), statistics are collected from all flagged contacts in 200 random configurations of spheres with packing density $24 \%$. The same behaviour is noted for all hyperparameter combinations.

Table 1
Mean squared deviation (MSD) and max difference of contact forces with different hyperparameter combinations ( $\hat{d}_{\text {try }}$, TOL). The differences are computed relative to the reference $\hat{d}_{\text {try }}=\hat{d}_{\text {try }}^{*}$ and TOL $=5 \cdot 10^{-6}$ over all three components of the contact force for all spheres in 200 random configurations.

| (a) MSD of contact force |  |  |  | (b) Max (relative) difference of contact force |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| $\hat{d}_{\text {try }} / \mathrm{TOL}$ | $1 \cdot 10^{-2}$ | $5 \cdot 10^{-3}$ | $5 \cdot 10^{-6}$ | $\hat{d}_{\text {try }} / \mathrm{TOL}$ | $1 \cdot 10^{-2}$ | $5 \cdot 10^{-3}$ | $5 \cdot 10^{-6}$ |
| $\hat{d}_{\text {try }}^{*}$ | $2.45 \cdot 10^{-6}$ | $6.63 \cdot 10^{-7}$ | Reference | $\hat{d}_{\text {try }}^{*}$ | 0.14 | 0.14 | Reference |
| $1.1 \hat{d}_{\text {try }}^{*}$ | $3.76 \cdot 10^{-3}$ | $9.60 \cdot 10^{-6}$ | $8.84 \cdot 10^{-6}$ | $1.1 \hat{d}_{\text {try }}^{*}$ | 0.15 | 0.16 | 0.16 |
| $1.2 \hat{d}_{\text {try }}^{*}$ | $3.86 \cdot 10^{-3}$ | $1.03 \cdot 10^{-5}$ | $9.72 \cdot 10^{-6}$ | $1.2 \hat{d}_{\text {try }}^{*}$ | 0.16 | 0.17 | 0.17 |

the time-step, and the stopping criterion in the contact force optimisation, TOL, where the iterative optimisation method is stopped if $\max _{i}\left|\lambda_{i}^{k+1}-\lambda_{i}^{k}\right|<$ TOL $\|\lambda\|_{\infty}$. The experiment is repeated with $\hat{d}_{\text {try }} \in\left\{\hat{d}_{\mathrm{try}}^{*}, 1.1 \hat{d}_{\mathrm{try}}^{*}, 1.2 \hat{d}_{\mathrm{try}}^{*}\right\}$, with $\hat{d}_{\mathrm{try}}^{*}$ defined in (36) and TOL $\in$ $\left\{10^{-2}, 5 \cdot 10^{-3}, 5 \cdot 10^{-6}\right\}$. As for the stopping criterion, one would expect that with a more restrictive TOL, the number of nonnegligible contact forces is reduced. In practice, the dependence on TOL for the magnitude of the contact force is however very small, as visualised in Fig. 8a. As the number of iterations required to solve (16) is increased with a stricter tolerance (iteration count not shown), it is hence preferable to choose TOL less strict. By viewing the statistics in Fig. 8a, one can also conclude that only the elements of $\lambda$ corresponding to particle pairs that would overlap without a contact correction get assigned a larger contact force and all the other magnitudes are small, for all tested combinations of hyperparameters. In Fig. 8b, the histogram shows the contact force magnitudes only for the pairs flagged with $\hat{d}_{\text {try }}>\hat{d}_{\text {try }}^{*}$ that are not flagged with $\hat{d}_{\text {try }}=\hat{d}_{\text {try }}^{*}$, relative to the magnitude of the external forces pushing particles together. Even if more particle pairs are flagged within the time-step using a larger $\hat{d}_{\text {try }}$, we show that all the extra contact forces computed with a larger buffer region are small in magnitude. In Table 1, the mean squared deviation and max relative difference in the three components of the contact force for any particle in 200 configurations are presented versus the setting with $\hat{d}_{\text {try }}=\hat{d}_{\text {try }}^{*}$ and TOL $=5 \cdot 10^{-6}$.

We end this section with a discussion on the choice of objective function. If we redo the random sphere experiment with TOL $=$ $5 \cdot 10^{-6}$ and $\hat{d}_{\text {try }} \in\left\{\hat{d}_{\text {try }}^{*}, 1.1 \hat{\mathrm{t}}_{\text {try }}^{*}, 1.2 \hat{d}_{\text {try }}^{*}\right\}$, but change the objective function to the approximated dissipative energy $\lambda^{T} \mathcal{D}^{T} \mathcal{M} \mathcal{D} \lambda$, the "extra" flagged particle pairs for larger $\hat{d}_{\text {try }}$ will get assigned a larger contact force magnitude, relative to choosing the 1 -norm objective function. The force magnitudes for these particle pairs are reported in Fig. 9 and can be compared to Fig. 8b.

### 3.2. Rods

### 3.2.1. Geometric considerations

Rods constituting of a central cylindrical part and semi-spherical caps at both ends are considered. The surface of the rod is everywhere one radius from its centre line segment and inter-particle distances $d_{i}^{\mathrm{p}}$ can hence be determined by computing the shortest distance between two line segments and subtracting $2 R_{\text {rod }}$, see [45,46]. The surface-node-to-particle distances $d_{i k}^{s}$ are determined by computing the shortest distance between each source grid node belonging to a grid on particle 1 and the line segment


Fig. 9. As in Fig. 8b, but using the objective function $\lambda^{T} \mathcal{D} \mathcal{M} \mathcal{D} \lambda$ and the strict stopping criterion TOL $=5 \cdot 10^{-6}$. A lot of the extra contact forces resulting from $\hat{d}_{\text {try }}>\hat{d}_{\text {try }}^{*}$ are not small in magnitude. We therefore conclude that this choice of objective function is not as robust as $\sum_{i} \lambda_{i}$. Results are in both cases normalised by the size of the external force, $\left\|\boldsymbol{f}_{i}\right\|$.
of particle 2. A new target surface node is introduced on particle 2 where the vector of shortest distance cuts the surface of the particle. This source/target pair is added to the list of surface points that will be used to form the elements in the matrix $\mathcal{D}$ in (28). The procedure is repeated for all points sufficiently close to the other particle on the resolved surfaces of both particles. Finally, the closest points of contact as defined by $d_{i}^{\mathrm{p}}$ are added to the list not to underestimate the closest distance between the particles, as for all grid node pairs $k, d_{i k}^{\mathrm{s}} \leq d_{i}^{\mathrm{p}}$.

### 3.2.2. A chain of rods

Chains of rods of length $L_{\mathrm{rod}}=0.5$ and radius $R_{\mathrm{rod}}=L_{\mathrm{rod}} / 20$ are considered, where for each chain, a unit direction vector $\boldsymbol{t} \in \mathbb{R}^{3}$ and rotation $(\theta, \phi)$ are drawn at random from the first orthant. The chain is constructed with the first rod placed at the origin with orientation coinciding with the $z$-axis. The consecutive rods are obtained from the previous by rotating the particle by $(\theta, \phi)$ and then translating the centre coordinate by $\beta(\hat{d}) t$ in the coordinate frame of the previous particle, with the constant $\beta(\hat{d})$ determining the magnitude of the translation such that the smallest distance between a pair of particles is $\hat{d}$. Example geometries are visualised in Fig. 10a.

Every rod in each chain is assigned a force in the direction of the next particle so that a collision is caused in the next time-step ( $\hat{d}$ is violated) with no contact handling. The given set of external forces on the chain configuration is then corrected with contact forces and torques. Statistics for the resulting particle-particle distances $d(\lambda)$, upon applying contact forcing, are reported versus $\hat{d}$ for a large range of $\hat{d}$ in Fig. 10b. Note that (i) the largest $\hat{d}$ values do not represent realistic choices for a dynamic simulation, but rather demonstrate the robustness of the method, and (ii) that the choices of $\hat{d}_{\text {try }}$ depend on $\Delta t$ and therefore do not follow the curve for $\hat{d}$. By comparing $d(\lambda)$ to the reference distance $\hat{d}$, it can be concluded that the minimum allowed distance $\hat{d}$ is respected for almost all particles in all chains. Statistics for $(d(\lambda)-\hat{d}) / \hat{d}$ are illustrated for all chains before and after contact in Fig. 10c. In Fig. 10d, histograms display probable contact force/torque magnitudes on any rod in any chain for fixed $\hat{d}$ relative to the magnitude of the external force triggering the contact avoiding algorithm. Contact forces are comparable in magnitude to the external forces.

For reproducibility, the time-step size is set to $\Delta t=0.05$ and the force magnitude is $\left\|f_{i}\right\|=2 \hat{d}^{1 / 3}$ (this particular choice was made to promote collisions for the entire range of $\hat{d}$ considered in the test).

### 3.2.3. Rods in a biaxial compression flow

Consider a 2D grid of 30 rods of length $L_{\text {rod }}=2$ and radius $R_{\text {rod }}=L_{\text {rod }} / 4$ arranged as in Figs. 11a-11b and affected by a background biaxial compression flow $\boldsymbol{\mathcal { U }}_{\mathrm{bg}}$ as indicated in the figure, given by

$$
\boldsymbol{U}_{\mathrm{bg}}=\left[\left(\boldsymbol{E} \boldsymbol{x}_{1}\right)^{T}, \mathbf{0},\left(\boldsymbol{E} \boldsymbol{x}_{2}\right)^{T}, \mathbf{0}, \ldots,\left(\boldsymbol{E} \boldsymbol{x}_{n}\right)^{T}, \mathbf{0}\right]^{T} \text {, with } \boldsymbol{E}=\left[\begin{array}{ccc}
\dot{\gamma} & 0 & 0  \tag{37}\\
0 & -\dot{\gamma} / 2 & 0 \\
0 & 0 & -\dot{\gamma} / 2
\end{array}\right] \text {, and } \dot{\gamma}=1 \text {. }
$$

We choose different minimum allowed distances $\hat{d}$ and a hierarchy of time-step sizes $\Delta t$ and discretise the dynamics of the system with forward Euler with one hyperparameter setting $(\hat{d}, \Delta t)$ at a time. Whenever the smallest allowed distance $\hat{d}$ is violated at the end of the time-step, contact forces and torques are computed for particles considered sufficiently close to contact at the trial time-step, as given by $\hat{d}_{\text {try }}$. Due to the contractile nature of the background flow, the contact avoiding algorithm will be triggered in almost every time-step, see Table 2. In Fig. 11c, the change in coordinate position per time-step is displayed for all particles until the particles start to diverge along the $x$-axis (where the background flow is diverging) and the simulation is stopped. The contact loading causes no drastic jumps in the particle trajectories, and a vast majority of the coordinate updates are very close to zero. Table 2 indicates that contacts are robustly handled for a long sequence of time-steps. The maximum deviation along the $x$-axis for the particles is displayed as function of time in Fig. 11e. A smaller time-step allows for a slightly smaller deviation than a larger time-step. Note however that for all hyperparameters, particles stay in the $y z$-plane.
 tween each consecutive pair of rods so that the inter-particle distance $\hat{d}$ is equal between every pair. Red arrows indicate external forces towards the centre of the next particle that deliberately cause collisions during the next time-step. Particle colours indicate the depth in the suspension.

(c) Bars show the relative difference in computed distances $d(\boldsymbol{\lambda})$ to $\hat{d}$ at the trial time-step without contact forces and upon applying contact forces, for each choice of $\hat{d}$. Whiskers display the minimum and maximum relative distance difference, box edges the 10th and 90th percentile of the relative difference, and the box centre lines display the median.

Fig. 10. Chains of 40 rods are forced to come into contact by externally applied forces towards the next particle in the chain. By finding optimal magnitudes of contact forces and torques, the minimum separation distance $\hat{d}$ can be kept. For each $\hat{d}$, the experiment is repeated with 100 randomly generated chains.

Table 2
Rods in a biaxial compression flow: Number of time-steps where the contact algorithm is triggered for the different time-step sizes $\Delta t$, together with the total number of steps taken. Here, $T_{S}$ is the time up until which the order parameter $S$ is displayed in Fig. 11d and $T_{\text {end }}$ the time when particles start to diverge along the $x$-axis (where the background flow is diverging).

|  | \# contact time-steps | \# contact time-steps | total \# steps |
| :--- | :--- | :--- | :--- |
| Time-step $\Delta t$ as | $\hat{d}=0.01 R_{\text {rod }}$ | $\hat{d}=0.1 R_{\text {rod }}$ |  |
| fraction of $\Delta t^{*}=0.02$ | $T_{\text {end }}=9\left(T_{S}=4.36\right)$ | $T_{\text {end }}\left(T_{S}\right)$ | $T_{\text {end }}\left(T_{S}\right)$ |
| $\Delta t^{*}$ | $410(178)$ | $412(180)$ | $450(218)$ |
| $\Delta t^{*} / 2$ | $819(357)$ | $824(362)$ | $900(436)$ |
| $\Delta t^{*} / 4$ | $1639(715)$ | $1647(723)$ | $1800(872)$ |

If all symmetries of the problem were kept, the particles would not only stay in the same plane, but also keep the alignment with the $z$-axis. The alignment can be quantified with the Onsager order parameter $S$ defined by [51]


Fig. 11. A grid of 30 rods affected by a biaxial compression flow. Contact forces and torques are robustly computed whenever the minimum allowed distance $\hat{d}$ is violated. The symmetry of the configuration is broken, but this happens after a large number of time-steps. Despite the challenging setting with the background flow pushing particles together, the particles stay in the original 2D-grid for a relatively long time.

$$
\begin{equation*}
S=\frac{1}{N} \sum_{i}^{N}\left\{3 / 2\left(e_{z} \cdot s_{i}\right)^{2}-1 / 2\right\} \tag{38}
\end{equation*}
$$

with $s_{i}$ the unit vector representing the symmetry axis for particle $i$. If $S=1$, the particles are perfectly aligned with the $z$-axis and if $S=-1 / 2$, particles are all perpendicular to the $z$-axis. The order parameter is visualised for two choices of the parameter $\hat{d}$ in Fig. 11d. We expect the order parameter to be $S(t)=1$ throughout the simulation, but numerically this only holds up until some point in time when the alignment is broken for some particles and $S(t)$ decreases. With varying choices of $\hat{d}$ and $\Delta t$, the same qualitative behaviour can be noted for the order parameter; The particles are initially ordered symmetrically and the background flow is symmetric, then, particles first start to come in contact in the vertical direction and later both vertically and horizontally. Due to contact forces and torques not being perfectly symmetric, the alignment is broken. This is a consequence of the general fact while-timestepping that after a large number of steps, errors will accumulate. The divergence of the order-parameter however happens after a large number of time-steps for all $(\hat{d}, \Delta t)$, from which we can conclude that the contact handling is robust.

### 3.3. Boomerangs

### 3.3.1. Geometric considerations

We consider boomerang particles of the type displayed in Fig. 1b. Such a boomerang is constructed by dividing a rod of aspect ratio $L_{\mathrm{rod}} / R_{\text {rod }}=8$ and radius $R_{\mathrm{rod}}=0.5$ as described in Section 3.2.1 into two equal pieces, each with a cylindrical part and a semi-spherical cap at one end. One of these half-rods is placed vertically and one placed horizontally so that the corners of the cut area of the two pieces touch. A circle is rotated around this point so that the two rod pieces are joined into one particle. The centre coordinate $\boldsymbol{x}$ is defined to be the mass centre of the particle.

(a) An example configuration (colours indicating depth in the suspension).

(b) The minimum distances upon applying contact forces and torques all respect the minimum allowed distance $\hat{d}$. For each configuration, the corresponding $\hat{d}_{\mathrm{try}}$ is also displayed, determining the $N_{c}$ particle pairs flagged to be part of the contact optimisation.
(c) Inter-particle distances for the $N_{c}$ contact pairs before and after applying contact forces and torques, with statistics collected from all 200 configurations. Note that $\hat{d}_{\mathrm{try}}$ is large relative to $\hat{d}$, see (b).

Fig. 12. Contact avoidance for 200 random configurations of 40 boomerangs.

For the barrier energy and the direction of contact forces and torques in $\mathcal{D}$, the contact distance $d^{s}$, based on surface-node-toparticle distances, is used, as the particle geometry is non-convex. For this purpose, the centre curve of each particle is considered, given by two line segments and a quarter of a circle joining the two segments. For each point on the discretised surface of particle 1 , we may easily compute the shortest distance to the centre curve of particle 2 , subtracting one particle radius to give the shortest distance between the surfaces. The corresponding source/target pair of surface points is added to the list of surface nodes used to construct the matrix $\mathcal{D}$ in (28). The procedure is repeated with reversed numbering of the two particles in the pair. Note that for small $\hat{d}$, a fine representation of the surface has to be used.

### 3.3.2. A random suspension of boomerangs

Systems of 40 randomly positioned boomerangs in a cube of length $L=12$ are considered, as exemplified in Fig. 12a. The boomerangs are generated one at the time with a uniformly sampled centre coordinate $x$ and quaternion $q$ such that the distance to any other particle is larger than $\delta=10^{-2}$. We perform the same type of test as for the spheres in Section 3.1: For each generated configuration of 40 boomerangs, the minimum separation distance is computed and $\hat{d}$ is chosen to be this distance. The boomerangs are then assigned forces and torques randomly sampled from a sphere with $\left\|\mathcal{F}_{\text {ext }}\right\|=100$ so that in a trial time-step with $\Delta t=0.01$, some particles violate the constraints for the surface-node-to-particle distances. For these particle pairs, correcting contact forces and torques are computed. Among all the particle pairs flagged to potentially be assigned a contact correction, the smallest distance obtained with a contact correction at time $t+\Delta t$ is reported vs $\hat{d}$ in Fig. 12b. All distances for the flagged pairs are reported before and after applying contact forces and torques in Fig. 12c. It can be concluded that contact forces and torques do not push particles apart unnecessarily far. Note also that the parameter $\hat{d}_{\text {try }}$ is chosen large relative to $\hat{d}$ not to miss any collisions. Despite the nonconvex particle shape and the risk of particles getting stuck in locked configurations, $\hat{d}$ can be maintained for all the boomerangs in all of the 200 configurations.

## 4. A discussion on the computational cost

Despite the increased cost in any time-step with contacts, a lot can be gained in terms of computational cost by applying the contact avoiding algorithm presented in this paper, as a much larger time-step can be considered for dynamical simulations than what would be allowed in a simulation without contact handling - time-step sizes used for the numerical experiments in this work are excessively large, but the contact algorithm is still robust.

The main cost of finding optimal contact force and torque magnitudes is the need for determining the matrix-matrix product $\mathcal{M D}$ in the computation of the gradient of the barrier energy with respect to the vector of contact force magnitudes, which is needed in the interior point method used for solving the optimisation problem in (16). The matrix-matrix product $\mathcal{M D}$ has to be computed and stored once per time-step where contact occurs, as $\mathcal{M}$ depends on $\boldsymbol{Q}_{t}$ and $\mathcal{D}$ depends on $\boldsymbol{Q}_{t+\Delta t}^{*}$ - there is no dependence on $\lambda$. For the multiblob method, the mobility matrix $\boldsymbol{\mathcal { M }}$ can easily be computed explicitly if the total number of particles is small. For larger particle systems, evaluating $\mathcal{M D}$ would amount to solving $N_{c}$ Stokes mobility problems, the cost of which is determined by $\hat{d}_{\text {try }}$ that will set $N_{c}$, which of course also depends on the particle density in the system. One Stokes mobility problem also has to be solved for every evaluation of the action of the contact loading, $\boldsymbol{\mathcal { D }} \lambda$, that enters in the barrier energy. The number of such solves depends on the number of iterations to solve the optimisation problem (16) with the interior-point method, which, in turn, depends on the particle type, the time-step size and a few hyperparameters as discussed in Section 2.3. Choices that need to be made are e.g. how to regularise the barrier function, i.e. how to handle negative distances between boundaries, and what stopping criterion to pick when solving for minimum contact magnitudes in (16). From numerical experiments in this paper, it is expected that the number of mobility solves $\mathcal{M D} \mathcal{D}$ at least is in the range $10-20$. The number of iterations is larger if the stopping criterion is very
strict and smaller if the stopping criterion is less strict, which could be allowed by choosing a slightly larger buffer region $\hat{d}$ than required for sufficient accuracy in $\mathcal{M}$, as dictated by the Stokes solver. We may compare the number of iterations to what is reported in the literature for the LCP technique in the work of Yan and collaborators, but at the same time emphasise that the iteration count will vary depending on the setting: Yan et al. report $\approx 20$ extra Stokes solves per time-step for spheres in [36] and 5-10 iterations for dilute suspensions of spherocylinders in [14], but $O(1000)$ close to the random-close-packing limit.

The first order interior-point method combined with the constrained minimisation problem considered in this work also has benefits over second order methods, e.g. based on Newton's method as applied to the non-constrained IPC-formulation [27]. In contrast to an IPC-formulation, we are not dependent on second order derivatives of the distances with respect to the particle coordinates and also avoid solving large linear systems - other than the Stokes mobility problem - while iteratively updating the solution vector in the optimisation problem. The matrix in the linear system in an IPC-formulation of the problem would be a function of $\boldsymbol{\mathcal { M }}$, but also depend on the gradient of external forces and torques and on the Hessian of the distances with respect to both particle centre coordinates and quaternions. The structure of this matrix is not trivially investigated for the general case a priori, and the matrix has to be projected onto the cone of non-negative matrices for Newton's method to converge. Moreover, fast matrix-vector techniques do not apply as for the Stokes mobility problem (applying $\mathcal{M \mathcal { F }}$, for some vector $\mathcal{F}$ ) [1]. Because of these difficulties, the current formulation coupled to a first order method, such as the interior-point method, is especially beneficial in very large particle systems, where we neither want to compute $\boldsymbol{\mathcal { M }}$ explicitly nor want to repeatedly solve a large linear system that is a non-trivial function of $\mathcal{M}$. ${ }^{4}$

## 5. Conclusions

We have presented an optimisation procedure to guarantee non-overlapping configurations of 3D particles in an unbounded Stokesian fluid. The method is based on a barrier formulation where non-overlapping constraints are rewritten as a barrier energy, constrained to be zero for non-overlapping configurations. Numerical examples are provided for spheres, rod-like particles with semi-spherical caps and boomerangs. The numerical examples show the performance of the contact algorithm for dense suspensions of particles deliberately pushed together by external forces or a background flow. In all examples, the highlighted property of the proposed method is its ability to keep a set minimum separation distance $\hat{d}$. Discrete complementarity is obtained at the trial timestep between the force/torque magnitude vector $\lambda$ and the modified barrier energy (with parameter $\hat{d}_{\text {try }}$ ). We have numerically shown that the method is robust for different choices of $\hat{d}$, particle shapes and flow scenarios. The magnitudes of external forces, background flows and time-step sizes are in this work deliberately chosen to force particles to come too close to each other and in a realistic simulation, contact will occur less frequently and $\hat{d}$ is supposed to be small. Strengths of the method are its simple formulation, its ability of handling non-convex particles, the no-collision guarantee of the formulation and the minimum impact on the system by contact forces, as these are balanced and their magnitudes are minimised. By solving for force/torque magnitudes explicitly, as we do here, a penalisation parameter can be omitted, which otherwise has to be carefully tuned. Moreover, the size of the optimisation problem is much more moderate in the number of primal variables (at least for moderately crowded systems) than if all particle coordinates are solved for implicitly, as in Incremental Potential Contact (IPC) [27]. Only one dual variable has to be solved for in a primal-dual approach.

For future work, we identify two directions:

1. Approximation of the action of the mobility matrix: There is a global coupling between all particles in the suspension in the mobility matrix, $\mathcal{M}$. Motivated by the fact that close interactions are dominating, one option is to build an approximation to the action of $\mathcal{M}$ for matrix-vector multiplies $\boldsymbol{\mathcal { M }}$ constructed only from all pairs of particles in contact, instead of applying the full matrix. An issue is however that a contact force/torque vector determined with such an approximation not necessarily has to lead to a non-overlapping configuration, as in the true system, all particles have an impact on the hydrodynamic interactions and all these interactions are not captured if $\mathcal{M}$ is modified. A workaround could be to introduce a small buffer for the minimum allowed distance so that $\hat{d} \rightarrow \hat{d}(1+\delta)$ or to use the approximated solution vector for contact force/torque magnitudes as an initial guess for iterations with the full mobility matrix. Along similar lines, separated clusters of particles could be considered to compute contact forcing locally.
2. A continuous contact force potential: In this work, a discrete version of the contact force potential is considered, computed from all points on the particle surfaces sufficiently close to each other. One could also consider a continuously defined potential for two particles in contact, where the barrier energy at the trial time-step is expressed in terms of integrals over the particle surfaces close to contact. The asymmetry of the contact forces and torques noted in the experiment with a biaxial compression flow in Section 3.2.3 has two explanations: To start with, the discrete surface grid on the particles results in different barrier energies for seemingly similar relative rotations and particle-particle distances for different colliding pairs. The reason is that a different number of grid nodes may be flagged for collision depending on rotations of the particles around their own axis. Furthermore, the shortest particle-particle distance $d^{p}$ is sensitive to slight differences in orientation when particles are close to parallel. One benefit of a continuous contact force potential could be to better preserve symmetries. We leave this approach for future work starting in 2D.
[^4]We end with a note on the optimisation algorithm: In this work, the inbuilt optimisation procedure fmincon in Matlab has been used to solve the optimisation problem for the contact force/torque magnitude vector $\lambda$, employing a highly tuned interior point method. The fact that a standard solver in Matlab can be used off-the-shelf is a strength of the method. We are satisfied with a possibly local optimum as long as a zero barrier energy is reached, implying sufficiently separated particles.

## CRediT authorship contribution statement

Anna Broms: Conceptualization, Data curation, Formal analysis, Investigation, Methodology, Software, Validation, Visualization, Writing - original draft. 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

The authors thank several people for discussions that have improved this work: Aleksandar Donev for sharing insights on dissipated energy, Anders Forsgren for discussions on the choice of optimisation algorithm, Georg Stadler for discussing alternative ways of setting up the non-overlap constraints, Mattias Sandberg for constructive input on the numerical examples, and Fredrik Lundell and Zhouyang Ge for their insights on contacts in physical systems. We also thank the anonymous reviewers for several useful comments that have improved the paper. Finally, we acknowledge the support from the Swedish Research Council: grant no. 2019-05206 and the research environment grant INTERFACE (biomaterials), no. 2016-06119. This work was done in connection with the Linné Flow Centre at KTH.

## References

[1] 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).
[2] A. Broms, M. Sandberg, A.-K. Tornberg, A locally corrected multiblob method with hydrodynamically matched grids for the Stokes mobility problem, J. Comput. Phys. 487 (2023) 112172.
[3] S. Delong, F. Balboa Usabiaga, A. Donev, Brownian dynamics of confined rigid bodies, J. Chem. Phys. 143 (2015).
[4] 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).
[5] 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.
[6] 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.
[7] A.M. Fiore, J.W. Swan, Fast Stokesian dynamics, J. Fluid Mech. 878 (2019) 544-597.
[8] J.R. Melrose, R.C. Ball, The pathological behaviour of sheared hard spheres with hydrodynamic interactions, Europhys. Lett. 32 (1995) 535-540.
[9] R.C. Ball, J.R. Melrose, Lubrication breakdown in hydrodynamic simulations of concentrated colloids, Adv. Colloid Interface Sci. 59 (1995) 19-30.
[10] A.A. Catherall, J.R. Melrose, R.C. Ball, Shear thickening and order-disorder effects in concentrated colloids at high shear rates, J. Rheol. 44 (2000) 1-25.
[11] J.F. Morris, Shear thickening of concentrated suspensions: recent developments and relation to other phenomena, Annu. Rev. Fluid Mech. 52 (2020) 121-144.
[12] L. af Klinteberg, A.K. Tornberg, Fast Ewald summation for Stokesian particle suspensions, Int. J. Numer. Methods Fluids 76 (2014) 669-698.
[13] J. Wang, E. Nazockdast, A. Barnett, An integral equation method for the simulation of doubly-periodic suspensions of rigid bodies in a shearing viscous flow, J. Comput. Phys. 424 (2021).
[14] 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).
[15] M.S.D. Wykes, J. Palacci, T. Adachi, L. Ristroph, X. Zhong, M.D. Ward, J. Zhang, M.J. Shelley, Dynamic self-assembly of microscale rotors and swimmers, Soft Matter 12 (2016) 4584-4589.
[16] B.J. Berne, P. Pechukas, Gaussian model potentials for molecular interactions, J. Chem. Phys. 56 (1972) 4195-4205.
[17] S. Varga, G. Jackson, Study of the pitch of fluids of electrostatically chiral anisotropic molecules: mean-field theory and simulation, Mol. Phys. 104 (2006) 3681-3691.
[18] 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).
[19] D. Baraff, Issues in computing contact forces for non-penetrating rigid bodies, Algorithmica 10 (1993) 292-352.
[20] Anitescu, Cremer, Potra, Formulating 3D contact dynamics problems, Mech. Struct. Mach. 24 (1996) 405-437.
[21] A. Tasora, M. Anitescu, A fast NCP solver for large rigid-body problems with contacts, friction, and joints, Multibody Syst. Dyn. (2008) 45-55.
[22] A. Tasora, D. Negrut, M. Anitescu, Large-scale parallel multi-body dynamics with frictional contact on the graphical processing unit, Proc. Inst. Mech. Eng., Proc., Part K, J. Multi-Body Dyn. 222 (2008) 315-326.
[23] M. Tur, F.J. Fuenmayor, P. Wriggers, A mortar-based frictional contact formulation for large deformations using Lagrange multipliers, Comput. Methods Appl. Mech. Eng. 198 (2009) 2860-2873.
[24] P. Wriggers, Computational Contact Mechanics, second ed., Springer, 2006.
[25] J.M. Snyder, Interactive tool for placing curved surfaces without interpenetration, in: Proceedings of the ACM SIGGRAPH Conference on Computer Graphics, 1995, pp. 209-218.
[26] D. Harmon, D. Panozzo, O. Sorkine, D. Zorin, Interference-aware geometric modeling, ACM Trans. Graph. 30 (2011) 1.
[27] M. Li, Z. Ferguson, T. Schneider, T. Langlois, D. Zorin, D. Panozzo, C. Jiang, D.M. Kaufman, Incremental potential contact: intersection- and inversion-free, large-deformation dynamics, ACM Trans. Graph. 39 (2020).
[28] Z. Ferguson, M. Li, T. Schneider, F. Gil-Ureta, T. Langlois, C. Jiang, D. Zorin, D.M. Kaufman, D. Panozzo, Intersection-free rigid body dynamics, ACM Trans. Graph. 40 (2021) 1-16.
[29] M. Li, D.M. Kaufman, C. Jiang, Codimensional incremental potential contact, ACM Trans. Graph. 40 (2021).
[30] S. Yamamoto, T. Matsuoka, A method for dynamic simulation of rigid and flexible fibers in a flow field, J. Chem. Phys. 98 (1993) 644-650.
[31] S. Yamamoto, T. Matsuoka, Dynamic simulation of fiber suspensions in shear flow, J. Chem. Phys. 102 (1995) 2254-2260.
[32] D. Das, D. Saintillan, Electrohydrodynamic interaction of spherical particles under Quincke rotation, Phys. Rev. E 87 (2013) 1-14.
[33] B. Delmotte, E. Climent, F. Plouraboué, A general formulation of Bead Models applied to flexible fibers and active filaments at low Reynolds number, J. Comput. Phys. 286 (2015) 14-37.
[34] 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.
[35] L. Lu, Parallel contact-aware algorithms for large-scale direct blood flow simulations, Ph.D. thesis, Courant Institute of Mathematical Sciences, New York University, 2019.
[36] W. Yan, E. Corona, D. Malhotra, S. Veerapaneni, M. Shelley, A scalable computational platform for particulate Stokes suspensions, J. Comput. Phys. 416 (2020) 109524.
[37] L. Bystricky, S. Shanbhag, B. Quaife, Stable and contact-free time stepping for dense rigid particle suspensions, Int. J. Numer. Methods Fluids 92 (2020) 94-113.
[38] W. Yan, S. Ansari, A. Lamson, M.A. Glaser, R. Blackwell, M.D. Betterton, M. Shelley, Toward the cellular-scale simulation of motor-driven cytoskeletal assemblies, eLife (2022) e74160.
[39] R. Kohl, E. Corona, V. Cheruvu, S. Veerapaneni, Fast and accurate solvers for simulating Janus particle suspensions in Stokes flow, Adv. Comput. Math. 49 (2023).
[40] L. Lu, M.J. Morse, A. Rahimian, G. Stadler, D. Zorin, Scalable simulation of realistic volume fraction red blood cell flows through vascular networks, in: International Conference for High Performance Computing, Networking, Storage and Analysis, 2019.
[41] L. Bystricky, Contact-free Simulations of Rigid Particle Suspensions Using Boundary Integral Equations, Ph.D. thesis, Florida State University, 2018.
[42] L. Lu, A. Rahimian, D. Zorin, Contact-aware simulations of particulate Stokesian suspensions, J. Comput. Phys. 347 (2017) 160-182.
[43] A. Gubbiotti, M. Chinappi, C.M. Casciola, Confinement effects on the dynamics of a rigid particle in a nanochannel, Phys. Rev. E 100 (2019) 1-17.
[44] S. Boyd, L. Vandenberghe, Convex optimization, https://doi.org/10.1142/9789814412520_0010, 2009.
[45] V.J. Lumelsky, On fast computation of distance between line segments, Inf. Process. Lett. 21 (1985) 55-61.
[46] Ondrej, Fast shortest distance between two line segments (in N dimensions), MATLAB Central File Exchange, https://www.mathworks.com/matlabcentral/ fileexchange/49502-fast-shortest-distance-between-two-line-segments-in-n-dimensions. (Accessed 4 May 2022).
[47] A. Fischer, A special Newton-type optimization method, Optimization 24 (1992) 269-284.
[48] Y.H. Dai, R. Fletcher, Projected Barzilai-Borwein methods for large-scale box-constrained quadratic programming, Numer. Math. 100 (2005) 21-47.
[49] R. Fletcher, On the Barzilai-Borwein method, in: L. Qi, K. Teo, X. Yang (Eds.), Optimization and Control with Applications, Springer US, Boston, MA, 2005, pp. 235-256.
[50] E. Guazzelli, J.F. Morris, A Physical Introduction to Suspension Dynamics, Cambridge Texts in Applied Mathematics, vol. 45, Cambridge University Press, Cambridge, 2012.
[51] H. Doi, K.Z. Takahashi, K. Tagashira, J. ichi Fukuda, T. Aoyagi, Machine learning-aided analysis for complex local structure of liquid crystal polymers, Sci. Rep. 9 (2019) 1-12.
[52] P.E. Gill, W. Murray, M.A. Saunders, SNOPT: an SQP algorithm for large-scale constrained optimization, SIAM Rev. 47 (2005) 99-131.


[^0]:    * Corresponding author.

    E-mail address: annabrom@kth.se (A. Broms).
    https://doi.org/10.1016/j.jcp.2023.112648
    Received 19 March 2023; Received in revised form 5 October 2023; Accepted 16 November 2023
    Available online 22 November 2023
    0021-9991/© 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]:    ${ }^{1}$ For multiblob particles of a more general geometry, one could construct a potential from an HGO-potential centred on each individual blob building up the particle.

[^2]:    ${ }^{2}$ The quaternion vector takes the form $\boldsymbol{q}_{i}:=\left(s_{i}, \boldsymbol{p}_{i}\right)$, where the scalar, $s$, specifies a rotation magnitude and $\boldsymbol{p} \in \mathbb{R}^{3}$ specifies a vector around which the rotation takes place.

[^3]:    ${ }^{3}$ One could also consider to set $\hat{d}$ in relation to any inhomogeneities on the surfaces of the physical particles modelled by our smooth particles.

[^4]:    4 An option for larger particle systems could potentially be to solve the problem by employing the sparse nonlinear optimiser provided in the SNOPT package [52]. Care however has to be taken for how second order derivatives are computed.

