FLORIDA STATE UNIVERSITY

COLLEGE OF ARTS AND SCIENCES

CONTACT-FREE SIMULATIONS OF RIGID PARTICLE SUSPENSIONS USING BOUNDARY INTEGRAL EQUATIONS

By

LUKAS BYSTRICKY

A Dissertation submitted to the
Department of Scientific Computing
in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy

2018

Lukas Bystricky defended this dissertation on July 16th, 2018.
The members of the supervisory committee were:

Bryan Quaife

Professor Co-Directing Dissertation

Sachin Shanbhag

Professor Co-Directing Dissertation

Nick Cogan

University Representative

Chen Huang

Committee Member

Nick Moore

Committee Member

The Graduate School has verified and approved the above-named committee members, and certifies that the dissertation has been approved in accordance with university requirements.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

ABSTRACT

In many composite materials, rigid fibers are distributed throughout the material to tune the mechanical, thermal, and electric properties of the composite. The orientation and distribution of the fibers play a critical role in the properties of the composite. Many composites are processed as a liquid molten suspension of fibers and then solidified, holding the fibers in place. Once the fiber orientations are known, theoretical models exist that can predict properties of the composite. Modeling the suspended fibers in the liquid state is important because their ultimate configuration depends strongly on the flow history during the molten processing.

Continuum models, such as the Folgar-Tucker model, predict the evolution of the fibers' orientation in a fluid. These models are limited in several ways. First, they require empirical constants and closure relations that must be determined *a priori*, either by experiments or detailed computer simulations. Second, they assume that all the fibers are slender bodies of uniform length. Lastly, these methods break down for concentrated suspensions. For these reasons, it is desirable in certain situations to model the movement of individual fibers explicitly. This dissertation builds upon recent advances in boundary integral equations to develop a robust, accurate, and stable method that simulates fibers of arbitrary shape in a planar flow.

In any method that explicitly models the individual fiber motion, care must be taken to ensure numerical errors do not cause the fibers to overlap. To maintain fiber separation, a repulsion force and torque are added when required. This repulsion force is free of tuning parameters and is determined by solving a sequence of linear complementarity problems to ensure that the configuration does not have any overlap between fibers. Numerical experiments demonstrate the stability of the method for concentrated suspensions.

CHAPTER 1

INTRODUCTION

Many natural and man-made materials can be classified as particle suspensions. The occurrence of suspensions spans many fields, including biofluids, foodstuff, and manufacturing processes. To give three examples, blood is a suspension of platelets and red and white blood cells, mayonnaise is a suspension of oil drops, and pulp used in paper manufacturing is a suspension of fibers [45]. In all of these examples, hydrodynamics plays a critical role. In particular, the particles are not only carried by the flow, but affect the flow. This is different from suspensions involving granular media, for example atmospheric pollutants, where the particle interactions are negligible, and do not affect the flow. This dissertation focuses on viscous suspensions of non-deformable rigid bodies such as fibers.

Rigid body suspensions in viscous fluids are important for manufacturing composite materials, where rigid fibers are distributed throughout the material. The orientation and distribution of these fibers can be used to tune the mechanical, thermal, or electrical properties of the composite. For example, the composite is strongest in the direction of maximum fiber alignment, and weakest in the direction of minimum fiber alignment [26]. Operations such as injection molding, extrusion, or casting are used to process molten suspensions. The liquid suspension is then solidified, holding the fibers in place. Once the fiber orientations are known, mathematical models can predict the mechanical properties of the suspension [1]. Modeling the fibers in the liquid sate is important because the distribution and orientation of the fibers depend strongly on the flow history of the melt during processing. In addition, the rheological properties of the suspension, which govern the flow, depend themselves on the suspended fibers. Thus there exists a relationship between the particle distribution and orientations and the fluid flow.

The study of suspensions has a long history. In his 1906 Ph.D. thesis, Einstein [19] calculated the effective viscosity of a dilute suspension of spherical particles in a fluid of viscosity μ. He found that if the suspension is dilute enough so that hydrodynamic interactions play no role, the effective

viscosity of the suspension is

$$\mu_{\text{eff}} = \mu(1 + 2.5\phi),$$

where ϕ is the particle volume fraction. This relationship is only valid for very dilute suspensions, those with $\phi < 0.05$. For suspensions with a higher volume fraction, hydrodynamic interactions between particles are important and higher-order terms in ϕ are needed to estimate the effective viscosity. For an extensional flow, Batchelor [8] computed the extensional viscosity of a suspension of spheres assuming pair interactions are dominant,

$$\mu_{\text{eff}} = \mu(1 + 2.5\phi + 6.95\phi^2).$$

Even in this relatively simple case, the shear viscosity of a suspension of spheres demonstrates shear-thickening, a non-Newtonian behavior.

The contribution of non-spherical particles to the viscosity of a suspension is more complicated. In a shear flow, for instance, non-spherical rigid bodies contribute less to the effective viscosity when they are aligned with the flow direction. Thus their contribution is not static and their dynamics must be considered. The motion of a single ellipsoid fiber in shear flow was analyzed by Jeffery [35]. In $\mathbb{R}^2$, an ellipse with length ℓ and diameter d suspended in a shear flow with shear rate $\dot{\gamma}$ rotates with period $(\pi/\dot{\gamma})(\lambda + \lambda^{-1})$, where λ is the aspect ratio ℓ/d. The period of the "Jeffery orbit" increases with λ, and the amount of time the fiber spends aligned with the flow direction also increases with λ. Given suitable corrections, the trajectories of fibers of different elongated shapes, such as cylinders, can also be described as Jeffery orbits.

The concentration of the fibers plays a critical role in the properties of rigid body suspension. We quantify the concentration of a suspension of fibers by defining ν to be the average number of fibers per unit volume. If $\nu < 1/\ell^3$, the suspension is in the dilute regime and interactions between fibers are rare. In this case the fiber trajectories can be quantitatively described by Jeffery orbits.

Outside the dilute regime, interactions between fibers are more frequent and must be considered. To define the orientation of a fiber, we assign to each fiber a unit vector $\mathbf{p}$ that points in the direction of its semi-major axis (Figure 1.1). In the semi-dilute regime, $1/\ell^3 < \nu < 1/d\ell^2$, Batchelor related the average stress tensor $\boldsymbol{\sigma}$ to the distribution of fiber orientations $\mathbf{p}$ and the rate of strain tensor $\mathbf{e} = (\nabla\mathbf{u} + (\nabla\mathbf{u})^T)/2$. Letting $\Psi(\mathbf{p}, \mathbf{x})$ be the probability of a fiber at location $\mathbf{x}$ being oriented in

the direction $\mathbf{p}$, we define the ensemble averages

$$a_{ij} = \int p_i p_j \Psi(\mathbf{p}, \mathbf{x}) \, \mathrm{d}\mathbf{p}, \qquad a_{ijk\ell}(\mathbf{x}) = \int p_i p_j p_k p_\ell \Psi(\mathbf{p}, \mathbf{x}) \, \mathrm{d}\mathbf{p}.$$

In the absence of Brownian motion and assuming purely hydrodynamic interactions between fibers (i.e. no external forces and torques), and a slender body approximation ($\lambda \gg 1$), the stress tensor can be approximated as

$$\sigma_{ij} = 2\mu e_{ij} + \nu\zeta a_{ijk\ell} e_{k\ell}, \tag{1.1}$$

where ζ is a drag coefficient that depends on the size and concentration of the fibers and the fluid viscosity [8].

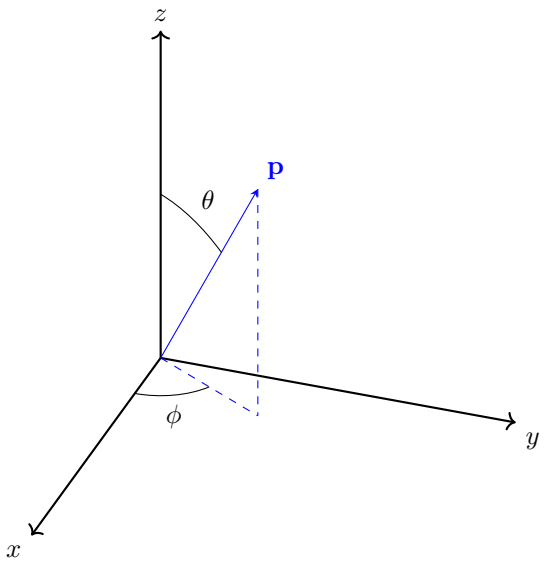

Figure 1.1: Sketch of the orientation vector $\mathbf{p}$.

Computational simulations of fiber suspensions can be either implicit or explicit. Implicit models treat the suspension as a continuum, while explicit models track each fiber individually. In an implicit approach, a suitable fluid model (Navier-Stokes, Stokes) uses the stress tensor (1.1) and combines it with the Folgar-Tucker model [26, 34] that governs the evolution of a_{ij}

$$\frac{D}{Dt} a_{ij} = \frac{1}{2}\Omega_{ik} a_{kj} + \frac{1}{2} a_{ik}\Omega_{kj} + \frac{1}{2}\lambda \left(e_{ik} a_{kj} + a_{ik} e_{kj} - 2 e_{k\ell} a_{ijk\ell} \right) + C_I ||\mathbf{e}|| (2\delta_{ij} - 6 a_{ij}).$$

Here $\mathbf{\Omega} = (\nabla \mathbf{u} - (\nabla \mathbf{u})^T)/2$ is the vorticity tensor and C_I is an empirically derived interaction coefficient that depends on the volume fraction. The limitations of this approach are:

- it requires closure relations and fiber interaction coefficients based on empirical data,
- all the rods must be uniform and slender,
- the model is invalid in the concentrated regime.

In the concentrated regime, $\nu > 1/d\ell^2$, the interactions between fibers dominate and the contribution to the total stress can only be computed by adding the force from each fiber individually [5,48]. To do this, explicit models are required. Explicit models are also useful in the dilute and semi-dilute regimes to model irregularly shaped bodies, or to compute interaction constants and tensor closures to be used in the Folgar-Tucker model. The goal of this dissertation is to develop techniques that are capable of accurately and stably simulating dense suspensions of rigid bodies. This method will be computationally efficient, robust, and will resolve rigid bodies of arbitrary shape over a wide range of concentrations.

1.1 Related Work

Many explicit models represent rigid bodies as prolate ellipsoids [5], sets of connected beads [36, 80], rods [47, 71], or slender bodies [7, 22, 76]. Among explicit methods, we can define two general groups. The first models the fluid phase and updates the bodies by computing the velocity of the fluid on the surface of the particles. The immersed boundary method [56], level set methods [18], lattice Boltzmann methods [42, 43], smoothed particle hydrodynamics [59], and dissipative particle hydrodynamics [58] all fall within this category. An advantage of these methods is that fluid inertia effects can be included by modeling the fluid equations with the Navier-Stokes equations. Moreover, if appropriate, the solvent can be non-Newtonian. A disadvantage of explicit methods is that it requires computing the fluid velocity at all points inside the fluid domain, something that adds significant computational expense, but may not be of interest. In addition, these methods often struggle with resolving unbounded or periodic domains.

In problems where the fluid is Newtonian and inertia is negligible, the problem simplifies considerably. In this case, the fluid equations become the Stokes equations, and potential theory [44] allows us to compute the rigid body velocities without solving for the fluid velocity inside the fluid domain. This allows us to determine the motion of the bodies solely based on the configuration of the particles and the background flow. Stokesian dynamics and boundary integral equations both

make use of potential theory to compute the motions of the particles without directly computing the velocity in the bulk solvent.

As mentioned, the hydrodynamic interactions between bodies must be considered when simulating suspensions. If two bodies are well-separated, their hydrodynamic interactions are well-approximated with a multipole expansion of a few terms. This means that the hydrodynamic interactions of well-separated bodies are almost independent of the body's shape. For nearly touching bodies, however, the interaction is much stronger and depends strongly on the body's shape. That said, if the bodies have simple shapes, for example spheres or ellipsoids, the interactions of nearly touching bodies can be precomputed using lubrication theory. Stokesian dynamics [11, 16] combines the multipole representation for bodies that are far apart with lubrication theory to compute interactions between nearly touching bodies. This has been shown to be very effective for a wide range of problems, including those with small Peclet numbers where Brownian motion needs to be considered.

For general shaped rigid bodies however, lubrication theory is not computationally tractable, since the number of possible hydrodynamic interactions of nearly touching bodies that would need to be precomputed is loo large. Instead of Stokesian dynamics, boundary integral equations (BIEs) can be used. BIEs reduce the Stokes equations to an integral over the boundary of the rigid bodies (and solid walls). This leads to a dimension reduction when compared to explicit methods that discretize the entire fluid domain. Since BIEs require only a discretization of the boundary, they can naturally resolve problems with complicated geometries, including unbounded ones. In the context of the Stokes equations, solutions derived from BIEs automatically satisfy the incompressibility condition and the far field condition. This is in contrast to many other methods, for example the finite element method. BIEs have been successfully applied to various kinds of suspensions, including rigid bodies [13, 17, 76], vesicles [63, 64, 68], and drops [73]. In addition, BIEs can be rigorously analyzed using the Fredholm alternative and the spectral theorem. This theory guides numerical methods whose convergence properties are provably optimal. Of course, the BIE formulation must be approximated by applying discretization methods. One discretization technique applies quadrature to a collocation scheme. By using appropriate quadrature, high-order or even spectral accuracy can be achieved. The limitation of BIEs is that the resulting linear system is dense. In contrast, methods that discretize the entire domain result in larger, sparse linear systems.

Nonetheless, by using efficient iterative solvers and fast summation techniques, an $\mathcal{O}(N)$ solver for the dense linear system is possible, where N is the number of discretization points on the boundary of the domain.

Power and Miranda [60, 61] developed an integral equation representation for the Stokes equations. Their formulation is particularly nice since its discretization is high-order, invertible, and compatible with fast algorithms so that it can be efficiently solved. In particular, the condition number of the discrete linear system is independent of the size of the system N. The Power and Miranda representation has been used to simulate suspensions of rigid bodies [13, 17, 76].

The Stokes equations prohibit contact between bodies in finite time, however, numerical errors may cause overlaps between bodies. When using a model that explicitly represents rigid bodies, care must be taken when advancing in time to avoid collisions and overlaps. An algorithm that prevents collisions between rigid bodies, but allows for a large time step size is desirable. Adaptive time stepping and local refinement can help [41, 65], but particularly in concentrated suspensions, adaptive time stepping may require an excessively small time step to ensure contact is avoided.

Repulsion forces [25, 49, 51] are another approach to prevent contact. Popular forces are variants of a Morse or Lennard-Jones potential that grows as a high-order polynomial as two bodies approach [25, 49]. Spring based repulsion forces [37, 85] have also been used. Methods such as these are inherently heuristic and require the choice of tuning parameters. Furthermore, they introduce stiffness, and therefore still require a small time step to maintain stability. Lastly, they still do not guarantee that collisions are avoided. We adopt a method outlined in [51] that guarantees that each time step is collision-free. This method enforces a constraint on the variational form of the Stokes equations. This constraint appears as a repulsion force in the resulting Euler-Lagrange equations (i.e. the regular Stokes equations), and is free of tuning parameters and guarantees that contact is avoided.

1.2 Contributions

Lu *et al.* [51] use a locally-implicit time stepper coupled with a contact algorithm to simulate suspensions of vesicles and rigid bodies. Locally-implicit time steppers treat the inter-particle hydrodynamic interactions explicitly by lagging them from the previous time step. This method yields a block diagonal system to solve at every time step. Unfortunately, due to nearly touching

bodies, the locally-implicit time stepper requires a small step size or a large minimum separation distance to maintain stability. In this dissertation, we couple the contact algorithm [51] with a globally-implicit time stepper. Globally-implicit time steppers treat all hydrodynamic interactions implicitly. This results in a dense linear system to solve at every time step. While each globally-implicit time step is more computationally expensive than a locally-implicit time step, it allows larger time steps and a smaller minimum separation distance between particles. In particular, for certain problems, the locally-implicit time stepper is not stable, even for time steps on the order of 10^{-7}. In contrast, globally-implicit time steppers are stable for an acceptable time step size, a characteristic that is necessary to simulate dense suspensions.

We use this method to investigate concentrated suspensions, alignment angles, and the effective viscosity of suspensions in confined geometries. Compared to previous methods, our method better resolves the interactions between closely touching bodies, and the result is a stable time stepping method that avoids contact. We investigate the effect of the repulsion force on the time reversibility of the suspension and demonstrate that these repulsion forces cause the particles to jump streamlines and therefore break reversibility. In addition, we make qualitative comparisons between our method and other methods and experiments from the literature.

1.3 Limitations

The main limitation of this model is that it is developed in two dimensions. To be useful for composite manufacturing, a three-dimensional model is required. Three-dimensional models for boundary integral equations are well-developed [17,55], though they are more difficult to implement and computationally more expensive. Moreover, they require many more unknowns, meaning that far fewer particles can be simulated in three dimensions at a fixed number of discretization points.

To compute the shear viscosity of a suspension, we have chosen to model the torque exerted on a Couette apparatus. In reality though, wall effects have a large impact on the particle trajectories. For this reason, unbounded periodic simulations, perhaps involving Lees-Edwards boundary conditions [46], are preferred. In the context of BIEs, this boundary condition has not been investigated, however periodic simulations involving boundary integral equations have been developed [2, 67].

The computation of repulsion forces requires the solution of a nonlinear complementarity problem (NCP). This NCP is linearized and this results in a sequence of linear complementarity problems

(LCPs). From a mathematical standpoint, there is no theory that guarantees that these LCPs has a unique solution. This might cause the NCP iterations to stall or diverge, which is something that we observe for concentrated suspensions in bounded domains. In addition, even if the LCPs do have a unique solution, many LCP iterations may be required to converge to the NCP solution. Since each LCP solution requires us to solve the Stokes equations, this can be very computationally expensive.

CHAPTER 2

STOKES EQUATIONS

In this chapter we review the Stokes equations, the governing equations for low Reynolds number flow. These equations are appropriate for particulate flows, since the small length scales, the slow velocity, and the high viscosity of the solvent means that the Reynolds number is small. Using the Stokes equations instead of the more general Navier-Stokes equations provides significant simplifications. The Stokes equations do not contain a nonlinear inertial term, hence the fluid velocity can be written in terms of the boundary data when there are no non-conservative forcing terms. One of the benefits of formulating the Stokes equations as a boundary integral equation is that it reduces the two-dimensional problem to a one-dimensional boundary integral. This leads to a method that naturally resolves complicated moving boundaries.

We will be modeling mobile rigid bodies suspended in a Stokesian fluid. These bodies interact hydrodynamically with each other and with solid walls. The suspended bodies will be rigid, meaning that they cannot change shape; they can only translate and rotate. This allows us to describe the position of a rigid body using only its center and an orientation angle. Its velocity can likewise be specified by a translational and a rotational velocity.

2.1 Fluid Equations

We are interested in modeling suspensions of rigid bodies, and this requires a good description of the state of the solvent. Given a domain $V \subset \mathbb{R}^2$ with boundary S, the state of a fluid can be described by the incompressible Navier-Stokes equations,

$$\frac{\partial \mathbf{u}}{\partial t} + \mathbf{u} \cdot \nabla \mathbf{u} = -\frac{1}{\rho}\nabla p + \nu \Delta \mathbf{u} + \mathbf{f}, \qquad \mathbf{x} \in V,\ t \geq 0, \tag{2.1a}$$

$$\nabla \cdot \mathbf{u} = 0, \qquad \mathbf{x} \in V,\ t \geq 0, \tag{2.1b}$$

where ρ is the (constant) density, p is the pressure, $\mathbf{u}$ is the velocity, ν is the kinematic viscosity, and $\mathbf{f}$ is an external body force acting on the fluid. Equation (2.1a) is a statement of the conservation of momentum, and (2.1b), the incompressibility condition, is a statement of conservation of mass. To

close the formulation, boundary conditions and an initial condition for $\mathbf{u}$ are required. We consider Dirichlet boundary conditions, along with an arbitrary initial condition

$$\begin{aligned} \mathbf{u}(\mathbf{x}, t) &= \mathbf{g}(\mathbf{x}, t), \qquad \mathbf{x} \in S, t \geq 0 \\ \mathbf{u}(\mathbf{x}, 0) &= \mathbf{u}_0(\mathbf{x}), \qquad \mathbf{x} \in V. \end{aligned}$$

To satisfy global conservation of mass, the boundary condition $\mathbf{g}$ must satisfy the no-flux compatibility condition

$$\int_S \mathbf{g} \cdot \mathbf{n} \, \mathrm{d}S = 0,$$

where $\mathbf{n}$ is the unit normal on S. Since we are restricting ourselves to planar flows, the velocity $\mathbf{u}$ is a vector in $\mathbb{R}^2$.

Following well-established techniques for non-dimensionalization, we pick a characteristic length scale ℓ, a characteristic velocity scale U, a characteristic time scale τ, and make the following substitutions:

$$\mathbf{x}^* = \mathbf{x}/\ell, \qquad \mathbf{u}^* = \mathbf{u}/U, \qquad t^* = t/\tau, \qquad p^* = p\ell/(\rho\nu U), \qquad \mathbf{f}^* = \mathbf{f}/|\mathbf{f}|. \tag{2.2}$$

This leads to the non-dimensionalized Navier-Stokes equations,

$$\begin{aligned} St\frac{\partial \mathbf{u}^*}{\partial t^*} + \mathbf{u}^* \cdot \nabla^* \mathbf{u}^* &= -\frac{1}{Re}\nabla^* p^* + \frac{1}{Re}\Delta^* \mathbf{u}^* + \frac{1}{Fr^2}\mathbf{f}^*, \\ \nabla \cdot \mathbf{u}^* &= 0, \end{aligned}$$

where $St = \ell/(U\tau)$ is the Strouhal number, $Fr = U/\sqrt{\ell|\mathbf{f}|}$ is the Froude number, and $Re = \ell U/\nu$ is the Reynolds number. If $Re \ll 1$ and $St \ll 1/Re$ then the terms with $1/Re$ dominate and the remaining terms can be neglected. This gives the nondimensional steady incompressible Stokes equations,

$$-\Delta \mathbf{u} + \nabla p = \frac{1}{Fr^2}\mathbf{f}, \qquad \mathbf{x} \in V, \tag{2.3a}$$

$$\nabla \cdot \mathbf{u} = 0, \qquad \mathbf{x} \in V, \tag{2.3b}$$

$$\mathbf{u} = \mathbf{u}_b, \qquad \mathbf{x} \in S. \tag{2.3c}$$

For notational clarity we have dropped the asterisk superscript, however, $\mathbf{u}$, p, $\mathbf{f}$, and all the derivatives are non-dimensionalized according to (2.2). For the remainder of this dissertation,

whenever we speak of the Stokes equations, we mean the incompressible, steady, nondimensional, Stokes equations (2.3). The Stokes equations are the governing equations for small scale, slow moving, viscous flows. These are precisely the kind of flows encountered in most particulate flows.

Defining the stress tensor $\boldsymbol{\sigma}$ to be

$$\boldsymbol{\sigma} = -p\mathbf{I} + \frac{1}{2}\left(\nabla\mathbf{u} + (\nabla\mathbf{u})^T\right),$$

the momentum balance (2.3a) can be written as

$$\nabla \cdot \boldsymbol{\sigma} = \frac{1}{Fr}\mathbf{f}, \qquad \mathbf{x} \in V.$$

The stress tensor has a static contribution from the pressure and a dynamic contribution from viscous forces and is often written as,

$$\boldsymbol{\sigma} = -p\mathbf{I} + \mathbf{e},$$

where $\mathbf{e} = (\nabla\mathbf{u} + (\nabla\mathbf{u})^T)/2$ is the strain rate tensor and represents the contribution of viscous forces to the total stress.

In many practical cases, $\mathbf{f}$ is a conservative field and can therefore be expressed as the gradient of a scalar field and included in the pressure gradient, thus leaving a homogeneous equation. In this case, the momentum balance for the Stokes equations is

$$\nabla \cdot \boldsymbol{\sigma} = 0, \qquad \mathbf{x} \in V.$$

The Stokes equations represent a significant simplification of the Navier-Stokes equations. By neglecting the nonlinear $\mathbf{u}\cdot\nabla\mathbf{u}$ term and the time dependent $\partial\mathbf{u}/\partial t$ term, we obtain a set of linear, time-independent, elliptic equations. Because of the linearity, the superposition principle applies. This makes the Stokes equations much easier to analyze mathematically than the Navier-Stokes equations. For example, it is known that given appropriate boundary conditions, a solution to the three-dimensional Stokes equations exists and is unique (up to a constant for the pressure). By contrast, existence and uniqueness of the three-dimensional incompressible Navier-Stokes equations is an unresolved problem, and is one of the Millennium Prize problems [15].

2.2 Rigid Body Motion

This dissertation is concerned with rigid (i.e. non-deforming) bodies. When such a body moves, it must do so without changing shape, meaning that it can only translate and rotate. For a body

in $\mathbb{R}^3$ with a center of rotation $\mathbf{c}$, all rigid body motions can be expressed as

$$\mathbf{u}^{\mathrm{RBM}}(\mathbf{x}) = \mathbf{U} + \boldsymbol{\omega} \times (\mathbf{x} - \mathbf{c}), \tag{2.4}$$

where $\mathbf{U} \in \mathbb{R}^3$ is a constant translational velocity and $\boldsymbol{\omega} \in \mathbb{R}^3$ is a constant angular velocity. In $\mathbb{R}^2$, $\boldsymbol{\omega} = (0, 0, \omega)$, so a rigid body motion is

$$\mathbf{u}^{\mathrm{RBM}} = \mathbf{U} + \omega(\mathbf{x} - \mathbf{c})^{\perp},$$

where $\mathbf{x}^{\perp} = \langle -x_2, x_1 \rangle$. Once we know the translational and rotational velocities of a body undergoing a rigid body motion, equation (2.4) gives the velocity at any point inside that body.

Proposition 2.2.1 shows that the stress field associated with rigid body motion is zero everywhere. The associated proof uses the standard third-order Levi-Civita tensor to denote the cross product in $\mathbb{R}^3$, however the result holds in $\mathbb{R}^2$, where $\omega_1 = \omega_2 = 0$.

Proposition 2.2.1. *The stress tensor of a rigid body motion is zero.*

Proof. Twice the strain rate $2\mathbf{e} = (\nabla \mathbf{u}^{\mathrm{RBM}} + (\nabla \mathbf{u}^{\mathrm{RBM}})^T)$ of a rigid body motion satisfies

$$\begin{aligned}
2e_{ij} &= \partial_i \left(U_j + \epsilon_{jk\ell}\omega_k(x_\ell - c_\ell)\right) + \partial_j \left(U_j + \epsilon_{ik\ell}\omega_k(x_\ell - c_\ell)\right) \\
&= \epsilon_{jk\ell}\partial_i \left(\omega_k(x_\ell - c_\ell)\right) + \epsilon_{ik\ell}\partial_j \left(\omega_k(x_\ell - c_\ell)\right) \\
&= \epsilon_{jk\ell}\omega_k\delta_{i\ell} + \epsilon_{ikl}\omega_k\delta_{j\ell} \\
&= \epsilon_{jki}\omega_k + \epsilon_{ikj}\omega_k.
\end{aligned}$$

Since, by definition, $\epsilon_{jki} = -\epsilon_{ikj}$, it follows that $e_{ij} = 0$.

By equation (2.3a), $\Delta \mathbf{u}^{\mathrm{RBM}} = \nabla p$. Since $\mathbf{u}^{\mathrm{RBM}}$ is linear in $\mathbf{x}$, $\Delta \mathbf{u}^{\mathrm{RBM}} = \nabla p = 0$. Thus, the pressure of a rigid body motion is constant, and typically is taken to be zero. Therefore, the total stress of a rigid body motion, $\boldsymbol{\sigma} = -p\mathbf{I} + \mathbf{e}$, is zero. □

Each rigid body is subject to a net force and torque. For a rigid body with boundary S, unit normal $\mathbf{n}$, and center of rotation $\mathbf{c}$, the hydrodynamic component of the force and torque $\mathbf{F}$ and L, respectively, are

$$\mathbf{F} = \int_S \mathbf{f} \, \mathrm{d}S, \qquad L = \int_S (\mathbf{x} - \mathbf{c})^{\perp} \cdot \mathbf{f} \, \mathrm{d}S,$$

where $\mathbf{f} = \boldsymbol{\sigma} \cdot \mathbf{n}$ is the fluid traction acting on the body. In order to study how rigid bodies move under the influence of forces and torques, we need to solve the equations of motion,

$$m\frac{d\mathbf{U}}{dt} = \mathbf{F}, \qquad I\frac{d\omega}{dt} = L, \tag{2.5}$$

where m is the mass of the body and I is its rotational inertia. Assuming the body has a constant density ρ, then $I = \rho \int_A |\mathbf{x} - \mathbf{c}|^2 \, \mathrm{d}A$, where A is the cross-sectional area of the particle. Scaling time in (2.5) by ℓ/U and the force by $\nu U \ell/\rho$ we have

$$Re\frac{m}{\rho\ell^3}\frac{d\mathbf{U}}{dt} = \mathbf{F}, \qquad Re\frac{I}{\rho\ell^5}\frac{d\omega}{dt} = L.$$

Thus when Re $\ll 1$, the bodies undergo no acceleration. This is the basis for the *quasistatic* approximation. The fact that the bodies are not accelerating means that (2.5) is not needed, and that the body center $\mathbf{c}$ and angle θ evolve according to

$$\frac{d\mathbf{c}}{dt} = \mathbf{U}, \qquad \frac{d\theta}{dt} = \omega.$$

The quasistatic approximation allows us to solve a sequence of steady Stokes equations, even as the rigid bodies move and the geometry changes. The fluid is assumed to instantaneously adjust to the new geometry.

2.3 Integral Representation

In the homogeneous case (when $\mathbf{f}$ is conservative), potential theory and Green's identity allow us to represent the solution to the Stokes equations as an integral of the fluid velocity and traction on S. This leads to a dimension reduction, and once we know the velocity and traction on the boundary, the solution at any point inside V is represented as a boundary integral. The major advantage of this dimension reduction is that it simplifies discretizing complicated geometries.

The first step to develop an integral equation formulation is to compute the fundamental solution of (2.3). This is a pair of functions $\mathbf{u}(\mathbf{x})$ and $p(\mathbf{x})$ satisfying

$$\nabla \cdot \sigma = \Delta\mathbf{u} - \nabla p = -\mathbf{F}\delta(\mathbf{x} - \mathbf{x}_0), \qquad \mathbf{x} \in \mathbb{R}^2, \tag{2.6a}$$

$$\nabla \cdot \mathbf{u} = 0, \qquad \mathbf{x} \in \mathbb{R}^2, \tag{2.6b}$$

where $\delta(\mathbf{x} - \mathbf{x}_0)$ is the Dirac delta function centered at $\mathbf{x}_0$. We note that the choice of sign on the right hand side of (2.6a) is arbitrary and an equivalent formulation exists if a positive sign is

used. A detailed discussion of distributions, including the Dirac delta function, is found in many textbooks, e.g. [29]. For our purposes, however, (2.6a) means that:

1. $\nabla \cdot \boldsymbol{\sigma} = 0$ for all $\mathbf{x} \neq \mathbf{x}_0$,
2. $\int_{\mathbb{R}^2} \nabla \cdot \boldsymbol{\sigma} \, \mathrm{d}V = -\mathbf{F}$.

In $\mathbb{R}^2$

$$\mathbf{u}(\mathbf{x}) = \frac{\mathbf{F} \cdot \mathcal{G}(\mathbf{x}, \mathbf{x}_0)}{4\pi}, \qquad p(\mathbf{x}) = \frac{\mathbf{F} \cdot \mathcal{P}(\mathbf{x}, \mathbf{x}_0)}{4\pi},$$

satisfies (2.6) [62], where

$$\mathcal{G}_{ij}(\mathbf{x}, \mathbf{x}_0) = -\delta_{ij} \ln r + \frac{r_i r_j}{r^2}, \qquad \mathcal{P}_j(\mathbf{x}, \mathbf{x}_0) = \frac{2r_j}{r^2} + P_j^\infty,$$

and $\mathbf{r} = \mathbf{x} - \mathbf{x}_0$, $r = |\mathbf{r}|$, $\mathcal{G}(\mathbf{x}, \mathbf{x}_0)$ is called the Oseen tensor, and $\mathcal{G}(\mathbf{x}, \mathbf{x}_0)/(4\pi)$ is called the Green's dyadic. The pressure only appears as a gradient in the Stokes equations and is therefore only determined up to the constant P_j^∞. By combining the velocity and the pressure, the corresponding stress tensor is

$$\sigma_{ij}(\mathbf{x}) = \frac{T_{ijk}(\mathbf{x}, \mathbf{x}_0) F_k}{4\pi},$$

where

$$T_{ijk}(\mathbf{x}, \mathbf{x}_0) = -4 \frac{r_i r_j r_k}{r^4}.$$

The solution of any eliptic homogeneous PDE can be represented in terms of integrals along the boundary of the domain of the solution and its derivatives. An example is Green's third identity for the Laplace equation. If $w(\mathbf{x})$ is a harmonic function, then

$$w(\mathbf{x}) = \int_S w(\mathbf{y}) \frac{\partial}{\partial \mathbf{n}} \left(\frac{1}{2\pi} \log |\mathbf{x} - \mathbf{y}| \right) \, \mathrm{d}s(\mathbf{y}) - \frac{1}{2\pi} \int_S \log |\mathbf{x} - \mathbf{y}| \frac{\partial w}{\partial \mathbf{n}} \, \mathrm{d}s(\mathbf{y}), \qquad \mathbf{x} \in V.$$

The solution to the Stokes equations can be written in terms of boundary integrals of $\mathbf{u}$ and the surface traction $\mathbf{f} = \boldsymbol{\sigma} \cdot \mathbf{n}$, where $\mathbf{n}$ is the unit normal pointing out of the fluid domain. By applying the Lorenz reciprocal theorem [39, 44, 62]

$$\mathbf{u}(\mathbf{x}) = -\int_S \frac{f_i(\mathbf{y}) \mathcal{G}_{ij}(\mathbf{x}, \mathbf{y})}{4\pi} \, \mathrm{d}S(\mathbf{y}) + \int_S u_i(\mathbf{y}) T_{ijk}(\mathbf{x}, \mathbf{y}) n_k(\mathbf{y}) \, \mathrm{d}S(\mathbf{y}), \qquad \mathbf{x} \in V. \tag{2.7}$$

Equation (2.7) expresses $\mathbf{u}(\mathbf{x})$ for $\mathbf{x} \in V$ in terms of $\mathbf{u}$ and $\mathbf{f}$ on S.

The first term in (2.7),

$$\mathcal{S}[\mathbf{f}](\mathbf{x}) = \int_S \frac{f_i(\mathbf{y})\mathcal{G}_{ij}(\mathbf{x},\mathbf{y})}{4\pi}\, \mathrm{d}S(\mathbf{y}), \tag{2.8}$$

is known as the *single-layer potential*. This represents the velocity field generated by the force distribution $\boldsymbol{\sigma}(\mathbf{y}) \cdot \mathbf{n}(\mathbf{y})$ along S. The second term,

$$\mathcal{D}[\mathbf{u}](\mathbf{x}) = \int_S u_i(\mathbf{y}) T_{ijk}(\mathbf{x},\mathbf{y}) n_k(\mathbf{y})\, \mathrm{d}S(\mathbf{y}), \tag{2.9}$$

is known as the *double-layer potential*. Thus (2.7) can be written as

$$\mathbf{u}(\mathbf{x}) = -\mathcal{S}[\mathbf{f}](\mathbf{x}) + \mathcal{D}[\mathbf{u}](\mathbf{x}), \qquad \mathbf{x} \in V.$$

We will on occasion write the single-layer and double-layer potentials as

$$\mathcal{S}[\mathbf{f}](\mathbf{x}) = \int_S f_i(\mathbf{y}) W_{ij}(\mathbf{x},\mathbf{y})\, \mathrm{d}S(\mathbf{y}), \qquad \mathcal{D}[\mathbf{u}](\mathbf{x}) = \int_S u_i(\mathbf{y}) K_{ij}(\mathbf{x},\mathbf{y})\, \mathrm{d}S(\mathbf{y}),$$

where $\mathbf{W}$ and $\mathbf{K}$ are the kernels of the single- and double-layer potentials, respectively, given by

$$W_{ij}(\mathbf{x},\mathbf{y}) = \frac{\mathcal{G}_{ij}(\mathbf{x},\mathbf{y})}{4\pi}, \tag{2.10a}$$

$$K_{ij}(\mathbf{x},\mathbf{y}) = T_{ijk}(\mathbf{x},\mathbf{y}) n_k(\mathbf{y}). \tag{2.10b}$$

The point $\mathbf{y} \in S$ in the layer potentials is known as the source point and the point $\mathbf{x} \in V \cup S$ is known as the target point. While equation (2.7) holds for any target point $\mathbf{x} \in V$, it is false if $\mathbf{x} \in S$. This is due to a jump in the double-layer potential. If we let $\mathbf{x} \in S$ and assume the boundary has a continuous tangent vector, the double-layer potential satisfies the following conditions for $\mathbf{x} \in S$:

$$\mathcal{D}[\mathbf{u}](\mathbf{x})_{(i)} := \lim_{\epsilon \to 0} \mathcal{D}[\mathbf{u}](\mathbf{x} - \epsilon \mathbf{n}(\mathbf{x})) = -\frac{1}{2}\mathbf{u}(\mathbf{x}) + \mathcal{D}[\mathbf{u}](\mathbf{x}), \tag{2.11a}$$

$$\mathcal{D}[\mathbf{u}](\mathbf{x})_{(e)} := \lim_{\epsilon \to 0} \mathcal{D}[\mathbf{u}](\mathbf{x} + \epsilon \mathbf{n}(\mathbf{x})) = +\frac{1}{2}\mathbf{u}(\mathbf{x}) + \mathcal{D}[\mathbf{u}](\mathbf{x}). \tag{2.11b}$$

However, the single-layer potential does not have such a jump and satisfies:

$$\mathcal{S}[\mathbf{f}](\mathbf{x})_{(i)} := \lim_{\epsilon \to 0} \mathcal{S}[\mathbf{f}](\mathbf{x} - \epsilon \mathbf{n}(\mathbf{x})) = \mathcal{S}[\mathbf{f}](\mathbf{x}),$$

$$\mathcal{S}[\mathbf{f}](\mathbf{x})_{(e)} := \lim_{\epsilon \to 0} \mathcal{S}[\mathbf{f}](\mathbf{x} + \epsilon \mathbf{n}(\mathbf{x})) = \mathcal{S}[\mathbf{f}](\mathbf{x}).$$

Therefore, the jumps in the layer potentials at $\mathbf{x} \in S$ are

$$[\![\mathcal{S}[\mathbf{f}](\mathbf{x})]\!] := \mathcal{S}[\mathbf{f}](\mathbf{x})_{(e)} - \mathcal{S}[\mathbf{f}](\mathbf{x})_{(i)} = \mathbf{0},$$

$$[\![\mathcal{D}[\mathbf{u}](\mathbf{x})]\!] := \mathcal{D}[\mathbf{u}](\mathbf{x})_{(e)} - \mathcal{D}[\mathbf{u}](\mathbf{x})_{(i)} = \mathbf{u}(\mathbf{x}).$$

We note that the jump in the traction of the single-layer is $\mathbf{f}$, and this is often used to balance the fluid force with the elastic force of deformable (non-rigid) bodies [63, 64, 68]. In contrast, the traction of the double-layer is continuous through S. This fact is used in Proposition 2.3.1.

Letting $\mathbf{x} \in S$ and using the jump condition of the double-layer potential (2.11a), we have the boundary integral equation (BIE)

$$\mathcal{S}[\mathbf{f}](\mathbf{x}) + \mathcal{D}[\mathbf{u}](\mathbf{x}) = \frac{1}{2}\mathbf{u}(\mathbf{x}), \qquad \mathbf{x} \in S. \tag{2.12}$$

Assuming that we know the velocity on S, (2.12) allows us to compute the traction on S and vice versa. Thus (2.12) is the Stokes analog of the Dirichlet to Neumann map for Laplace's equation. Solving (2.12) followed by applying (2.7) is known as a direct BIE formulation.

Instead of using $\mathbf{u}$ to compute $\mathbf{f}$ (or vice-versa) and then applying (2.7), we use an indirect integral equation formulation. This is done by representing $\mathbf{u}$ with an arbitrary, non-physical, density function of the single- or double-layer potentials. For example, we could represent $\mathbf{u}$ as a single-layer potential with density $\boldsymbol{\zeta}(\mathbf{x})$,

$$\mathbf{u}(\mathbf{x}) = \mathcal{S}[\boldsymbol{\zeta}](\mathbf{x}), \qquad \mathbf{x} \in V. \tag{2.13}$$

Recalling that $[\![\mathcal{S}[\boldsymbol{\zeta}](\mathbf{x})]\!] = \mathbf{0}$, and letting $\mathbf{x} \to S$ gives the indirect BIE

$$\mathcal{S}[\boldsymbol{\zeta}](\mathbf{x}) = \mathbf{g}(\mathbf{x}) \qquad \mathbf{x} \in S. \tag{2.14}$$

That is, (2.14) is solved for the density function $\boldsymbol{\zeta}(\mathbf{x})$ that guarantees the boundary condition $\mathbf{u}(\mathbf{x}) = \mathbf{g}(\mathbf{x})$, $\mathbf{x} \in S$ is satisfied. Then $\mathbf{u}(\mathbf{x})$ is evaluated for $\mathbf{x} \in V$ with (2.13). Likewise, we could represent the velocity as a double-layer potential with an arbitrary function density $\boldsymbol{\eta}(\mathbf{x})$,

$$\mathbf{u}(\mathbf{x}) = \mathcal{D}[\boldsymbol{\eta}](\mathbf{x}), \qquad \mathbf{x} \in V. \tag{2.15}$$

Recalling that $[\![\mathcal{D}[\boldsymbol{\eta}](\mathbf{x})]\!] = \boldsymbol{\eta}(\mathbf{x})$ and letting $\mathbf{x} \to S$ (from inside V) gives the indirect BIE

$$-\frac{1}{2}\boldsymbol{\eta}(\mathbf{x}) + \mathcal{D}[\boldsymbol{\eta}](\mathbf{x}) = \mathbf{g}(\mathbf{x}), \qquad \mathbf{x} \in S. \tag{2.16}$$

As in the single-layer equation (2.14), the density function $\boldsymbol{\eta}$ ensures that the boundary condition for $\mathbf{u}$ is satisfied. We solve (2.16) for the density function $\boldsymbol{\eta}(\mathbf{x})$ and then use the density function in (2.15) to evaluate the $\mathbf{u}(\mathbf{x})$ for $\mathbf{x} \in V$.

Regardless of whether we use a single-layer or double-layer ansatz for $\mathbf{u}$ we must solve a BIE. The Fredholm alternative and spectral theorem allow us to analyze the solvability and stability of the BIEs (2.14) and (2.16). This is discussed in the following section.

2.3.1 Fredholm Equations

A general Fredholm integral equation is

$$\lambda z(\mathbf{x}) - \int_S K(\mathbf{x},\mathbf{y}) z(\mathbf{y}) \, \mathrm{d}S(\mathbf{y}) = g(\mathbf{x}), \qquad \mathbf{x} \in S, \tag{2.17}$$

where λ is a given constant, $g : S \to \mathbb{R}$ is continuous, and $K(\mathbf{x},\mathbf{y}) : S \times S \to \mathbb{R}$ is continuous and at worst weakly singular, i.e.,

$$\int_S \left(\int_S |K(\mathbf{x},\mathbf{y})| \, \mathrm{d}S(\mathbf{y}) \right) \, \mathrm{d}S(\mathbf{x}) < \infty.$$

Typically, the function $g(\mathbf{x})$ is a boundary condition, $K(\mathbf{x},\mathbf{y})$ is related to the fundamental solution of the PDE of interest, and λ depends on the choice of K. We are interested in solving for $z : S \to \mathbb{R}$. If $\lambda = 0$, (2.17) is called a first-kind Fredholm integral equation, otherwise it is a second-kind Fredholm integral equation. Therefore, (2.14) is a first-kind Fredholm integral equation and (2.16) is a second-kind Fredholm integral equation. Fredholm integral equations may be written in the abstract notation

$$\lambda z - \mathcal{K} z = g. \tag{2.18}$$

For example, the double-layer BIE (2.16) can be written in this form, with $\lambda = -1/2$ and $\mathcal{K} = -\mathcal{D}$.

We now discuss the behavior and solvability of (2.14) and (2.16). To begin, some background on compact operators in necessary. The following follows closely to [39] where the discussion is framed in the context of BIEs; more details on compact operators can be found in standard texts on functional analysis, e.g. [29]. Because of the integrability of $K(\mathbf{x},\mathbf{y})$, $\mathcal{K}$ is a compact linear operator on $\mathrm{L}^2(S)$. Therefore, by the spectral theorem, the eigenvalues of $\mathcal{K}$ cluster only at the origin, and it can be approximated to any precision with a finite rank linear operator (i.e. a matrix). This will play a fundamental role when we consider discretizations of BIEs in Chapter 3.

Consider a compact operator on a Hilbert space H ($\mathrm{L}^2(S)$ is a Hilbert space). A compact operator is any operator that maps a bounded set to a relatively compact set. Every finite dimensional bounded set is relatively compact. Compact operators form a subset of bounded operators, meaning that every compact operator is necessarily bounded. It follows that identity operator I is compact if and only if H is finite dimensional, since in an infinite dimensional space, the identity maps bounded sets that are not relatively compact to themselves. Multiplication by a bounded operator from either the left or the right preserves compactness.

From this, we can deduce that an infinite dimensional compact operator $\mathcal{K}$ cannot have a bounded inverse. To show this, we observe that if $\mathcal{K}^{-1}$ is bounded, then $I = \mathcal{K}\mathcal{K}^{-1}$ must also be compact, since multiplication by a bounded operator preserves compactness. This contradicts the statement that the identity is not compact on infinite dimensional spaces. Therefore, the inverse of a first-kind Fredholm integral equation is unbounded, meaning that small changes in the data g lead to arbitrarily large changes to the solution z. At a numerical level, the unboundedness results in a matrix with an increasing condition number as the discretization is refined.

The inverse of a second-kind integral equation, however can be bounded. The Fredholm alternative states that either a second-kind Fredholm integral equation of the form (2.18) has a unique solution for any g, or the corresponding homogeneous equation

$$\lambda z - \mathcal{K}z = 0$$

has nontrivial solutions. This is analogous to the invertible matrix theorem of linear algebra. If the homogeneous equation has nontrivial solutions, then these solutions are null functions of (2.18) and also eigenfunctions of the operator $\mathcal{K}$ corresponding to the eigenvalue λ, since the null functions satisfy

$$\mathcal{K}z = \lambda z.$$

If λ is not an eigenvalue of $\mathcal{K}$, then the null space of (2.18) is empty and its inverse is bounded. Therefore, second-kind equations are well-posed, but care must be taken if λ is an eigenvalue of $\mathcal{K}$.

To determine if (2.18) has a nontrivial null space, we exploit even more symmetries between second-kind Fredholm equations and square matrices. To do this we first define the adjoint, which plays the role of the matrix transpose for compact operators. The adjoint of a compact operator $\mathcal{K}$ is the unique operator $\mathcal{K}^*$ satisfying

$$\langle \mathcal{K}x, y\rangle = \langle x, \mathcal{K}^* y\rangle, \qquad \text{for all } x, y \in \mathrm{H}(S),$$

where $\langle\, \cdot, \cdot\, \rangle$ is the inner product associated with H. If $\mathcal{K}$ is compact, then so is $\mathcal{K}^*$.

If $\lambda \neq 0$ is an eigenvalue, then the corresponding eigenfunctions are null functions of (2.18), and they span the null space $N(\lambda I - \mathcal{K})$. As with square matrices, we have the relation

$$R(\lambda I - \mathcal{K}) = N(\lambda^* I - \mathcal{K}^*)^{\perp},$$

where λ^* is the complex conjugate of λ. Since (2.18) has a solution if and only if $g \in R(\lambda I - \mathcal{K})$, it follows that for a solution to exist g must be orthogonal to the null space of the adjoint equation $\lambda^* I - \mathcal{K}^*$.

Therefore if a null function v for (2.18) exists (for $\lambda \neq 0$), it must satisfy two properties

(a) v must be an eigenfunction for $\mathcal{K}$, corresponding to eigenvalue λ;

(b) v cannot be in $N(\lambda^* I - \mathcal{K}^*)^{\perp}$.

If one or more null functions exist, then to solve (2.18) for an arbitrary right hand side g, we must address the null functions. We discuss modifications to (2.18) that do this in the following sections.

2.3.2 Exterior Flow

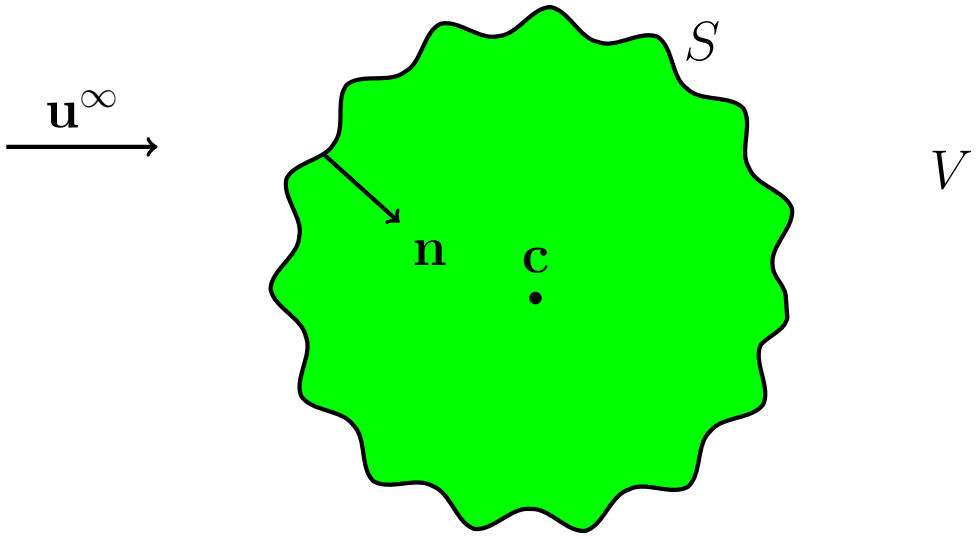

Figure 2.1: Example of an exterior flow problem past a stationary or mobile rigid obstacle. The fluid domain is outside the green region. The normal vector $\mathbf{n}$ points outside the fluid domain and into the obstacle. The point $\mathbf{c}$ can be any point inside the obstacle; in the case of a mobile particle, it is the center of rotation.

We begin by considering an exterior (unbounded) flow around a single rigid obstacle with boundary S centered at $\mathbf{c}$ (Figure 2.1). The obstacle disturbs the prescribed background flow $\mathbf{u}^\infty(\mathbf{x})$ by $\mathbf{u}^D(\mathbf{x})$, meaning that the velocity $\mathbf{u}$ at point $\mathbf{x}$ is $\mathbf{u}(\mathbf{x}) = \mathbf{u}^\infty(\mathbf{x}) + \mathbf{u}^D(\mathbf{x})$. The governing

equations for $\mathbf{u}^D$ are the Stokes equations, along with a no-slip boundary condition on $\mathbf{u}$,

$$\Delta \mathbf{u}^D = \nabla p^D, \qquad \mathbf{x} \in V, \tag{2.19a}$$

$$\nabla \cdot \mathbf{u}^D = 0, \qquad \mathbf{x} \in V, \tag{2.19b}$$

$$\mathbf{u}^D = \mathbf{U} + \omega(\mathbf{x} - \mathbf{c})^\perp - \mathbf{u}^\infty(\mathbf{x}), \qquad \mathbf{x} \in S, \tag{2.19c}$$

where we have used the fact that S is a rigid boundary. The conditions in the far field are

$$\mathbf{u}^D - \mathbf{F} \ln r = O(1), \qquad p^D = O(1), \qquad \text{as } r \to \infty,$$

where $\mathbf{F}$ is the net force on the particle. For $\mathbf{u}^D$ to be bounded as $||\mathbf{x}|| \to \infty$, we require that $\mathbf{F} = \mathbf{0}$. This is known as the two-dimensional Stokes paradox. The Stokes paradox can be understood by looking at the fundamental solution $\mathcal{G}(\mathbf{x}, \mathbf{x}_0)$ for the two-dimensional Stokes equations, which grows as $\ln r$. In the case of multiple obstacles (to be discussed later), for a bounded solution to exist at infinity, we require that the total force on all rigid bodies is zero.

The flow outside a rigid obstacle can be represented uniquely with the single-layer potential [62]. Matching the limiting value of the single-layer potential as $\mathbf{x}$ approaches S with the boundary condition (2.19c), gives the indirect BIE

$$\mathcal{S}[\boldsymbol{\zeta}](\mathbf{x}) = \mathbf{U} + \omega(\mathbf{x} - \mathbf{c})^\perp - \mathbf{u}^\infty(\mathbf{x}), \qquad \mathbf{x} \in S. \tag{2.20}$$

If the translational velocity $\mathbf{U}$ and rotational velocity ω are known, then there is a density function $\boldsymbol{\zeta}$ that satisfies (2.20). This is known as the *resistance problem*. If, on the other hand, the net force $\mathbf{F}$ and net torque L on the obstacle are prescribed, but not the velocity, then (2.20) has more unknowns ($\boldsymbol{\zeta}$, $\mathbf{U}$ and ω) than equations. In this case, we relate the density function to the net force and torque

$$\int_S \boldsymbol{\zeta} \, \mathrm{d}S = \mathbf{F}, \qquad \int_S \boldsymbol{\zeta} \cdot (\mathbf{x} - \mathbf{c})^\perp \, \mathrm{d}S = L. \tag{2.21}$$

Equations (2.20) and (2.21) form an invertible system for the density and the translational and rotational velocities of the obstacle. We note that for a rigid body, the density function $\boldsymbol{\zeta}$ is actually the traction $\mathbf{f}$. This is known as the *mobility problem*. The problem with representing the velocity using a single-layer potential is that the resulting boundary integral equation is ill-posed since its inverse is unbounded (Section 2.3.1).

Alternatively, we can represent the velocity as a double-layer potential. Matching the limiting value of the double-layer potential with the boundary conditions (2.19c), and applying the jump condition (2.11a) gives the indirect BIE

$$-\frac{1}{2}\boldsymbol{\eta}(\mathbf{x}) + \mathcal{D}[\boldsymbol{\eta}](\mathbf{x}) = \mathbf{U} + \omega(\mathbf{x}-\mathbf{c})^{\perp} - \mathbf{u}^{\infty}(\mathbf{x}), \qquad \mathbf{x} \in S, \tag{2.22}$$

which is a Fredholm integral equation of the second-kind for the density function $\boldsymbol{\eta}$. Note that because of the direction of the normal we have used the interior jump condition. If this equation is invertible, then this would be a well-posed problem. Unfortunately, the formulation has a nontrivial null space. Therefore, by the Fredholm alternative solutions are not unique and may not even exist for an arbitrary right hand side.

To show that this formulation has a null space, it suffices to show that $1/2$ is an eigenvalue of the double-layer potential. This is demonstrated by recalling equation (2.12). Letting $\mathbf{u}$ be a rigid body motion on S, then the traction $\mathbf{f} = \boldsymbol{\sigma} \cdot \mathbf{n}$ is zero (by Proposition 2.2.1), and

$$\mathcal{D}[\mathbf{u}](\mathbf{x}) = \frac{1}{2}\mathbf{u}(\mathbf{x}), \qquad \mathbf{x} \in S.$$

Therefore, $1/2$ is an eigenvalue of $\mathcal{D}$, and the homogeneous version of (2.22) has nontrivial solutions, these being rigid body motion density functions. In $\mathbb{R}^2$, rigid body motions are spanned by two translational components and one rotational component,

$$\mathbf{v}_1 = (1,0), \qquad \mathbf{v}_2 = (0,1), \qquad \mathbf{v}_3 = (x_2 - c_2, -(x_1 - c_1)).$$

In fact, for flow around an obstacle, these are the only eigenvectors of the double-layer potential [38, 39, 44, 61].

Since $-1/2 + \mathcal{D}$ has a nontrivial null space, we expect its range to not cover all solutions of the Stokes equations. We now show this is the case in $\mathbb{R}^3$, but the result also holds in $\mathbb{R}^2$.

Proposition 2.3.1. *The double-layer potential can only represent flows that generate no net force or torque on S.*

Proof. Consider fluid inside a domain V with surface S. Given two velocity fields that satisfy the Stokes equations, $\mathbf{v}$ and $\mathbf{v}'$, and corresponding surface traction, $\mathbf{f}$ and $\mathbf{f}'$, the Lorenz reciprocal theorem states [39]

$$\int_S \mathbf{v} \cdot \mathbf{f}' \, \mathrm{d}S = \int_S \mathbf{v}' \cdot \mathbf{f} \, \mathrm{d}S. \tag{2.23}$$

Letting $\mathbf{v}'$ be a rigid body motion about $\mathbf{c} = \mathbf{0}$, i.e. $\mathbf{v}' = \mathbf{U} + \boldsymbol{\omega} \times \mathbf{x}$ and $\mathbf{f}' = \mathbf{0}$,

$$\int_S \mathbf{v}' \cdot \mathbf{f} \, \mathrm{d}S = \int_S (\mathbf{U} + \boldsymbol{\omega} \times \mathbf{x}) \cdot \mathbf{f} \, \mathrm{d}S = 0.$$

In index notation,

$$\int_S (U_i f_i + \epsilon_{ijk} \omega_j x_k f_i) \, \mathrm{d}S = 0.$$

Recalling that $f_i = \sigma_{\ell i} n_\ell$, we apply the divergence theorem to convert the surface integral into a volume integral

$$\int_S (U_i \sigma_{\ell i} n_\ell + \epsilon_{ijk} \omega_j x_k \sigma_{\ell i} n_\ell) \, \mathrm{d}S = \int_V (U_i \partial_\ell \sigma_{\ell i} + \epsilon_{ijk} \omega_j \partial_\ell (x_k \sigma_{\ell i})) \, \mathrm{d}V = 0.$$

Applying the product rule to the last term yields

$$\begin{aligned} 0 &= \int_V (U_i \partial_\ell \sigma_{\ell i} + \epsilon_{ijk} \omega_j (\delta_{k\ell} \sigma_{\ell i} + x_k \partial_\ell \sigma_{\ell i})) \, \mathrm{d}V \\ &= \int_V (U_i \partial_\ell \sigma_{\ell i} + \epsilon_{ijk} \omega_j (\sigma_{ki} + x_k \partial_\ell \sigma_{\ell i})) \, \mathrm{d}V. \end{aligned}$$

By the symmetry of σ, $\int_V \epsilon_{ijk} \omega_j \sigma_{ki} = 0$. Factoring out $\partial_\ell \sigma_{\ell i}$ from the remaining terms,

$$\int_V (U_i + \epsilon_{ijk} \omega_j x_k) \partial_\ell \sigma_{\ell i} \, \mathrm{d}V = 0.$$

Since this must hold for any $\mathbf{U}$ and $\boldsymbol{\omega}$ and over an arbitrary volume V, it follows that

$$\partial_\ell \sigma_{\ell i} = \nabla \cdot \boldsymbol{\sigma} = \mathbf{0}, \qquad \mathbf{x} \in V, \tag{2.24}$$

so V cannot contain any point forces.

Next, the net force and torque on S are

$$\mathbf{F} = \int_S \mathbf{f} \, \mathrm{d}S, \qquad \mathbf{L} = \int_S \mathbf{f} \times \mathbf{x} \, \mathrm{d}S.$$

In index notation after applying the divergence theorem

$$\begin{aligned} F_i &= \int_S \sigma_{ji} n_j \, \mathrm{d}S = \int_V \partial_j \sigma_{ji} \, \mathrm{d}V, \\ L_i &= \int_S \epsilon_{ijk} x_j \sigma_{\ell k} n_\ell \, \mathrm{d}S = \epsilon_{ijk} \int_V \partial_\ell (x_j \sigma_{\ell k}) \, \mathrm{d}V. \end{aligned}$$

By (2.24), the force is zero. Applying the product rule to the torque

$$\begin{aligned} L_i &= \epsilon_{ijk} \int_V (\delta_{j\ell} \sigma_{\ell k} + x_j \partial_\ell \sigma_{\ell k}) \, \mathrm{d}V \\ &= \epsilon_{ijk} \int_V (\sigma_{jk} + x_j \partial_\ell \sigma_{\ell k}) \, \mathrm{d}V. \end{aligned}$$

Again by (2.24) the second part of the volume integral vanishes. The other part of the volume integral,

$$\epsilon_{ijk} \int_V \sigma_{jk} \, \mathrm{d}V,$$

is also zero because of the symmetry of $\boldsymbol{\sigma}$. Since the traction of the double-layer potential is continuous through S, and there is no force or torque generated from the interior flow, it follows that the force and torque of the double-layer representation of an exterior flow is zero. This means that the net force and torque on S must be zero if the flow is represented using only a double-layer potential. □

Proposition 2.3.1 characterizes the velocity fields that cannot be represented with a double-layer potential. Therefore, we must modify or complete the double-layer potential representation so that velocity fields that induce a net force and torque on S are included. We use the completion method of Power and Miranda [38, 60, 61] where terms that induce an arbitrary net force and torque on S are explicitly added to the double-layer representation. These terms are the Stokeslet, $\mathbf{S}[\mathbf{F}, \mathbf{c}](\mathbf{x})$ and the rotlet, $\mathbf{R}[L, \mathbf{c}](\mathbf{x})$,

$$\mathbf{S}[\mathbf{F}, \mathbf{c}](\mathbf{x}) = \left(-\delta_{ij} \ln r + \frac{r_i r_j}{r^2} \right) F_i, \qquad \mathbf{R}[L, \mathbf{c}](\mathbf{x}) = \frac{L \epsilon_{ij} r_j}{r^2},$$

where, in this context, $\mathbf{r} = \mathbf{x} - \mathbf{c}$ and $r = |\mathbf{x} - \mathbf{c}|$. Figure 2.2 shows examples of a Stokeslet and a rotlet velocity field. The Stokeslet is the fundamental solution of the Stokes equations and represents the velocity induced by a point force of strength $\mathbf{F}$. The rotlet represents the velocity induced by a point torque of strength L. On any surface S enclosing the point $\mathbf{c}$, the net torque of a Stokeslet is $\mathbf{F}$, while the net torque of a rotlet is L. On any surface not enclosing $\mathbf{c}$, the net force and torque from the Stokelet and the rotlet are zero.

For an exterior flow around a surface S, the completed double-layer potential is

$$\mathbf{u}^{\mathrm{DLP}}(\mathbf{x}) = \mathcal{D}[\boldsymbol{\eta}](\mathbf{x}) + \mathbf{S}[\mathbf{F}, \mathbf{c}](\mathbf{x}) + \mathbf{R}[L, \mathbf{c}](\mathbf{x}), \tag{2.25}$$

where $\mathbf{c}$ is a point enclosed by S. This completed double-layer potential has a full range and can represent any solution of the Stokes equations [60, 61]. Matching (2.25) to the boundary condition (2.19c) and applying the jump condition (2.11a) as before yields

$$-\frac{1}{2}\boldsymbol{\eta}(\mathbf{x}) + \mathcal{D}[\boldsymbol{\eta}](\mathbf{x}) + \mathbf{S}[\mathbf{F}, \mathbf{c}](\mathbf{x}) + \mathbf{R}[L, \mathbf{c}](\mathbf{x}) = \mathbf{U} + \omega(\mathbf{x} - \mathbf{c})^{\perp} - \mathbf{u}^{\infty}, \qquad \mathbf{x} \in S. \tag{2.26}$$

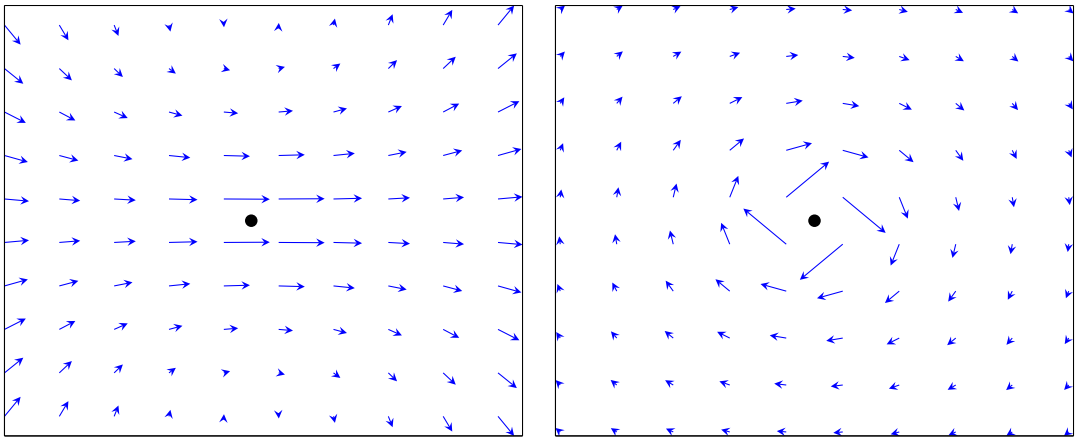

Figure 2.2: The Stokeslet (left) and rotlet (right) velocity fields. The Stokeslet has strength $\mathbf{F} = (1, 0)$ and the rotlet has strength $L = 1$.

To close the system, the Power and Miranda completion relates the net force and torque to the density function by

$$\int_S \boldsymbol{\eta} \, \mathrm{d}S = \mathbf{F}, \qquad \int_S \boldsymbol{\eta} \cdot (\mathbf{x} - \mathbf{c})^{\perp} \, \mathrm{d}S = L. \tag{2.27}$$

Equations (2.26) and (2.27) can be used to solve:

(a) *the mobility problem*, where $\mathbf{F}$ and L are specified and $\mathbf{U}$, ω and $\boldsymbol{\eta}$ are solved for; or

(b) *the resistance problem*, where $\mathbf{U}$ and ω are specified and $\mathbf{F}$, $\mathbf{L}$ and $\boldsymbol{\eta}$ are solved for.

Note that (2.26) along with the closure (2.27) is still a second-kind Fredholm integral equation, so the benefits of the second-kind formulation are not lost by this completion.

2.3.3 Interior Flow

We now turn our attention to flow inside an interior simply-connected domain. Such a geometry is sketched in Figure 2.3. On S, we have the boundary condition on the velocity $\mathbf{u} = \mathbf{g}$. We have already seen that the single-layer representation is ill-posed, so we will not consider that representation here. As with the exterior flow, we represent the solution at any point $\mathbf{x} \in V$ as a double-layer potential. Doing so leads to the second-kind BIE,

$$-\frac{1}{2}\boldsymbol{\eta}(\mathbf{x}) + \mathcal{D}[\boldsymbol{\eta}](\mathbf{x}) = \mathbf{g}(\mathbf{x}), \qquad \mathbf{x} \in S. \tag{2.28}$$

Where again we are using the limiting value coming from the interior of the domain (2.11a).

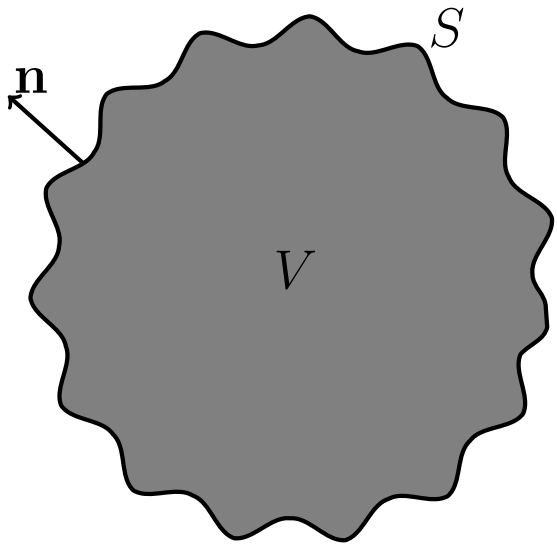

Figure 2.3: Example of a bounded, simply connected domain. The fluid region is colored in gray. The normal points out of the fluid domain.

Unlike the exterior problem, the rigid-body densities are not null functions of (2.28), since they are not eigenfunctions corresponding to eigenvalue 1/2. However, this formulation does still have a nontrivial null space [38, 44, 60].

Proposition 2.3.2. *The double-layer potential for an interior flow has an eigenvector that is not orthogonal to* $\mathbf{n}$*, the outward unit normal on* S*.*

Proof. To investigate the null space of (2.28), we look at the corresponding homogeneous adjoint problem

$$-\frac{1}{2}\psi_i(\mathbf{x}) + \int_S K_{ji}(\mathbf{y},\mathbf{x})\psi_i(\mathbf{y})\ \mathrm{d}S(\mathbf{y}) = 0. \tag{2.29}$$

Letting $\psi_i = n_i$,

$$\begin{aligned}
-\frac{1}{2}n_i(\mathbf{x}) + \int_S K_{ji}(\mathbf{y},\mathbf{x})n_i(\mathbf{y})\ \mathrm{d}S(\mathbf{y}) &= -\frac{1}{2}n_i(\mathbf{x}) - \int_S \frac{4r_j r_i r_k n_k(\mathbf{x})}{r^4} n_i(\mathbf{y})\ \mathrm{d}S(\mathbf{y}) \\
&= -\frac{1}{2}n_i(\mathbf{x}) - \int_S \frac{4r_j r_i r_k n_i(\mathbf{y})}{r^4} n_k(\mathbf{x})\ \mathrm{d}S(\mathbf{y}) \\
&= -\frac{1}{2}n_i(\mathbf{x}) - \int_S K_{ik}(\mathbf{x},\mathbf{y})n_k(\mathbf{x})\ \mathrm{d}S(\mathbf{y}) \\
&= -\frac{1}{2}n_i(\mathbf{x}) - n_k(\mathbf{x})\int_S K_{ik}(\mathbf{x},\mathbf{y})\ \mathrm{d}S(\mathbf{y}) \\
&= -\frac{1}{2}n_i(\mathbf{x}) + \frac{1}{2}n_k(\mathbf{x})\delta_{ik} = 0,
\end{aligned}$$

where we have used the identity [44]

$$\int_S K_{ij}\ \mathrm{d}S = -\frac{1}{2}\delta_{ij}.$$

By the Fredholm alternative, for a solution of (2.28) to exist, the boundary condition $\mathbf{g}(\mathbf{x})$ must be orthogonal to $\mathbf{n}$. In other words

$$\int_S \mathbf{g}\cdot\mathbf{n}\ \mathrm{d}S = 0.$$

Note that this is the no-flux condition required to satisfy the incompressibility condition $\nabla\cdot\mathbf{u} = 0$. □

It can be shown that $\mathbf{n}$ is the only null function of the adjoint problem [38, 44, 60]. For simply-connected domains in $\mathbb{R}^2$, this orthogonality condition is in fact always satisfied as it is the no-flux condition that guarantees global conservation of mass. Now that we have determined the null space of the adjoint equation, we modify the BIE (2.28) so that a unique solution is guaranteed. We use Wieland deflation [38] to remove the null function by adding the operator

$$\mathcal{N}_0[\boldsymbol{\eta}](\mathbf{x}) = \mathbf{n}(\mathbf{x})\int_S \mathbf{n}\cdot\boldsymbol{\eta}\ \mathrm{d}S, \qquad \mathbf{x}\in S$$

to (2.28). The final system is

$$-\frac{1}{2}\boldsymbol{\eta}(\mathbf{x}) + \mathcal{D}[\boldsymbol{\eta}](\mathbf{x}) + \mathcal{N}_0[\boldsymbol{\eta}](\mathbf{x}) = \mathbf{g}(\mathbf{x}), \qquad \mathbf{x}\in S, \tag{2.30}$$

which has a unique solution $\boldsymbol{\eta}$ satisfying (2.28). Furthermore the density function $\boldsymbol{\eta}$ will be orthogonal to $\mathbf{n}$.

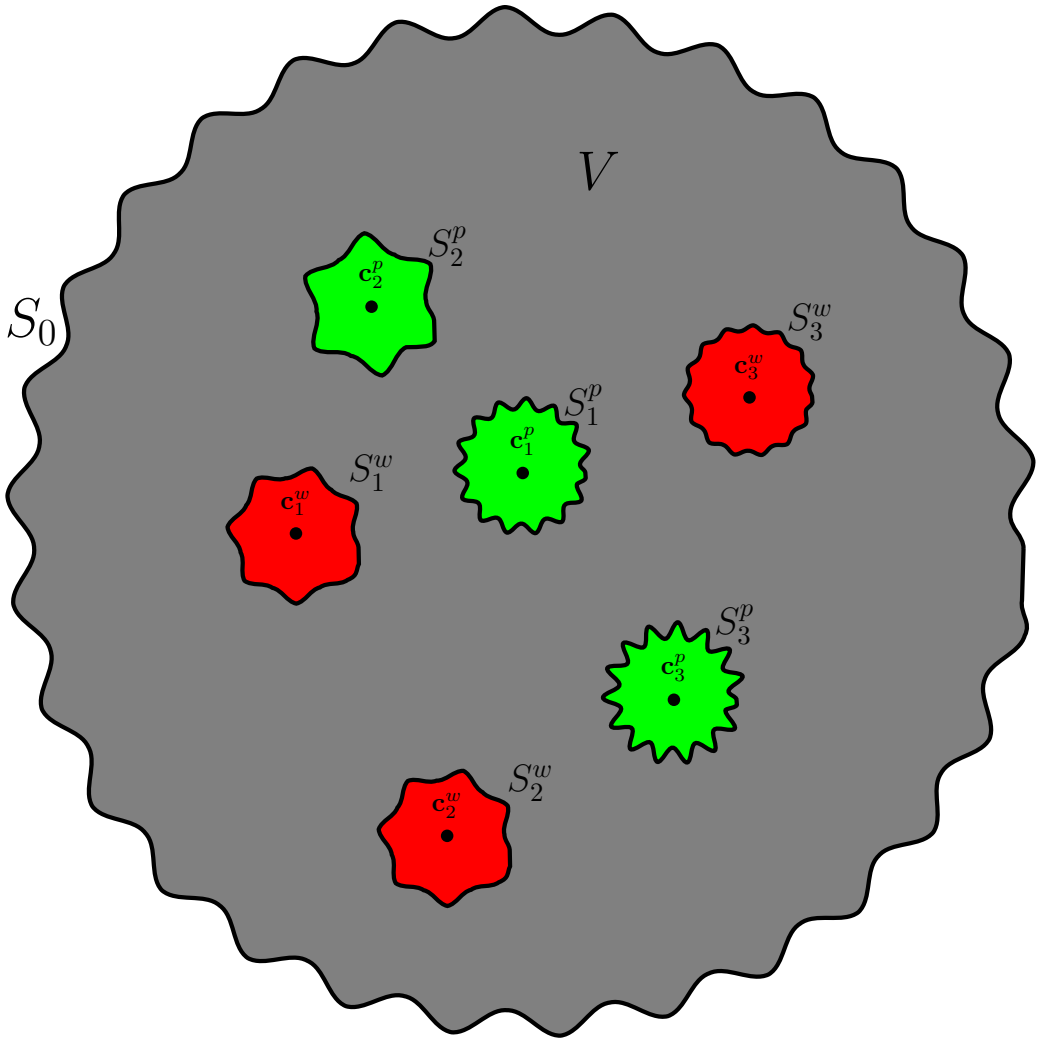

Figure 2.4: Example of a bounded multiply-connected domain. The fluid domain is colored in gray. Obstacles colored in green are mobile bodies and objects colored in red are fixed solid walls.

2.3.4 Boundary Integral Formulation of Rigid Body Suspensions

We now consider a multiply-connected domain consisting of mobile rigid bodies, interior solid walls, and a bounding solid wall. A sketch of a bounded multiply-connected domain is shown in Figure 2.4. Consider a suspension of n_p mobile rigid bodies and n_w interior solid walls. The domain is enclosed by a solid boundary S_0. For each rigid body V_k^p, $k = 1, \ldots, n_p$, we prescribe the net force and torque $\mathbf{F}_k^p$ and L_k^p, respectively. The translational and rotational velocities of the rigid

bodies $\mathbf{U}_k$ and ω_k are unknowns. For each solid wall V_ℓ^w, $\ell = 1, \ldots, n_w$, the velocity at any point on its boundary is given, and we solve for the net force and torque $\mathbf{F}_\ell^w$ and L_ℓ^w. Each rigid body is bounded by a surface S_k^p and has a center of rotation $\mathbf{c}_k^p$. The union of all the rigid body surfaces is denoted by S^p. Each solid wall is denoted by S_ℓ^w and is centered at $\mathbf{c}_\ell^w$. The union of all the wall surfaces is denoted by S^w.

We represent the velocity as a sum of double-layer potentials around each of the components of the fluid boundary. This results in a nontrivial null space that we address by combining the techniques described for exterior (Section 2.3.2) and interior (Section 2.3.3) problems. The double-layer representation has a single null function from the outer wall (which is not orthogonal to the normal on the outer wall) and $3(n_p + n_w)$ null functions from the rigid bodies and solid walls, those being density functions corresponding to rigid body motions [38, 39, 44, 61]. Therefore for $\mathbf{x} \in V$ we make the ansatz

$$\mathbf{u}(\mathbf{x}) = \mathcal{D}[\boldsymbol{\eta}](\mathbf{x}) + \sum_{\ell=1}^{n_w} \left(\mathbf{S}[\mathbf{F}_\ell^w, \mathbf{c}_\ell^w](\mathbf{x}) + \mathbf{R}[L_\ell^w, \mathbf{c}_\ell^w](\mathbf{x}) \right) + \sum_{k=1}^{n_p} \left(\mathbf{S}[\mathbf{F}_k^p, \mathbf{c}_k^p](\mathbf{x}) + \mathbf{R}[L_k^p, \mathbf{c}_k^p](\mathbf{x}) \right), \qquad \mathbf{x} \in V.$$

The density function $\boldsymbol{\eta}$ is defined on the entire surface, i.e. $S_0 \cup S^p \cup S^w$. On the solid walls and the bounding wall, we prescribe a Dirichlet boundary condition, $\mathbf{u}(\mathbf{x}) = \mathbf{g}(\mathbf{x})$ and on the rigid bodies we prescribe the net force and torque $\mathbf{F}_k^p$ and L_k^p, respectively. Then, the canonical equations for the mobility and resistance problem [38, 39] are:

$$\begin{aligned} -\frac{1}{2}\boldsymbol{\eta}(\mathbf{x}) + \mathcal{D}[\boldsymbol{\eta}](\mathbf{x}) + \sum_{\ell=1}^{n_w} \left(\mathbf{S}[\mathbf{F}_\ell^w, \mathbf{c}_\ell^w](\mathbf{x}) + \mathbf{R}[L_\ell^w, \mathbf{c}_\ell^w](\mathbf{x}) \right) + \mathcal{N}_0[\boldsymbol{\eta}](\mathbf{x}) \\ = \mathbf{g}(\mathbf{x}) - \sum_{k=1}^{n_p} \left(\mathbf{S}[\mathbf{F}_k^p, \mathbf{c}_k^p](\mathbf{x}) + \mathbf{R}[L_k^p, \mathbf{c}_k^p](\mathbf{x}) \right), \qquad \mathbf{x} \in S_0, \end{aligned} \tag{2.31a}$$

$$\begin{aligned} -\frac{1}{2}\boldsymbol{\eta}(\mathbf{x}) + \mathcal{D}[\boldsymbol{\eta}](\mathbf{x}) + \sum_{\ell=1}^{n_w} \left(\mathbf{S}[\mathbf{F}_\ell^w, \mathbf{c}_\ell^w](\mathbf{x}) + \mathbf{R}[L_\ell^w, \mathbf{c}_\ell^w](\mathbf{x}) \right) \\ = \mathbf{g}(\mathbf{x}) - \sum_{k=1}^{n_p} \left(\mathbf{S}[\mathbf{F}_k^p, \mathbf{c}_k^p](\mathbf{x}) + \mathbf{R}[L_k^p, \mathbf{c}_k^p](\mathbf{x}) \right), \qquad \mathbf{x} \in S^w, \end{aligned} \tag{2.31b}$$

$$\begin{aligned} -\frac{1}{2}\boldsymbol{\eta}(\mathbf{x}) + \mathcal{D}[\boldsymbol{\eta}](\mathbf{x}) + \sum_{\ell=1}^{n_w} \left(\mathbf{S}[\mathbf{F}_\ell^w, \mathbf{c}_\ell^w](\mathbf{x}) + \mathbf{R}[L_\ell^w, \mathbf{c}_\ell^w](\mathbf{x}) \right) - \mathbf{U}_k - \omega(\mathbf{x} - \mathbf{c}_k^p)^\perp \\ = -\sum_{k=1}^{n_p} \left(\mathbf{S}[\mathbf{F}_k^p, \mathbf{c}_k^p](\mathbf{x}) + \mathbf{R}[L_k^p, \mathbf{c}_k^p](\mathbf{x}) \right), \qquad \mathbf{x} \in S^p. \end{aligned} \tag{2.31c}$$

While the left hand side of (2.31) has a trivial null space, there are still more unknowns than equations. To close the system, we proceed as in the case of a single rigid body by associating the density function with the net force and torque on each obstacle:

$$\int_{S_\ell^w} \boldsymbol{\eta}\, \mathrm{d}S = \mathbf{F}_\ell^w, \qquad \int_{S_\ell^w} \boldsymbol{\eta} \cdot (\mathbf{x} - \mathbf{c}_\ell^w)^\perp \, \mathrm{d}S = L_\ell^w, \qquad \ell = 1, \ldots, N_w, \tag{2.32a}$$

$$\int_{S_k^p} \boldsymbol{\eta}\, \mathrm{d}S = \mathbf{F}_k^p, \qquad \int_{S_k^p} \boldsymbol{\eta} \cdot (\mathbf{x} - \mathbf{c}_k^p)^\perp \, \mathrm{d}S \ = L_k^p, \qquad k = 1, \ldots, N_p. \tag{2.32b}$$

Equations (2.31) and (2.32) form an invertible second-kind Fredholm integral equation. They can be solved for the translational and rotational velocities of the rigid bodies and the net forces and torques on the fixed solid walls. Note that we still require $\int_{S_0 \cup S^w} \mathbf{u} \cdot \mathbf{n}\, \mathrm{d}S = 0$ for global conservation of mass.

Unbounded domains are treated in a similar manner. In this case, for the disturbance velocity to decay at infinity, the total net force on all the obstacles must be zero. For this reason, we only consider unbounded flow past mobile rigid bodies where we prescribe the net force

$$\sum_{k=1}^{n_p} \mathbf{F}_k^p = \mathbf{0}.$$

The canonical equations in this case become

$$\begin{aligned} -\frac{1}{2}\boldsymbol{\eta}(\mathbf{x}) + \mathcal{D}[\boldsymbol{\eta}](\mathbf{x}) - \mathbf{U}_k - \omega(\mathbf{x} - \mathbf{c}_k^p)^\perp \\ = -\mathbf{u}^\infty - \sum_{k=1}^{n_p} \left(\mathbf{S}[\mathbf{F}_k^p](\mathbf{x}, \mathbf{c}_k^p) + \mathbf{R}[L_k^p](\mathbf{x}, \mathbf{c}_k^p) \right), \qquad \mathbf{x} \in S_k^p, \end{aligned} \tag{2.33}$$

along with the closure condition given in (2.32b).

2.3.5 Computing Pressure and Stresses

We are interested in characterizing rheological properties of the suspension. This requires computing the pressure and stress of the flow. Fortunately, these quantities (and others, such as the vorticity) can be computed as a post-processing step after we have calculated the density function, and the forces and torques on the solid walls and rigid bodies.

At a point $\mathbf{x} \in V$, the pressure of the double-layer potential [63] is

$$p^D[\boldsymbol{\eta}](\mathbf{x}) = \frac{1}{\pi} \int_S \frac{1}{r^2} \left(1 - 2\frac{\mathbf{r} \otimes \mathbf{r}}{r^2} \right) \mathbf{n} \cdot \boldsymbol{\eta}\, \mathrm{d}S, \qquad \mathbf{x} \in V. \tag{2.34}$$

At a point $\mathbf{x} \in V$, the strain rate tensor of the double-layer potential [63, 66] is

$$e_{ij}^D[\boldsymbol{\eta}](\mathbf{x}) = \frac{1}{2\pi}\int_S \frac{1}{r^4}\Big(r_k n_k r_\ell \eta_\ell \Big(\delta_{ij} - \frac{8}{r^2} r_i r_j\Big) + r_k \eta_k (n_i r_j + n_j r_i) + r_k n_k (\eta_i r_j + r_j \eta_i)\Big)\,\mathrm{d}S, \qquad \mathbf{x} \in V. \tag{2.35}$$

We also require the pressure and stresses of the Stokeslets and rotlets,

$$p^S(\mathbf{x}) = \sum_{i=1}^{n_w} \frac{\mathbf{r}\cdot\mathbf{F}_i^w}{2\pi r^2} + \sum_{i=1}^{n_p} \frac{\mathbf{r}\cdot\mathbf{F}_i^p}{2\pi r^2} \qquad p^R(\mathbf{x}) = 0, \tag{2.36}$$

as well as the strain rate of the Stokeslets and rotlets,

$$\mathbf{e}^S(\mathbf{x}) = \sum_{i=1}^{n_w} \frac{\mathbf{r}\cdot\mathbf{F}_i^w}{4\pi r^2}\left(\mathbf{I} - \frac{2}{r^2}\mathbf{r}\otimes\mathbf{r}\right) + \sum_{i=1}^{n_p} \frac{\mathbf{r}\cdot\mathbf{F}_i^p}{4\pi r^2}\left(\mathbf{I} - \frac{2}{r^2}\mathbf{r}\otimes\mathbf{r}\right), \qquad \mathbf{x} \in V \cup S, \tag{2.37a}$$

$$\mathbf{e}^R(\mathbf{x}) = \sum_{i=1}^{n_w} \frac{L_i^w}{r^4}\left(\mathbf{r}\otimes\mathbf{r}^\perp + \mathbf{r}^\perp\otimes\mathbf{r}\right) + \sum_{i=1}^{n_p} \frac{L_i^p}{r^4}\left(\mathbf{r}\otimes\mathbf{r}^\perp + \mathbf{r}^\perp\otimes\mathbf{r}\right), \qquad \mathbf{x} \in V \cup S. \tag{2.37b}$$

At a point $\mathbf{x} \in S$, the limiting values of (2.34) and (2.35) must be considered and the resulting surface pressure and strain rate are

$$\mathbf{e}^D[\boldsymbol{\eta}](\mathbf{x})_{(i)} := +\mathcal{J}[\boldsymbol{\eta}](\mathbf{x}) + \mathbf{e}^D[\boldsymbol{\eta}](\mathbf{x}), \tag{2.38a}$$

$$\mathbf{e}^D[\boldsymbol{\eta}](\mathbf{x})_{(e)} := -\mathcal{J}[\boldsymbol{\eta}](\mathbf{x}) + \mathbf{e}^D[\boldsymbol{\eta}](\mathbf{x}), \tag{2.38b}$$

$$p^D[\boldsymbol{\eta}](\mathbf{x})_{(i)} := -\frac{\partial\boldsymbol{\eta}}{\partial\boldsymbol{\tau}}\cdot\boldsymbol{\tau}(\mathbf{x}) + p^D[\boldsymbol{\eta}](\mathbf{x}), \tag{2.38c}$$

$$p^D[\boldsymbol{\eta}](\mathbf{x})_{(i)} := +\frac{\partial\boldsymbol{\eta}}{\partial\boldsymbol{\tau}}\cdot\boldsymbol{\tau}(\mathbf{x}) + p^D[\boldsymbol{\eta}](\mathbf{x}), \tag{2.38d}$$

where

$$\mathcal{J}[\boldsymbol{\eta}](\mathbf{x}) = \frac{1}{2}\left(\frac{\partial\boldsymbol{\eta}}{\partial\boldsymbol{\tau}}\cdot\boldsymbol{\tau}\right)\begin{pmatrix} 2\tau_x\tau_y & \tau_y^2 - \tau_x^2 \\ \tau_y^2 - \tau_x^2 & -2\tau_x\tau_y \end{pmatrix}.$$

Equations (2.34), (2.35), (2.36), (2.37) along with the jump conditions (2.38) allow us to compute the strain rate and the pressure at any point $\mathbf{x} \in V \cup S$. Finally, the pressure and strain rate can be added together to form the total stress.

CHAPTER 3

NUMERICAL METHODS

The canonical equations (2.31) or (2.33), along with the closure formula (2.32) can be solved analytically only for very simple geometries. To simulate general rigid body suspensions we must discretize the equations and efficiently solve the resulting linear system. There are multiple spatial discretizations we could employ. One approach is a Galerkin-based boundary element method [32,39], the integral equation analogue of the finite element method. Here, the geometry's boundary is split into elements and the density function $\boldsymbol{\eta}$ is approximated by a linear combination of basis functions with support over a few elements. An alternative method, that we use, is the Nyström method (Section 3.1), where the value of the density function is approximated at a discrete set of points. The method begins by enforcing the boundary integral equation at a set of N points $\{\mathbf{x}_i\},\ i = 1, \ldots, N$. These points double as quadrature points to discretize the integral operator. This technique is discussed in detail for the Stokes equations in Section 3.2 and we demonstrate that the trapezoid rule is an excellent choice for the quadrature in most cases. However, when bodies are close, the quality of the trapezoid rule deteriorates and specialized quadrature must be applied (Section 3.3).

The resulting linear system is dense and changes at each time step, so a direct linear solver is not appropriate. However, an iterative solver is appropriate, since the behavior of the eigenvalues of the linear system guarantees a bound on the number of required Generalized Minimum Residual Method (GMRES) iterations (Section 3.4). Therefore, the overall cost to solve the $N \times N$ linear system is proportional to the cost of a single matrix-vector product. A dense matrix-vector product ordinarily requires $\mathcal{O}(N^2)$, however, the structure of our specific linear system allows us to take advantage of a fast summation method to reduce the matrix-vector product to an $\mathcal{O}(N)$ operation. Finally, a block-diagonal preconditioner is used to reduce the number of GMRES iterations.

The solution of the governing equations includes the translational velocity $\mathbf{U}$, and the rotational velocity ω of each rigid body. To advance the simulation in time, we must solve the ODEs

$$\frac{d}{dt}\mathbf{c} = \mathbf{U}(t), \qquad \frac{d}{dt}\theta = \omega(t),$$

where $\mathbf{c}$ and θ are the center of the rigid body and its inclination angle, respectively. The time stepping can be done in a variety of ways, and in Section 3.5 we investigate first-order forward Euler and second-order Adams-Bashforth.

3.1 Spatial Discretization

Consider the general boundary integral equation (BIE)

$$\lambda\eta(\mathbf{x}) + \int_S K(\mathbf{x},\mathbf{y})\eta(\mathbf{y})\ \mathrm{d}S(\mathbf{y}) = g(\mathbf{x}), \qquad \mathbf{x} \in S. \tag{3.1}$$

A Nyström method begins by choosing N collocation points $\{\mathbf{x}_i\}_{i=1}^N \subset S$ and enforcing the boundary integral equation at these points

$$\lambda\eta(\mathbf{x}_i) + \int_S K(\mathbf{x}_i,\mathbf{y})\eta(\mathbf{y})\ \mathrm{d}S(\mathbf{y}) = g(\mathbf{x}_i), \qquad i = 1,\ldots,N.$$

Next, by approximating the integral using the quadrature points $\{\mathbf{x}_i\}$ with corresponding weights $\{w_i\}$, the result is the $N \times N$ linear system

$$\lambda\eta(\mathbf{x}_i) + \sum_{j=1}^{N} K(\mathbf{x}_i,\mathbf{x}_j)\eta(\mathbf{x}_j)w_j = g(\mathbf{x}_i), \qquad i = 1,\ldots,N, \tag{3.2}$$

or in matrix notation

$$(\lambda\mathbf{I} + \mathbf{A})\,\hat{\boldsymbol{\eta}} = \hat{\mathbf{g}},$$

where

$$A_{ij} = K(\mathbf{x}_i,\mathbf{x}_j)w_j, \qquad \hat{g}_i = g(\mathbf{x}_i), \qquad \hat{\eta}_i = \eta(\mathbf{x}_i).$$

Note that the quadrature weights $\{w_i\}$ depend on the geometry of S. If the original integral equation is solvable for any g, then the resulting linear system will be full rank.

We must choose the collocation points $\{\mathbf{x}_i\}$ and weights $\{w_i\}$ in (3.2). This choice directly affects the accuracy of the quadrature routine. Let S be parameterized by $\boldsymbol{\phi}(s)$, $s \in [0, 2\pi]$. We use this parameterization to transform the boundary integral over S to one over the interval $[0, 2\pi]$,

$$\int_S f(\mathbf{y})\ \mathrm{d}S = \int_0^{2\pi} f(\boldsymbol{\phi}(s))|\boldsymbol{\phi}'(s)|\ \mathrm{d}s.$$

Since S is a closed boundary, $\boldsymbol{\phi}(s)$ is periodic over $[0, 2\pi]$. In addition, we assume that S is smooth, so $\boldsymbol{\phi} \in \mathbf{C}^\infty([0, 2\pi])$. The collocation points are chosen to be equally spaced in parameter space s, that is $\mathbf{x}_j = \boldsymbol{\phi}(s_j)$, where

$$s_j = (j-1)h, \qquad h = \frac{2\pi}{N}.$$

A particularly attractive choice for quadrature in this case is the trapezoid rule. That is

$$\int_S f(\mathbf{y})\ \mathrm{d}S \approx h\sum_{j=1}^{N} f(\boldsymbol{\phi}(s_j))|\boldsymbol{\phi}'(s_j)|.$$

The trapezoid rule is used because it is spectrally accurate when applied to smooth, periodic functions [79]. As a demonstration of the spectral accuracy for periodic functions, consider numerically integrating two functions, $f_1(x) = e^x$ and $f_2(x) = e^{\cos(x)}$ from 0 to 2π. The first integral is $e^{2\pi} - 1$, and the second is $2\pi I_1(0)$, where $I_1(x)$ is the first-order first-kind modified Bessel function. As shown in Table 3.1, by applying the trapezoid rule, we achieve second-order accuracy when evaluating the non-periodic integral, and spectral accuracy when evaluating the periodic integral.

Table 3.1: Convergence study using the trapezoid rule for a non-periodic and a periodic function over the interval $x \in [0, 2\pi]$. E_1 is the relative error of the N point trapezoid rule applied to $f(x) = e^x$, and E_2 is the relative error of the N point trapezoid rule applied to $f(x) = e^{\cos(x)}$. In the first case, the trapezoid rule is second-order accurate, while in the second case it is spectrally accurate.

N	E_1	E_2
4	1.98×10^{-1}	4.32×10^{-3}
8	5.09×10^{-2}	1.57×10^{-7}
16	1.28×10^{-2}	2.23×10^{-16}
32	3.21×10^{-3}	3.25×10^{-16}
64	8.03×10^{-4}	0

As a final step to applying quadrature, we require the arclength term $\boldsymbol{\phi}'(s_j)$. This is computed by using the Fourier representation

$$\boldsymbol{\phi}(s_j) = \sum_{k=1}^{N} \hat{\boldsymbol{\phi}}_k e^{i(k-1)s_k}, \qquad j = 1, \ldots, N.$$

Then

$$\boldsymbol{\phi}'(s_j) = \sum_{k=1}^{N} i(k-1)\hat{\boldsymbol{\phi}}_k e^{i(k-1)s_k}, \qquad j = 1, \ldots, N.$$

Computing the Fourier coefficients $\hat{\boldsymbol{\phi}}_k$ requires a discrete Fourier transform, and computing $\boldsymbol{\phi}'(s_j)$ requires an inverse discrete Fourier transform. Both of these operations can be done naïvely in $\mathcal{O}(N^2)$ operations. However, we accelerate each calculation to $\mathcal{O}(N \log N)$ operations by using the fast Fourier transform (FFT).

3.2 Discretization of Stokes Boundary Integral Equation

Discretizing the boundary integral equation (3.1) depends on the kernel K and the constant λ. Here, we focus on the Stokes double-layer potential boundary integral equation, where the integral kernel is given by (2.10b) and $\lambda = -1/2$. The kernel, density function, and boundary conditions are all vector-valued, but this does not change the discretization approach outlined in Section 3.1.

3.2.1 Interior Flow

Considering the Stokes equations in a bounded simply-connected domain, we have to solve (2.30) for the density function $\boldsymbol{\eta}$. Denoting $\boldsymbol{\eta}_i$ to be $\boldsymbol{\eta}(\boldsymbol{\phi}(s_i))$, the N-point trapezoid results in the linear system

$$-\frac{1}{2}\boldsymbol{\eta}_i + h\sum_{j=1}^{N}\mathbf{K}(\boldsymbol{\phi}(s_i),\boldsymbol{\phi}(s_j))\boldsymbol{\eta}_j|\boldsymbol{\phi}'(s_j)| + h\mathbf{n}(\boldsymbol{\phi}(s_i))\sum_{j=1}^{N}\mathbf{n}(\boldsymbol{\phi}(s_j))\boldsymbol{\eta}_j = \mathbf{g}(\boldsymbol{\phi}(s_i)), \qquad i = 1,\ldots,N,$$

where $h = 2\pi/N$ and $\mathbf{g}(\mathbf{x})$ is the prescribed Dirichlet boundary condition on the velocity.

At first glance the kernel $\mathbf{K}$ is singular when $i = j$, and it appears that specialized quadrature is required. However, the singularity is removable and its limiting value can be used instead of $\mathbf{K}(\boldsymbol{\phi}_i, \boldsymbol{\phi}_i)$. In particular, in $\mathbb{R}^2$ the kernel of double-layer potential has the limiting value

$$\lim_{\substack{\mathbf{y}\to\mathbf{x}\\ \mathbf{y}\in S}} K_{ij}(\mathbf{x},\mathbf{y}) = \frac{\kappa(\mathbf{x})\tau_i(\mathbf{x})\tau_j(\mathbf{x})}{2\pi}, \qquad \mathbf{x}\in S, \tag{3.3}$$

where $\kappa(\mathbf{x})$ is the curvature of S at $\mathbf{x}$ and $\boldsymbol{\tau}(\mathbf{x})$ is the tangent vector of S at $\mathbf{x}$. Replacing the diagonal terms in the discretization with the limiting value (3.3), the linear system we must solve is

$$\begin{aligned}\mathbf{g}(\boldsymbol{\phi}(s_i)) =& \left(-\frac{1}{2} + h\left(\frac{\kappa(\boldsymbol{\phi}(s_i))\boldsymbol{\tau}(s_i)\otimes\boldsymbol{\tau}(s_i)}{2\pi} + \mathbf{n}(\boldsymbol{\phi}(s_i))\otimes\mathbf{n}(\boldsymbol{\phi}(s_i))\right)\right)\boldsymbol{\eta}_i|\boldsymbol{\phi}'(s_i)| \\ &+ h\sum_{\substack{j=1\\ j\neq i}}^{N}\left(\mathbf{K}(\boldsymbol{\phi}(s_i),\boldsymbol{\phi}(s_j)) + \mathbf{n}(\boldsymbol{\phi}(s_i))\otimes\mathbf{n}(\boldsymbol{\phi}(s_j))\right)\boldsymbol{\eta}_j|\boldsymbol{\phi}'(s_j)|, \qquad i = 1,\ldots,N.\end{aligned} \tag{3.4}$$

This is a full rank dense linear system of size $2N\times 2N$, and its solution approximates the values of $\boldsymbol{\eta}$ at the collocation points on S with spectral accuracy. Techniques for efficiently solving this system are discussed in Section 3.4.

We have computed the density function values $\{\boldsymbol{\eta}(\mathbf{x}_i)\}_{i=1}^{N}$, but these values have no physical meaning. We are actually interested in computing the velocity $\mathbf{u}(\mathbf{x})$ for $\mathbf{x}\in V$. This can be

approximated with the trapezoid rule

$$\mathbf{u}(\mathbf{x}) = h \sum_{i=1}^{N} \mathbf{K}(\mathbf{x}, \boldsymbol{\phi}(s_i)) \boldsymbol{\eta}_i |\phi'(s_i)|, \qquad \mathbf{x} \in V. \tag{3.5}$$

For a fixed $\mathbf{x} \in V$, $\mathbf{K}(\mathbf{x}, \mathbf{y})$ is a smooth, periodic function. Therefore we expect the trapezoid rule to approximate $\mathbf{u}(\mathbf{x})$ with spectral accuracy, which we now demonstrate.

Consider the Stokes equations inside an ellipse with semi-major axis 2 and semi-minor axis 1 (Figure 3.1). Inside this geometry V we consider the Stokes equations with boundary conditions $\mathbf{g}(\mathbf{x}) = (x, -y),\ \mathbf{x} \in S$. This problem has the exact solution $\mathbf{u}(\mathbf{x}) = (x, -y),\ \mathbf{x} \in V$. We perform a convergence study by solving the linear system (3.4) at various levels of refinement and then computing $\mathbf{u}(\mathbf{x})$ using (3.5) at a set of target points $\mathbf{x} \in V$. The results in Table 3.2 show spectral accuracy, but some target points result in larger error constants. In particular, at a fixed resolution N, the error grows as $\mathbf{x}$ approaches S. This increase in the error occurs because the derivative of the double-layer kernel grows as the target point approaches the boundary. Because of this, the trapezoid rule requires an unfeasible number of points to accurately integrate the kernel. A near singular integration technique is needed in this case and is discussed in Section 3.3.

Table 3.2: We solve the linear system (3.4) with boundary conditions $\mathbf{g}(\mathbf{x}) = (x, -y)$ on the geometry shown in Figure 3.1. When evaluating the velocity using the computed double-layer density according to (3.5), the error between the computed solution and the exact solution is very small at $\mathbf{x} = (0, 0)$ with as few as 8 quadrature points. For $\mathbf{x} = (0, 1)$ and $\mathbf{x} = (0, 1.9)$ we see spectral accuracy, but the convergence is delayed for $\mathbf{x} = (0, 1.9)$. At $\mathbf{x} = (0, 1.99)$ even 512 points gives only a single digit of accuracy.

N	error at (0,0)	error at (0,1)	error at (0,1.9)	error at (0,1.99)
8	8.75×10^{-16}	2.80×10^{-1}	7.41×10^{0}	8.89×10^{1}
16	3.82×10^{-16}	8.45×10^{-3}	2.75×10^{0}	4.32×10^{1}
32	9.93×10^{-16}	1.01×10^{-6}	6.04×10^{-1}	2.07×10^{1}
64	1.95×10^{-16}	1.25×10^{-13}	2.90×10^{-2}	9.45×10^{0}
128	4.12×10^{-16}	3.90×10^{-16}	4.17×10^{-5}	3.84×10^{0}
256	4.23×10^{-16}	1.11×10^{-15}	4.59×10^{-11}	1.13×10^{0}
512	4.66×10^{-16}	2.41×10^{-16}	4.37×10^{-17}	1.30×10^{-1}

3.2.2 Exterior Flow

For an unbounded geometry, the discretization of the boundary integral equation is similar to the bounded case. We no longer require the $\mathcal{N}_0$ term, however, to complete the double-layer potential (Section 2.3.2), we must add on a Stokeslet and a rotlet with strength $\mathbf{F}$ and L, respectively. Then, these strengths must be related to the density function by (2.27). The discretized Stokes equations (2.26) in this case are

$$\begin{aligned}\Big(-\frac{1}{2}+h\Big(\frac{\kappa(\boldsymbol{\phi}(s_i))\boldsymbol{\tau}(s_i)\otimes\boldsymbol{\tau}(s_i)}{2\pi}\Big)\Big)\boldsymbol{\eta}_i|\boldsymbol{\phi}'(s_i)| + h\sum_{\substack{j=1\\ j\neq i}}^{N}\left(\mathbf{K}(\boldsymbol{\phi}(s_i),\boldsymbol{\phi}(s_j))\boldsymbol{\eta}_j|\boldsymbol{\phi}'(s_j)|\right)\\ +\mathbf{U}+\omega(\boldsymbol{\phi}(s_i)-\mathbf{c})^{\perp}+\mathbf{S}[\mathbf{F},\mathbf{c}](\boldsymbol{\phi}(s_i))+\mathbf{R}[L,\mathbf{c}](\boldsymbol{\phi}(s_i))=-\mathbf{u}^{\infty},\qquad i=1,\dots,N.\end{aligned} \tag{3.6}$$

The closure formulas (2.27) involve a periodic integral, so we discretize them also with the trapezoid rule

$$h\sum_{k=1}^{N}\boldsymbol{\eta}_k|\boldsymbol{\phi}'(s_k)|=\mathbf{F},\qquad h\sum_{k=1}^{N}\boldsymbol{\eta}\cdot(\boldsymbol{\phi}(s_k)-\mathbf{c})^{\perp}|\boldsymbol{\phi}'(s_k)|=L. \tag{3.7}$$

For a mobile rigid body, $\mathbf{F}$ and L are known and so are moved to the right hand side of (3.6). The result is the full rank linear system

$$\begin{pmatrix}-\frac{1}{2}\mathbf{I}+\mathbf{D} & \mathbf{B}\\ \mathbf{B}^T & \mathbf{0}\end{pmatrix}\begin{pmatrix}\boldsymbol{\eta}\\ \hat{\mathbf{U}}\end{pmatrix}=\begin{pmatrix}-\mathbf{u}^{\infty}\\ \hat{\mathbf{F}}\end{pmatrix},$$

where $\mathbf{D}\in\mathbb{R}^{2N\times 2N}$ represents the contributions of the double-layer potential to the velocity, $\hat{\mathbf{U}}$ is shorthand for the vector $(\mathbf{U},\omega)$, $\hat{\mathbf{F}}$ is shorthand for the vector $(\mathbf{F},L)$, and $\mathbf{B}^T\in\mathbb{R}^{3\times 2N}$ is the discrete closure relation (3.7). This leaves a $\mathbb{R}^{(2N+3)\times(2N+3)}$ dense linear system to solve.

Once we have computed $\{\boldsymbol{\eta}(\mathbf{x}_i)\}_{i=1}^{N}$, we can again use it along with $\mathbf{F}$ and L to compute the velocity with spectral accuracy at any point $\mathbf{x}\in V$ by using the trapezoid rule

$$\mathbf{u}(\mathbf{x})=\sum_{i=1}^{N}\mathbf{K}(\mathbf{x},\boldsymbol{\phi}(s_i))\boldsymbol{\eta}_i|\boldsymbol{\phi}'(s_i)|+\mathbf{S}[\mathbf{F},\mathbf{c}](\mathbf{x})+\mathbf{R}[L,\mathbf{c}](\mathbf{x}),\qquad \mathbf{x}\in V.$$

As discussed in Section 2.3.2, for a bounded solution to exist at infinity, $\mathbf{F}=\mathbf{0}$ and $L=0$. We have included them in this formulation for clarity. In the following section, we will discuss simulations involving more than one rigid body, and in those cases $\mathbf{F}$ and L need not be zero.

3.2.3 Multiply-Connected Geometry

For multiply-connected domains (Figure 2.4), we must consider not only the contributions between points on the same body (intra-body effects), but also contributions between points on

different bodies (inter-body effects). Consider a collection of mobile rigid bodies inside a bounded or an unbounded multiply-connected domain. Denote the number of rigid bodies as n_p and the number of interior fixed walls as n_w. Let rigid body k be parameterized by $\boldsymbol{\phi}^k(s)$ and wall k by $\boldsymbol{\varphi}^k(s)$, $s \in [0, 2\pi)$.

Denote the density function on rigid body k as $\boldsymbol{\eta}^k$ and on solid wall k as $\boldsymbol{\zeta}^k$, and consider evaluating the velocity at a point $\mathbf{x}$ on rigid body q, i.e. $\mathbf{x} \in S_q^p$. There are four terms that contribute to this:

1. The contribution to the double-layer potential due to rigid body q:

$$-\frac{1}{2}\boldsymbol{\eta}^q(\mathbf{x}) + \mathcal{D}[\boldsymbol{\eta}^q](\mathbf{x}),$$

2. The contribution from all other rigid bodies:

$$\sum_{\substack{k=1 \\ k\neq q}}^{n_p} \mathcal{D}[\boldsymbol{\eta}^k](\mathbf{x}),$$

3. The contribution from all solid walls (including the bounding wall S_0^w if it exists):

$$\sum_{k=0}^{n_w} \mathcal{D}[\boldsymbol{\zeta}^k](\mathbf{x}),$$

4. The contribution from Stokeslets and rotlets:

$$\sum_{k=1}^{n_p} \left(\mathbf{S}[\mathbf{F}_k^p, \mathbf{c}](\mathbf{x}) + \mathbf{R}[L_k^p, \mathbf{c}](\mathbf{x})\right) + \sum_{k=1}^{n_w} \left(\mathbf{S}[\mathbf{F}_k^w, \mathbf{c}](\mathbf{x}) + \mathbf{R}[L_k^w, \mathbf{c}](\mathbf{x})\right).$$

The velocity at a point $\mathbf{x} \in S_q^p$ is thus

$$\begin{aligned}\mathbf{u}(\mathbf{x}) = &-\frac{1}{2}\boldsymbol{\eta}^q(\mathbf{x}) + \sum_{k=1}^{n_p} \mathcal{D}[\boldsymbol{\eta}^k](\mathbf{x}) + \sum_{k=0}^{n_w} \mathcal{D}[\boldsymbol{\zeta}^k](\mathbf{x}) \\ &+ \sum_{k=1}^{n_p} \left(\mathbf{S}[\mathbf{F}_k^p, \mathbf{c}](\mathbf{x}) + \mathbf{R}[L_k^p, \mathbf{c}](\mathbf{x})\right) + \sum_{k=1}^{n_w} \left(\mathbf{S}[\mathbf{F}_k^w, \mathbf{c}](\mathbf{x}) + \mathbf{R}[L_k^w, \mathbf{c}](\mathbf{x})\right), \qquad \mathbf{x} \in S_q^p.\end{aligned}$$

Similarly, the velocity at a point $\mathbf{x} \in S_m^w$ is

$$\begin{aligned}\mathbf{u}(\mathbf{x}) = &-\frac{1}{2}\boldsymbol{\zeta}^m(\mathbf{x}) + \sum_{k=1}^{n_p} \mathcal{D}[\boldsymbol{\eta}^k](\mathbf{x}) + \sum_{k=0}^{n_w} \mathcal{D}[\boldsymbol{\zeta}^k](\mathbf{x}) \\ &+ \sum_{k=1}^{n_p} \left(\mathbf{S}[\mathbf{F}_k^p, \mathbf{c}](\mathbf{x}) + \mathbf{R}[L_k^p, \mathbf{c}](\mathbf{x})\right) + \sum_{k=1}^{n_w} \left(\mathbf{S}[\mathbf{F}_k^w, \mathbf{c}](\mathbf{x}) + \mathbf{R}[L_k^w, \mathbf{c}](\mathbf{x})\right), \qquad \mathbf{x} \in S_m^w.\end{aligned}$$

We now disctetize the canonical equations (2.31) using the trapezoid rule. The discretization depends on the location of the target point $\mathbf{x}$ and we consider the three cases: 1) $\mathbf{x} \in S_0^w$, 2) $\mathbf{x} \in S_k^w,\ k \neq 0$, and 3) $\mathbf{x} \in S_k^p$. Wall k is discretized as in Section 3.2.1, using N_w points equally spaced in the parameter s, $\{\boldsymbol{\varphi}^k(s_i)\}_{i=1}^{N_w}$ (with spacing $h_w = 2\pi/N_w$), and rigid body k is discretized with N_p points equally spaced in the parameter q (with spacing $h_p = 2\pi/N_p$). The discrete version of (2.31a) is

$$
\begin{aligned}
&\Big(-\frac{1}{2} + h_w \Big(\frac{\kappa^0(\boldsymbol{\varphi}^0(s_\ell))\boldsymbol{\tau}^0(\boldsymbol{\varphi}^0(s_\ell)) \otimes \boldsymbol{\tau}^0(\boldsymbol{\varphi}^0(s_\ell))}{2\pi} + \mathbf{n}(\boldsymbol{\phi}(s_\ell)) \otimes \mathbf{n}(\boldsymbol{\phi}(s_\ell)) \Big) \Big) \boldsymbol{\zeta}_\ell^0 |(\boldsymbol{\varphi}^0)'(s_\ell)| \\
&\quad + h_w \sum_{\substack{k=1 \\ k\neq \ell}}^{N_w} \Big(\mathbf{K}(\boldsymbol{\varphi}^0(s_\ell), \boldsymbol{\varphi}^0(s_k)) + \mathbf{n}(\boldsymbol{\varphi}^0(s_\ell)) \otimes \mathbf{n}(\boldsymbol{\varphi}^0(s_k)) \Big) \boldsymbol{\zeta}_k^0 |(\boldsymbol{\varphi}^0)'(s_k)| \\
&\quad + \sum_{m=1}^{n_w} \mathcal{D}[\boldsymbol{\zeta}^m](\boldsymbol{\varphi}^0(s_\ell)) \\
&\quad + \sum_{r=1}^{n_p} \mathcal{D}[\boldsymbol{\eta}^r](\boldsymbol{\varphi}^0(s_\ell)) \\
&\quad + \sum_{m=1}^{n_w} \Big(\mathbf{S}[\mathbf{F}_m^w, \mathbf{c}_m^w](\boldsymbol{\varphi}^0(s_\ell)) + \mathbf{R}[L_m^w, \mathbf{c}_m^w](\boldsymbol{\varphi}^0(s_\ell)) \Big) \\
&\quad = \mathbf{g}(\boldsymbol{\varphi}^0(s_\ell)) - \sum_{r=1}^{n_p} \Big(\mathbf{S}[\mathbf{F}_r^p, \mathbf{c}_r^p](\boldsymbol{\varphi}^0(s_\ell)) + \mathbf{R}[L_r^p, \mathbf{c}_r^p](\boldsymbol{\varphi}^0(s_\ell)) \Big), \qquad \ell = 1, \ldots, N_w.
\end{aligned}
\tag{3.8a}
$$

Intra-body contributions

Contributions from other walls

Contributions from rigid bodies

Rotlet and Stokeslet contributions

Since we are prescribing the force and torque on the rigid bodies, the contributions from the rigid body Stokeslets and rotlets are a known quantity and can be moved to the right hand side. The net force and torque on the interior walls is unknown however, so the contribution from the wall Stokeslets and rotlets remain on the left hand side.

For interior fixed walls, the discretization is almost identical, except that we no longer need the $\mathcal{N}_0$ term (see (2.31b)). The discretization on wall j is

$$\begin{aligned}
&\Big(-\frac{1}{2} + h_w \Big(\frac{\kappa^j(\boldsymbol{\varphi}^j(s_\ell))\boldsymbol{\tau}^j(\boldsymbol{\varphi}^j(s_\ell)) \otimes \boldsymbol{\tau}^j(\boldsymbol{\varphi}^j(s_\ell))}{2\pi} \Big) \Big) \boldsymbol{\zeta}^j_\ell |(\boldsymbol{\varphi}^j)'(s_\ell)| \\
&\quad + h_w \sum_{\substack{k=1 \\ k\neq \ell}}^{N_w} \Big(\mathbf{K}(\boldsymbol{\varphi}^j(s_\ell), \boldsymbol{\varphi}^j(s_k))s \Big) \boldsymbol{\zeta}^j_k |(\boldsymbol{\varphi}^j)'(s_k)| \quad \leftarrow \text{Intra-body contributions} \\
&\quad + \sum_{\substack{m=0 \\ m\neq j}}^{n_w} \mathcal{D}[\boldsymbol{\zeta}^m](\boldsymbol{\varphi}^j(s_\ell)) \quad \leftarrow \text{Contributions from other walls} \\
&\quad + \sum_{r=1}^{n_p} \mathcal{D}[\boldsymbol{\eta}^r](\boldsymbol{\varphi}^j(s_\ell)) \quad \leftarrow \text{Contributions from rigid bodies} \\
&\quad + \sum_{m=1}^{n_w} \Big(\mathbf{S}[\mathbf{F}^w_m, \mathbf{c}^w_m](\boldsymbol{\varphi}^j(s_\ell)) + \mathbf{R}[L^w_m, \mathbf{c}^w_m](\boldsymbol{\varphi}^j(s_\ell)) \Big) \quad \leftarrow \text{Rotlet and Stokeslet contributions} \\
&\quad = \mathbf{g}(\boldsymbol{\varphi}^j(s_\ell)) - \sum_{r=1}^{n_p} \Big(\mathbf{S}[\mathbf{F}^p_r, \mathbf{c}^p_r](\boldsymbol{\varphi}^j(s_\ell)) + \mathbf{R}[L^p_r, \mathbf{c}^p_r](\boldsymbol{\varphi}^j(s_\ell)) \Big), \qquad \ell = 1, \ldots, N_w.
\end{aligned} \tag{3.8b}$$

.

For rigid body j the translational velocity $\mathbf{U}_j$ and rotational velocity ω_j are unknowns. The velocity on the boundary S^p_j is $\mathbf{U}_j + \omega_j(\mathbf{x} - \mathbf{c}^p_j)^\perp$. This replaces the given boundary condition $\mathbf{g}(\mathbf{x})$ and put on the left hand side (see (2.31c)). The resulting discretization is

$$\begin{aligned}
&\Big(-\frac{1}{2} + h_p \Big(\frac{\kappa^j(\boldsymbol{\phi}^j(q_\ell))\boldsymbol{\tau}^j(\boldsymbol{\phi}^j(q_\ell)) \otimes \boldsymbol{\tau}^j(\boldsymbol{\phi}^j(s_\ell))}{2\pi} \Big) \Big) \boldsymbol{\eta}^j_\ell |(\boldsymbol{\phi}^j)'(q_\ell)| \\
&\quad + h_p \sum_{\substack{k=1 \\ k\neq \ell}}^{N_p} \Big(\mathbf{K}(\boldsymbol{\phi}^j(q_\ell), \boldsymbol{\phi}^j(q_k))s \Big) \boldsymbol{\eta}^j_k |(\boldsymbol{\phi}^j)'(q_k)| \quad \leftarrow \text{Intra-body contributions} \\
&\quad + \sum_{m=0}^{n_w} \mathcal{D}[\boldsymbol{\zeta}^m](\boldsymbol{\phi}^j(q_\ell)) \quad \leftarrow \text{Contributions from walls} \\
&\quad + \sum_{\substack{r=1 \\ r\neq j}}^{n_p} \mathcal{D}[\boldsymbol{\eta}^r](\boldsymbol{\phi}^j(q_\ell)) \quad \leftarrow \text{Contributions from other rigid bodies} \\
&\quad + \sum_{m=1}^{n_w} \Big(\mathbf{S}[\mathbf{F}^w_m, \mathbf{c}^w_m](\boldsymbol{\phi}^j(q_\ell)) + \mathbf{R}[L^w_m, \mathbf{c}^w_m](\boldsymbol{\phi}^j(q_\ell)) \Big) \quad \leftarrow \text{Rotlet and Stokeslet contributions} \\
&\quad - \mathbf{U}_j - \omega_j(\boldsymbol{\phi}^j(q^\ell) - \mathbf{c}^p_j)^\perp \\
&\quad = -\sum_{r=1}^{n_p} \Big(\mathbf{S}[\mathbf{F}^p_r, \mathbf{c}^p_r](\boldsymbol{\phi}^j(q_\ell)) + \mathbf{R}[L^p_r, \mathbf{c}^p_r](\boldsymbol{\phi}^j(q_\ell)) \Big), \qquad \ell = 1, \ldots, N_w.
\end{aligned} \tag{3.8c}$$

Note that we have not yet discretized the inter-body contributions. As we saw in Section 3.2.1 (and in particular Table 3.2), the error of the trapezoid rule is unacceptable if $\mathbf{x}$ is close to the

boundary of the domain. For multiply-connected domains, this is also problematic if two disconnected components of the geometry's boundary are sufficiently close. If the target point $\mathbf{x}$ is far from rigid body ℓ (or wall ℓ) then the double-layer potential due to that rigid body or wall can be accurately evaluated using the standard trapezoid rule,

$$\mathcal{D}[\boldsymbol{\zeta}^\ell](\mathbf{x}) \approx h_w \sum_{k=1}^{N_w} \mathbf{K}(\mathbf{x}, \boldsymbol{\varphi}^\ell(s_k))\boldsymbol{\zeta}_k^\ell |(\boldsymbol{\varphi}^\ell)'(s_k)|,$$
$$\mathcal{D}[\boldsymbol{\eta}^\ell](\mathbf{x}) \approx h_p \sum_{k=1}^{N_p} \mathbf{K}(\mathbf{x}, \boldsymbol{\phi}^\ell(q_k))\boldsymbol{\eta}_k^\ell |(\boldsymbol{\phi}^\ell)'(q_k)|.$$

If, however, the target point is close to another body, then a different integration technique is required. This is discussed in the following section.

3.3 Near Singular Integration

We have observed that the trapezoid rule does a poor job of approximating the double-layer potential when the target point is close to the boundary, which includes both solid walls and rigid bodies. This loss of accuracy is caused by a large derivative in the kernel of the double-layer potential when the distance between the target point and the boundary is small. In dense suspensions this situation is inevitable, and we must modify the quadrature to maintain stability. There are many near singular integration methods, including quadrature by expansion [3, 72], barycentric formulae [6, 31], or asymptotic based methods [55].

In Figure 3.1, we revisit the example in Section 3.2.1 and consider evaluating the velocity at $\mathbf{x} = (0, x_2)$. For a fixed x_2, the double-layer kernel is smooth. As seen in the right plots however, as the target point approaches S, the kernel's derivative becomes larger and larger. However, if $\mathbf{x} \in S$, the singularity is removable, the kernel the takes limiting value (3.3), and its derivative is no longer large. This is observed in the bottom right plot of Figure 3.1.

We take advantage of this by adopting the near singular integration technique described in [63, 84]. When compared to other near singular integration schemes, this scheme is simpler to implement, and is relatively non-intrusive to the code base, yet still delivers high-order accuracy. Assume that the boundary S is discretized with N evenly spaced points, with h being the maximum arclength spacing, and let $d(\mathbf{x}, S) = \inf_{\mathbf{y}\in S} ||\mathbf{x} - \mathbf{y}||$ be the distance between a target point $\mathbf{x}$ and S. We partition V into two regions: the far zone of S, $V_1 = \{\mathbf{x} \,|\, d(\mathbf{x}, S) \geq h\}$, and the near zone of S, $V_0 = \{\mathbf{x} \,|\, d(\mathbf{x}, S) < h\}$.

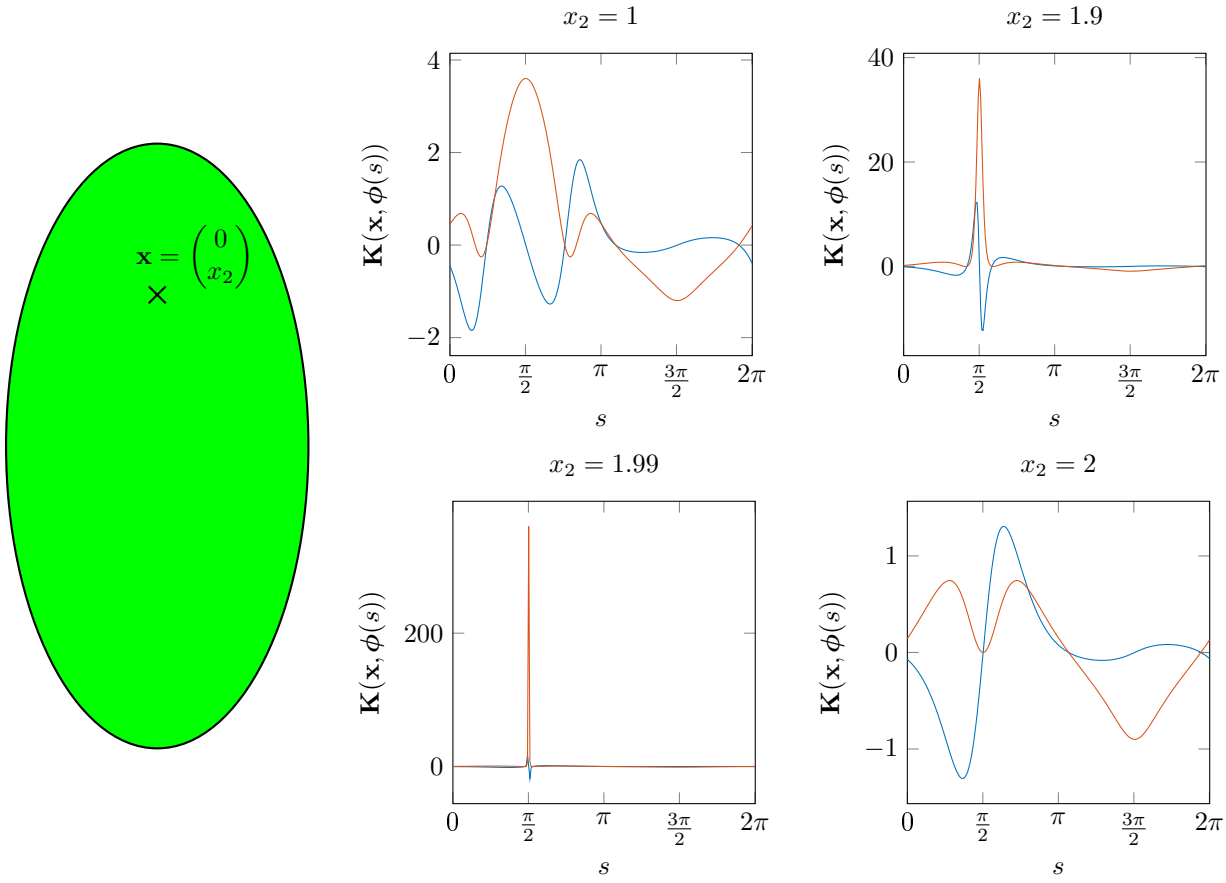

Figure 3.1: A scenario where near singular integration is needed. The point $\mathbf{x}$ is inside an ellipse with semi-major axis 2 and semi-minor axis 1. The two components of the kernel of the double-layer potential as a function of the parameterization variable s are plotted on the right for four values of x_2. Despite the large derivatives, the integrand is always C^∞. If x_2 is far from the boundary, then the derivatives are small and the kernel can be integrated accurately using the trapezoid rule. As x_2 approaches 2, the kernel becomes sharply peaked and requires more and more quadrature points for the trapezoid rule to maintain accuracy. However, when $x_2 = 2$, the target point is on the boundary and once again the trapezoid rule accurately approximates the integral.

For target points $\mathbf{x} \in V_1$, the kernel is sufficiently smooth for the trapezoid rule to be appropriate. However, if $\mathbf{x} \in V_0$ the error of the trapezoid rule is too large. Instead of applying the trapezoid rule, we first find the closest point on S to $\mathbf{x}$, i.e. the point $\mathbf{x}_0 \in S$ that minimizes $||\mathbf{x} - \mathbf{x}_0||^2$, using Newton's method. We then define the m interpolation points

$$\mathbf{x}_j = \mathbf{x}_0 + j\beta h \frac{\mathbf{x} - \mathbf{x}_0}{||\mathbf{x} - \mathbf{x}_0||} \qquad j = 1, \cdots m,$$

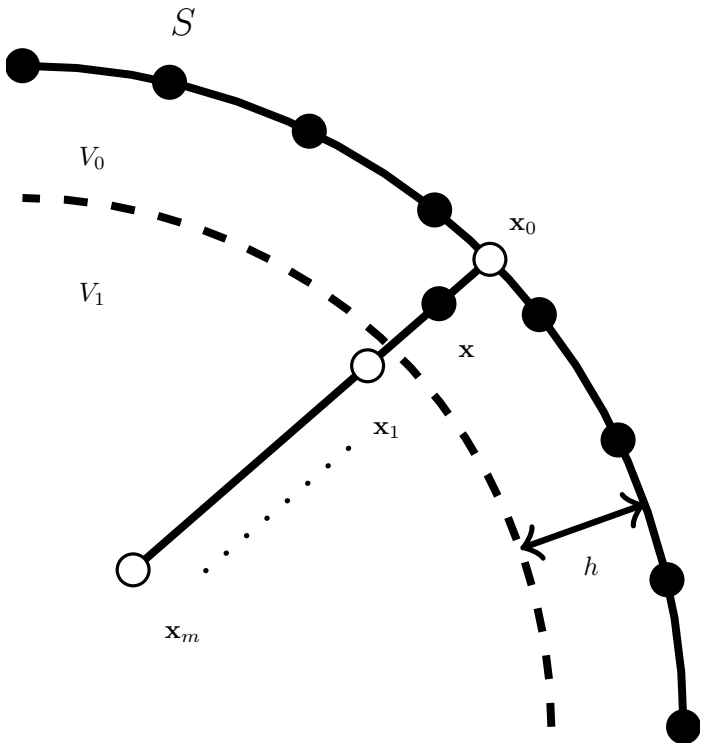

Figure 3.2: Schematic of the near singular integration technique used to evaluate the double-layer potential at a point $\mathbf{x} \in V_0$. By construction, the points $\{\mathbf{x}_i\}$, $i = 1, \ldots m$, are in V_1, where the layer potential can be accurately evaluated with the trapezoid rule. At $\mathbf{x}_0 \in S$ the double-layer potential is approximated with a local interpolant of the the DLP using collocation points on S. Then, a one-dimensional interpolation can be used to approximate $\mathcal{D}[\boldsymbol{\eta}](\mathbf{x})$.

where β is a constant slightly greater than one, and it guarantees that all interpolation points are in V_1. These points are shown in Figure 3.2. Next, we evaluate the double-layer potential at $\mathbf{x}_j$, $j = 0, \ldots, m$. Since $\{\mathbf{x}_i\}_{i=1}^m \in V_1$, the layer potential at these points can be accurately evaluated with the trapezoid rule. A local Lagrange interpolant of the smooth function $\mathcal{D}[\boldsymbol{\eta}](\mathbf{x})$, $\mathbf{x} \in S$ is used to evaluate the layer potential at $\mathbf{x}_0$. Finally, the value of the layer potential at the target point $\mathbf{x}$ is evaluated using a one-dimensional Lagrange interpolant. The accuracy of the method depends on m and N. Convergence rates and efficiently estimates are given in [63, 84].

We apply this scheme with $m = 5$ to the problem in Figure 3.1. As seen in Figure 3.3, the error is much smaller when compared to the trapezoid rule.

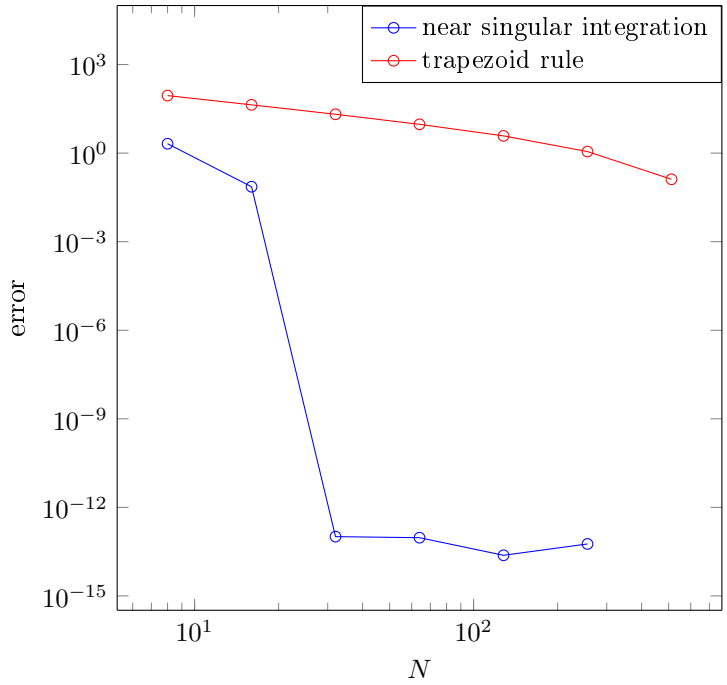

Figure 3.3: Convergence study for the near singular integration technique with $m = 5$ applied to the problem in Figure 3.1 with $x_2 = 1.99$. The near singular interpolation significantly outperforms the trapezoid quadrature for points in V_0.

3.4 Solving the Linear System

Consider a suspension of n_p rigid bodies in an unbounded domain. In symbolic notation, we can write the linear system associated with (2.31) and the closures (2.32) as

$$\begin{pmatrix} -\frac{1}{2}\mathbf{I} + \mathbf{D}_{11} & \dots & \mathbf{D}_{1n} & \\ \vdots & \ddots & \vdots & \mathbf{B} \\ \mathbf{D}_{n1} & \dots & -\frac{1}{2}\mathbf{I} + \mathbf{D}_{nn} & \\ & \mathbf{B}^T & & \mathbf{0} \end{pmatrix} \begin{pmatrix} \boldsymbol{\eta} \\ \hat{\mathbf{U}} \end{pmatrix} = \begin{pmatrix} -\mathbf{u}^{\infty} \\ -\hat{\mathbf{F}} \end{pmatrix}, \tag{3.9}$$

where $\mathbf{D}_{i,j}$ represents the hydrodynamic interactions between body i and body j and $\mathbf{B} \in \mathbb{R}^{2Nn_p \times 3n_p}$ is the block matrix

$$\mathbf{B} = \begin{pmatrix} 1 & 0 & -(y_1^1 - c_1) & & & & & & \\ 0 & 1 & (x_1^1 - c_1) & & & & & & \\ \vdots & \vdots & \vdots & & & & & & \\ 1 & 0 & -(y_{N_p}^1 - c_1) & & & & & & \\ 0 & 1 & (x_{N_p}^1 - c_1) & & & & & & \\ & & & 1 & 0 & -(y_1^2 - c_2) & & & \\ & & & 0 & 1 & (x_1^2 - c_2) & & & \\ & & & & & \ddots & & & \\ & & & & & & 1 & 0 & -(y_{N_p}^{n_p} - c_{n_p}) \\ & & & & & & 0 & 1 & (x_{N_p}^{n_p} - c_{n_p}) \end{pmatrix}. \tag{3.10}$$

Equation (3.9) is a dense linear system of size $N \times N$, where $N = 2(N_p n_p + N_w n_w) + 3(n_p + n_w)$. Solving this system directly using Gaussian elimination or an LU factorization would require $\mathcal{O}(N^3)$ operations. Traditionally this was seen as the major limitation of boundary integral equations. However, the matrices associated with second-kind Fredholm integral equations have a special properties that allow us to significantly reduce this cost.

Instead of a direct solver that requires building and storing the entire matrix, we use an iterative solver. The Generalized Minimal Residual method (GMRES) [69] is a Krylov method, and requires computing a sequence of matrix vector products to compute an approximate solution to the linear system. The total cost of an iterative solve is $\mathcal{O}(\mathrm{n}_{\mathrm{iter}} C)$, where n_{iter} is the number of iterations and C is the cost of a single matrix-vector product. An ordinary matrix-vector product requires $\mathcal{O}(N^2)$ operations, meaning the cost of a GMRES solve is $\mathcal{O}(n_{\mathrm{iter}} N^2)$. The following theorem allows us to estimate how quickly GMRES converges [78].

Theorem 3.4.1. *Let P_n be the following space of polynomials*

$$P_n := \{\textit{polynomials of degree} \leq n, \textit{ with } p(0) = 1\}.$$

Assume $\mathbf{A} \in \mathbb{C}^{m \times m}$ is diagonalizable, that is $\mathbf{A} = \mathbf{V}\mathbf{\Lambda}\mathbf{V}^{-1}$, where the columns of $\mathbf{V}$ are the eigenvectors of $\mathbf{A}$ and $\mathbf{\Lambda}$ is a diagonal matrix of the eigenvalues of $\mathbf{A}$. Then at GMRES iteration n, the residual r_n satisfies

$$\frac{||\mathbf{r}_n||}{||\mathbf{b}||} \leq \textit{cond}(\mathbf{V}) \inf_{p \in P_n} \sup_{z \in \Lambda(\mathbf{A})} |p(z)|,$$

where $\Lambda(\mathbf{A})$ *is the set of the eigenvalues of* $\mathbf{A}$ *and* $||\ \cdot\ ||$ *is the* ℓ^2 *norm. In other words the norm of the residual depends upon the condition number of* $\mathbf{V}$ *and the how closely we can construct a polynomial of degree* n*, such that* $p(0) = 1$ *and* $p(z) = 0$ *for every eigenvalue of* $\mathbf{A}$*.*

Since $\mathcal{K}$ is a compact operator, the eigenvalues of a discretization of the second-kind Fredholm integral equation $(\lambda - \mathcal{K})z = g$ cluster at λ. Figure 3.4 shows the eigenvalues of the discretized BIE for the problem in Figure 3.1. As N increases, the new eigenvalues cluster around $\lambda = -0.5$. Because of this, the number of GMRES iterations required to solve the system to a desired tolerance is independent of N [14]. Therefore, the cost of solving the linear system using GMRES is $\mathcal{O}(N^2)$.

The structure of the matrix $\mathbf{A}$ allows us to accelerate the solver in two ways. The rows in the symbolic representation (3.9) can be permuted into the new system $\mathbf{Ax} = \mathbf{b}$, where

$$\mathbf{A} = \begin{pmatrix} \mathbf{C}_{11} & \dots & \mathbf{C}_{1n} \\ \vdots & \ddots & \vdots \\ \mathbf{C}_{n1} & \dots & \mathbf{C}_{nn} \end{pmatrix}, \qquad \mathbf{b} = \begin{pmatrix} \boldsymbol{\eta}^1 \\ \hat{\mathbf{U}}_1 \\ \vdots \\ \boldsymbol{\eta}_n \\ \hat{\mathbf{U}}_n \end{pmatrix}, \qquad \mathbf{x} = \begin{pmatrix} -\mathbf{u}_1^\infty \\ -\hat{\mathbf{F}}_1 \\ \vdots \\ -\mathbf{u}_n^\infty \\ -\hat{\mathbf{F}}_n \end{pmatrix}, \tag{3.11}$$

and

$$\mathbf{C}_{ij} = \begin{pmatrix} -\frac{1}{2}\mathbf{I}\delta_{ij} + \mathbf{D}_{ij} & \mathbf{B}_{ij} \\ \mathbf{B}_{ij}^T & \mathbf{0} \end{pmatrix},$$

where $\mathbf{B}_{ij}$ is the (i, j)-th block in (3.10); this will be all zeros if $i \neq j$.

The singular values of $\mathbf{C}_{ij}$ are largest when $i = j$ (self-interactions). However, the singular values of the off-diagonal blocks decay very rapidly. Therefore, when applying GMRES to (3.11), the majority of the iterations are due to the diagonal blocks. Therefore we use a block-diagonal preconditioner, where the diagonal blocks $\mathbf{C}_{ii}$ are inverted directly. Thus we define the preconditioner matrix $\mathbf{P}$

$$\mathbf{P} = \begin{pmatrix} \mathbf{C}_{11} & \dots & \mathbf{0} \\ \vdots & \ddots & \vdots \\ \mathbf{0} & \dots & \mathbf{C}_{nn} \end{pmatrix}.$$

Instead of the original linear system, we now solve $\mathbf{P}^{-1}\mathbf{Ax} = \mathbf{P}^{-1}\mathbf{b}$.

To investigate the block-diagonal preconditioner, we consider the setup of nine circles depicted in Figure 3.6. After discretizing with $N_p = 16$ points per circle, we investigate the singular values of the blocks $\mathbf{C}_{11}$ (intra-body interactions), $\mathbf{C}_{12}$ (neighboring body interactions), and $\mathbf{C}_{19}$ (separated body interactions). These singular values are plotted in Figure 3.5. The singular values of the

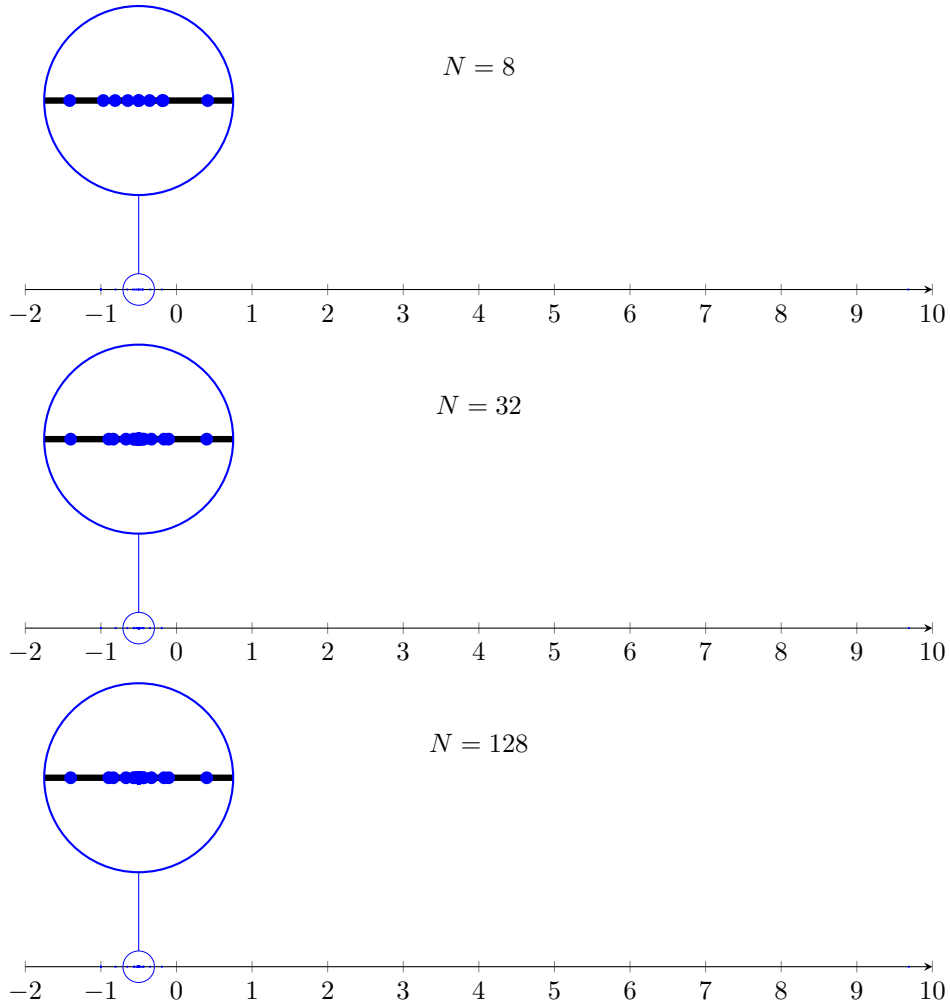

Figure 3.4: The eigenvalues of the discrete linear system associated with the problem shown in Figure 3.1 are plotted along the real axis. As N increases all the new eigenvalues cluster around $x = -0.5$. The is a single eigenvalue between 9 and 10.

diagonal block are much larger than the other two blocks. In particular, the off-diagonal singular values decay quickly to 10^{-8}. GMRES can quickly resolve the off-diagonal interactions, typically in very few iterations, however the diagonal interactions require many more iterations. This justifies the use of a block-diagonal preconditioner, where the diagonal blocks are inverted directly.

The eigenvalues of the unpreconditioned and the preconditioned system are shown in Figure 3.7. For the unpreconditioned system they cluster around $-1/2$, as expected. For the preconditioned system the diagonal blocks have been turned into identity matrices, so the eigenvalues cluster

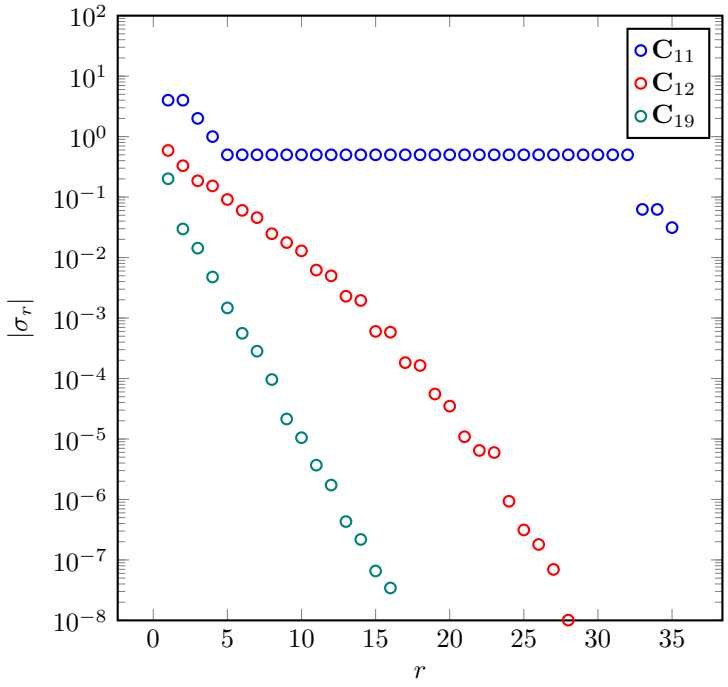

Figure 3.5: Singular values of the blocks $\mathbf{C}_{11}$, $\mathbf{C}_{12}$ and $\mathbf{C}_{19}$. For the diagonal block $\mathbf{C}_{11}$, the singular values do not decay quickly. For the off-diagonal blocks the singular values decay quickly, so GMRES can resolve these interactions in a small number of iterations.

around 1. The unpreconditioned system requires 26 GMRES iterations to solve to a tolerance of 10^{-12}, and the preconditioned system requires 16 iterations.

Theorem 3.4.1 helps us understand speedup the convergence of GMRES after applying the preconditioner. To estimate $\inf_{p \in P_n} \sup_{z \in \Lambda(\mathbf{A})} |p(z)|$, we perform a least squares fit of a polynomial of degree n through the point $(0, 1)$ and with roots at the eigenvalues of $\mathbf{A}$ (or $\mathbf{P}^{-1}\mathbf{A}$ in the preconditioned case). In Figure 3.8, we plot the maximum value this polynomial takes at one of the eigenvalues of the matrix $\mathbf{A}$ and $\mathbf{P}^{-1}\mathbf{A}$. We see that as n increases, $\sup_{z \in \Lambda(\mathbf{A})} |p(z)|$ decreases much faster for the preconditioned matrix than for the original matrix.

We can also speed up the matrix-vector products by noticing that double-layer kernel (2.10b) decays as r^{-1}. This means that far field points can be represented by a multipole expansion. The Fast Multipole Method (FMM) [27,28] uses these multipole expansions along with local expansions to reduce matrix multiplication from an $\mathcal{O}(N^2)$ procedure to $\mathcal{O}(N)$.

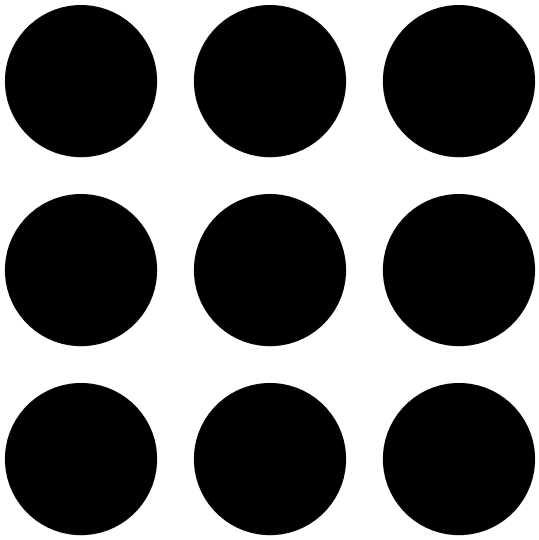

Figure 3.6: A setup to test our block-diagonal preconditioner. Here we have nine circular bodies placed in an unbounded background flow.

3.5 Time Stepping

The solution of (3.9) includes the translational and angular velocities of the bodies. The quasi-static approximation lets us update the centers and orientations according to the ordinary differential equations (ODEs):

$$\frac{\mathrm{d}}{\mathrm{d}t}\mathbf{c}_k^p = \mathbf{U}_k, \qquad k = 1, \ldots, n_p, \tag{3.12a}$$

$$\frac{\mathrm{d}}{\mathrm{d}t}\theta_k^p = \omega_k, \qquad k = 1, \ldots, n_p. \tag{3.12b}$$

After solving (3.12) for a specified time interval, the fluid instantaneously adjusts its velocity according to this new geometry. These ODEs (3.12) are quite simple and can be solved numerically by standard methods. We will investigate the forward Euler method

$$\mathbf{c}_k^p(t^{n+1}) = \mathbf{c}_k^p(t^n) + \Delta t \mathbf{U}_k(t^n), \qquad k = 1, \ldots, n_p,$$

$$\theta_k^p(t^{n+1}) = \theta_k^p(t^n) + \Delta t \omega_k(t^n), \qquad k = 1, \ldots, n_p,$$

and the second order Adams-Bathforth method

$$\mathbf{c}_k^p(t^{n+1}) = \mathbf{c}_k^p(t^n) + \frac{3\Delta t}{2}\mathbf{U}_k(t^n) - \frac{\Delta t}{2}\mathbf{U}_k(t^{n-1}), \qquad k = 1, \ldots, n_p,$$

$$\theta_k^p(t^{n+1}) = \theta_k^p(t^n) + \frac{3\Delta t}{2}\omega_k(t^n) - \frac{\Delta t}{2}\omega_k(t^{n-1}), \qquad k = 1, \ldots, n_p.$$

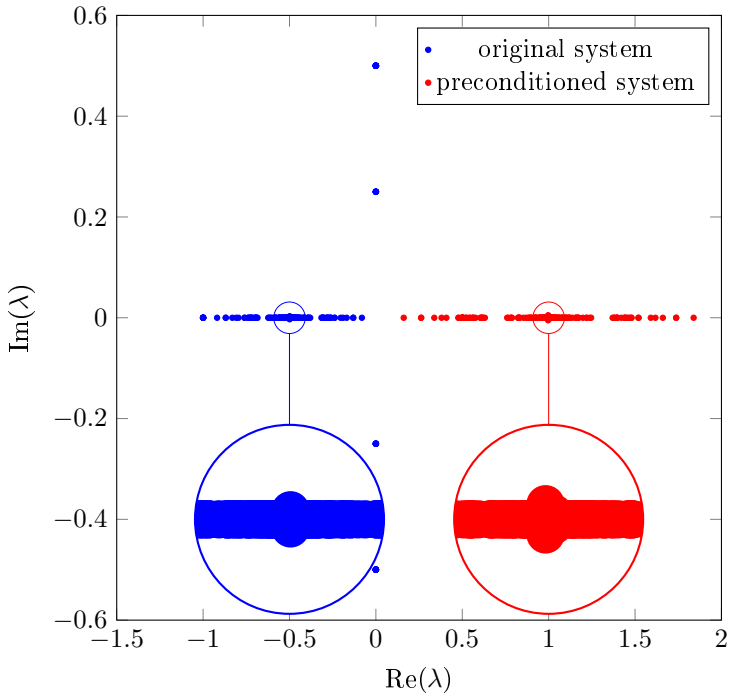

Figure 3.7: Eigenvalues of the linear system (3.11) for the geometry shown in Figure 3.6 using $N_p = 16$ points per circle. In the unpreconditioned cas, the eigenvalues cluster around $-1/2$, while in the preconditioned case they cluster around 1. Solving the linear system whose eigenvalues are the blue points requires 26 GMRES iterations, while solving the linear system whose eigenvalues are the red points requires 16 GMRES iterations.

Higher-order methods such as spectral deferred correction [64,65] or Runge-Kutta methods [2] can also be applied. However, when bodies get very close, even high-order time stepping may not be accurate enough to prevent non-physical overlap between rigid bodies. A method to avoid overlap between rigid bodies is discussed in Chapter 4.

3.5.1 Locally Implicit

The spatial discretization discussed above is known as *globally implicit* as the density function on all bodies are solved for simultaneously at each time step. An alternative discretization [51,68] lags the density function for inter-body interactions. This is known as a *locally implicit* discretization. Instead of a dense linear system, a locally implicit discretization leads to a block-diagonal linear system, and each block can be inverted directly. This linear system is computationally cheaper

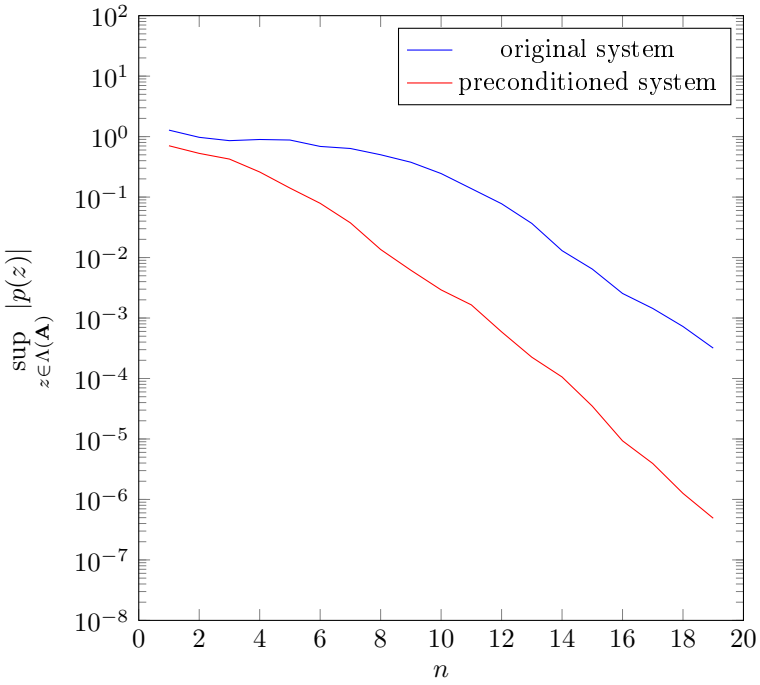

Figure 3.8: Plots of $\sup_{z\in\Lambda(\mathbf{A})} |p(z)|$ vs n for the original and preconditioned case with eigenvalues as shown in Figure 3.7. At $n = 20$, $\sup_{z\in\Lambda(\mathbf{A})} |p(z)| = 4.3 \times 10^{-7}$ for the preconditoned matrix, while $\sup_{z\in\Lambda(\mathbf{A})} |p(z)| = 3.2\times10^{-4}$ for the unpreconditioned matrix.

to solve, is trivial to parallelize, and works very well for dilute suspensions. However, as we will see, for dense suspensions locally implicit methods suffer from stability restrictions, particularly for concentrated suspensions. A globally implicit time stepper remains stable for larger time steps, and the added cost of solving a full dense linear system can be more than offset by the allowable larger time steps.

3.6 Examples

3.6.1 Jeffery Orbit

As mentioned in the introduction, the motion of a single ellipse in a shear flow has an analytic solution [35]. In $\mathbb{R}^2$, an elliptical rigid body with aspect ratio λ suspended in a shear flow with shear rate $\dot\gamma$ rotates with period $2\pi|\dot\gamma|(\lambda + \lambda^{-1})$ according to

$$\theta(t) = \tan^{-1}\left(\lambda^{-1}\tan\left(\frac{\lambda\dot\gamma t}{\lambda^2+1}\right)\right).$$

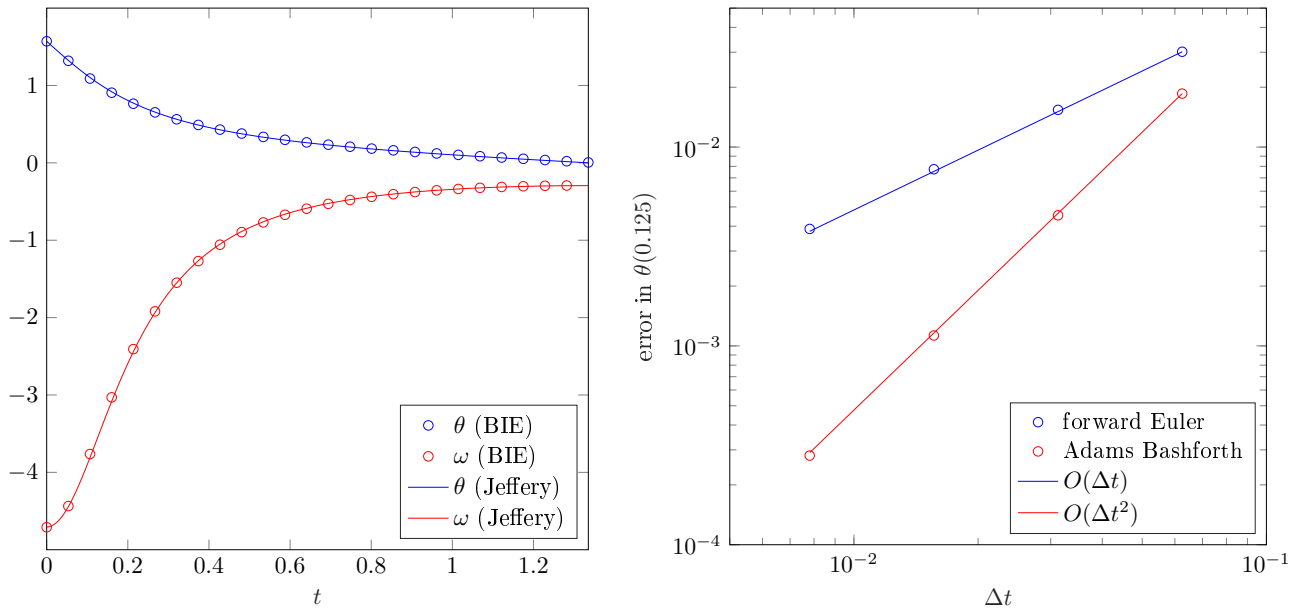

Figure 3.9: Numerical simulation of a single rigid body in shear flow. Left: The angle and angular velocity as a function of time (marks), compared to the analytic result (lines). Right: A convergence study of the error between the numerical approximation of the angle and the analytic angle after half a period. As expected the forward Euler method converges with first-order accuracy, while Adams-Bashforth converges with second-order accuracy.

To test the time stepping routine, Figures 3.9 shows a convergence study comparing the numerical solution to the exact solution. We find that forward Euler converges with accuracy $\mathcal{O}(\Delta t)$ and Adams-Bashforth converges with accuracy $\mathcal{O}(\Delta t^2)$.

3.6.2 Multiple Bodies

One of the advantages of BIEs over Stokesian dynamics is the ability to simulate suspensions of arbitrarily smooth bodies. Consider a rigid body with a boundary parameterized by

$$\boldsymbol{\phi}(s) = ((1 + a\cos(ks))\cos(s), (1 + a\cos(ks))\sin(s)), \qquad s \in [0, 2\pi). \tag{3.13}$$

Here $k \in \mathbb{Z}$ determines the number of bumps on the surface of the body, and $a \in [0, 1]$ the magnitude of these bumps. These boundaries are shown for multiple values of k and a in Figure 3.10.

To test our algorithms with multiple bodies, we perform a simulation of eight rigid bodies parameterized according to (3.13). We take $a = 0.2$ for all the bodies, and several values of k.

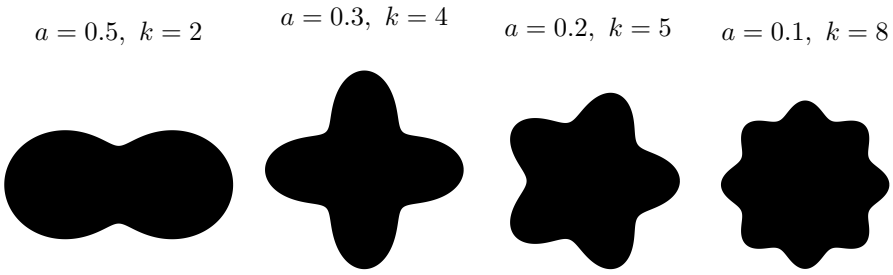

Figure 3.10: Plots of the parameterization (3.13) for various values of k and a.

Snapshots of the simulation are shown in Figure 3.11. As in the case of a single elliptical body, the rigid bodies rotate. However, hydrodynamic interactions push the bodies either above or below the stationary plane, so they also translate either to the right or to the left.

As the rigid bodies become close, near singular integration needs to be used if the discretization is not fine enough. Table 3.3 shows the error at the time horizon for a range of discretization levels compared to an overrefined solution generated with 512 points per rigid body. Taking Δt to be 0.4, without near singular integration, $N_p = 16$ is not enough to prevent overlap between rigid bodies. Even with near singular integration, $N_p = 8$ is not enough to prevent overlaps. However, using near singular integration we are able to prevent overlap with $N_p = 16$. For $N_p > 16$ we do not require near singular integration to prevent collisions. The error from this point decay spectrally, up to the tolerance of GMRES.

We can also use this test case to examine the efficiency of the block-diagonal preconditioner and the FMM. Tables 3.4 and 3.5 show the CPU times with and without the preconditioner using regular matrix-vector products and the FMM, respectively. In both cases the preconditioner significantly reduces the total number of required GMRES iterations. Even though constructing the block-diagonal preconditioner requires inverting eight $N_p \times N_p$ matrices per time step, this still results in computational savings. Without using the FMM, the CPU time scales as $\mathcal{O}(N_p^2)$ for large N, while the FMM reduces this to a method that is roughly $\mathcal{O}(N_p)$. Figure 3.12 illutstrates the timings in Tables 3.4 and 3.5.

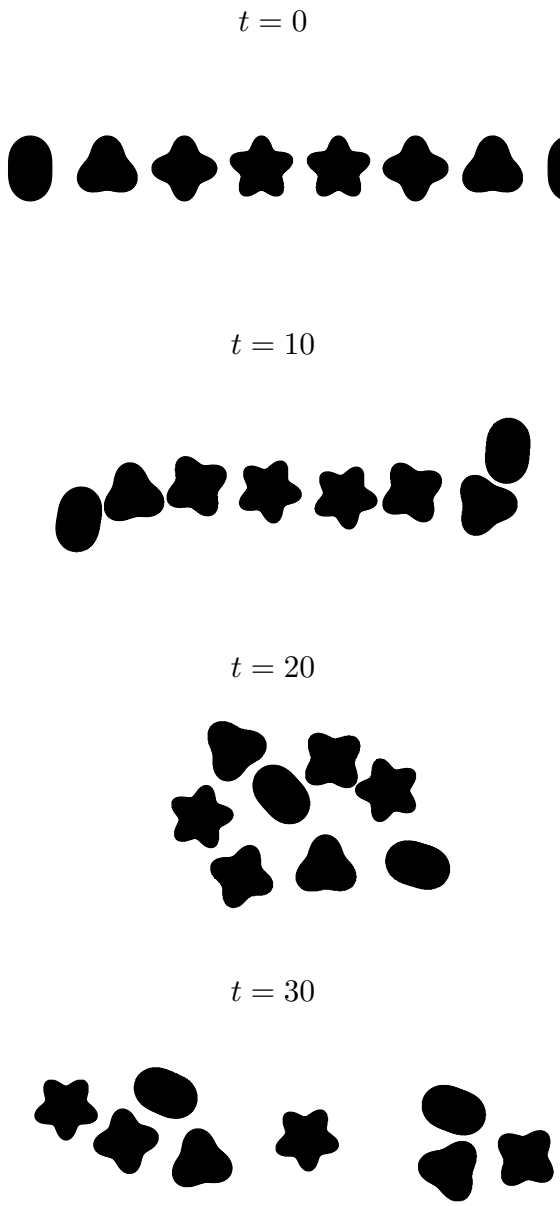

Figure 3.11: Snapshots of eight rigid bodeis in an unbounded shear flow. All the bodies are parameterized according to (3.13), with $a = 0.2$. From left to right, $k = 2, 3, 4, 5, 5, 4, 3, 2$. Instead of rotating in place like a single body, hydrodynamic interactions cause the bodies to shift vertically from $y = 0$ and therefore translate.

Table 3.3: Study of a refinement in N_p for eight bodies in an unbounded shear flow. Snapshots of the simulation are shown in Figure 3.11. The errors reported are the difference between the rigid body centers at the time horizon and an overrefined solution generated with $N_p = 512$ points. Simulations in bold indicate that a collision between rigid bodies occured before the time horizon. With a time step size of $\Delta t = 0.4$, and without near-singular integration (NSI), we require at least 32 points to avoid collisions between rigid bodies. Using NSI, we require only 16 points to avoid collisions. Note that for $N_p \geq 256$ the rigid bodies never get within an arclength of each other, so the NSI and the trapezoid rule solutions coincide.

N_p	error (without NSI)	error (with NSI)
8	$\mathbf{2.90 \times 10^1}$	$\mathbf{1.48 \times 10^1}$
16	$\mathbf{1.14 \times 10^2}$	1.06×10^1
32	2.25×10^0	3.63×10^0
64	3.30×10^{-2}	9.22×10^{-1}
128	7.53×10^{-4}	8.85×10^{-2}
256	1.55×10^{-9}	1.55×10^{-9}

Table 3.4: CPU timings for eight rigid bodies in an unbounded shear flow at various resolutions. The GMRES matrix-vector products are computed directly. As predicted, the number of GMRES iterations does not grow with N_p. For large N_p, the cost of the linear solve is dominant and the runtime scales as $\mathcal{O}(N_p^2)$. Using a block-diagonal significantly preconditioner decreases the number of required GMRES iterations each time step, and also reduces the CPU runtime.

N_p	No preconditioner		Preconditioner	
	# matvecs	CPU time (s)	# matvecs	CPU time (s)
8	9771	51	1832	17
16	10678	64	3129	27
32	11631	97	3169	31
64	9593	117	4187	66
128	8217	291	3312	137
256	7572	1075	3323	536
512	7260	6653	3390	3413

3.6.3 Rotors

We are also interested in modeling rigid, active, particles. Active particles are particles that move by a net force or torque. These can be living cells, for example, bacteria, or non-living matter subject to an external force or torque coming from, for example, a magnetic field. One class of

Table 3.5: CPU timings for eight rigid bodies in an unbounded shear flow at various resolutions. The GMRES matrix-vector products are accelerated using FMM. As predicted, the number of GMRES iterations does not grow with N_p. For large N_p, the cost of the linear solve is dominant and the runtime scales as $\mathcal{O}(N_p)$. Using a block-diagonal significantly preconditioner decreases the number of required GMRES iterations each time step, and also reduces the CPU runtime.

N_p	No preconditioner		Preconditioner	
	# matvecs	CPU time (s)	# matvecs	CPU time (s)
8	10801	168	4441	75
16	9857	200	3395	61
32	7693	144	3125	75
64	7367	221	3238	118
128	7346	436	3264	233
256	7357	1116	3317	593
512	7375	3245	3380	1692

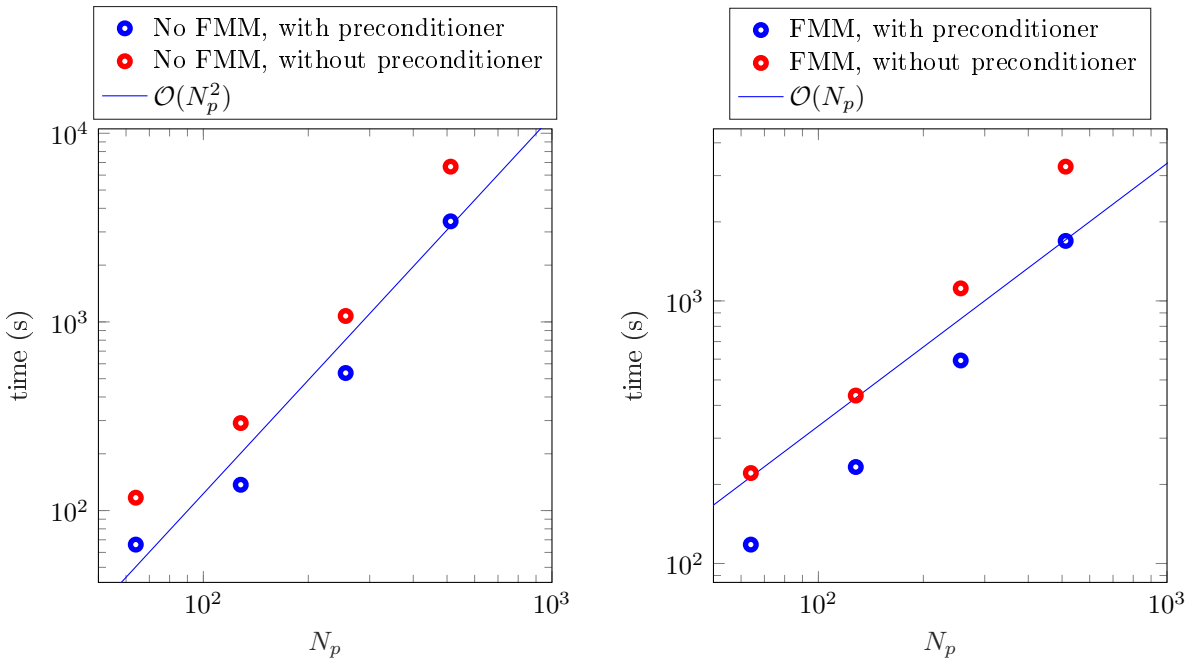

Figure 3.12: Code timings to advance eight rigid bodies in unbounded shear flow to the time horizon. For large enough N_p, the timing scales as $\mathcal{O}(N_p^2)$ when GMRES uses full matrix-vector products, while it scales as $\mathcal{O}(N_p)$ using the FMM accelerated GMRES. In both cases, preconditioning speeds up the computation by a constant factor.

active particles are rotors, whose motion is driven by an external torque. Rotors have been observed experimentally, where the torque comes from chemicals, light, or magnetic or electric fields. These particles can demonstrate complicated behaviors, from periodic and chaotic motion [52] to large scale collective organization [83]. To adopt our method to handle rotors, we must simply specify the net torque on each rotor in equation (2.31c).

In Figure 3.13, we simulate four circular bodies, each undergoing an identical positive (clockwise) net torque. If the four bodies are initially placed at the corners of a rectangle, they move in a periodic orbit around the center of the rectangle. If however, the bodies start in a random initial configuration, the demonstrate chaotic motion. This replicates a result from [52].

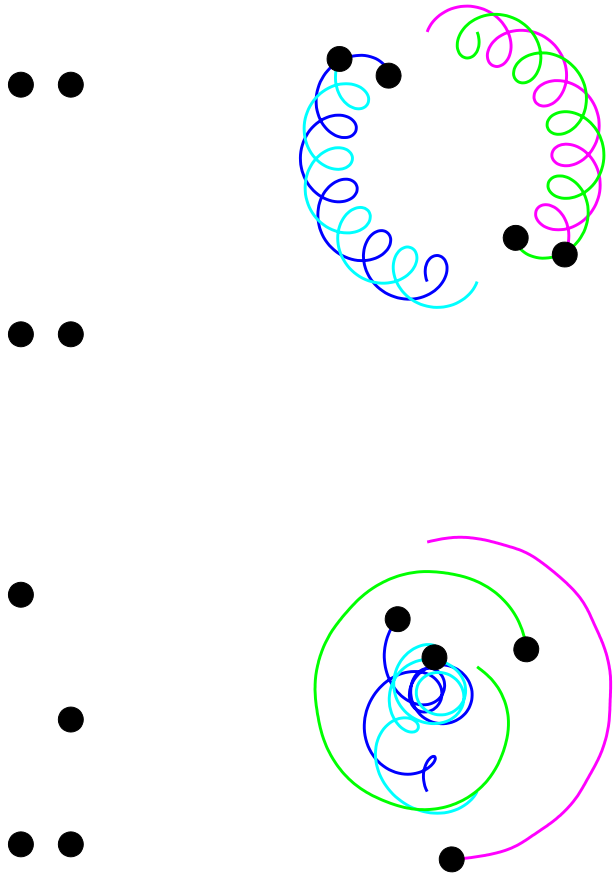

Figure 3.13: Four rotors in quiescent background flow. The particle are all undergoing a net torque in the same direction. This causes them to spin. As they spin they interact hydrodynamically with all the other particles. These hydrodynamic interactions can cause a periodic motion if the four particles are initially arranged in a rectangle, or chaotic motion if the particles are initially placed at random. This matches the results in [52].

CHAPTER 4

RESOLVING CONTACT

Complementarity problems appear in many scientific or engineering applications [23]. They have been used to model contact mechanics [4, 9, 74] and structural mechanics [54] among many other problems. Complementarity problems are problems that are constrained by a set of complementarity conditions, in other words a requirement that the product of two or more non-negative quantities is zero. For example, in contact mechanics the force between two objects is typically complementarity to the distance between them. That is, the force between two objects is zero, unless the distance between them is zero. This suggests a procedure for preventing contact between rigid bodies in our simulations.

1. Starting from a contact-free configuration, perform a single time step to obtain a new configuration.
2. If the new configuration is contact-free, accept this configuration and move to the next time step. Otherwise, for each rigid body pair in contact, apply a repulsion force until the distance between them is zero.
3. Once the distance between all rigid body pairs is at least zero, the configuration is contact-free. Accept this configuration and move to the next time step.

Before formalizing this procedure, some background on complementarity problems is provided. This discussion follows largely from [10] and [20].

4.1 Description and Properties of Complementarity Problems

A one-dimensional complementarity problem is the following. Given two variables $x, y \in \mathbb{R}$, the complementarity constraint states

$$y > 0 \Rightarrow x = 0, \qquad \text{or} \qquad x > 0 \Rightarrow y = 0.$$

This means that if one variable is positive then the other is zero. This can be stated as the constraint

$$x \geq 0, \qquad y \geq 0, \qquad xy = 0,$$

or in the more compact notation

$$0 \leq y \perp x \geq 0.$$

The space that satisfies this condition is any point on the positive x or postive y axis. If there is a linear relationship between x and y, then we have the linear complementarity problem (LCP)

$$y = ax + b, \qquad 0 \leq y \perp x \geq 0, \tag{4.1}$$

where $a, b \in \mathbb{R}$. Now the solution space is the intersection of the $y = ax + b$ and the positive x and y axes. Depending on the values of a and b, this problem admits either zero, one, two, or infinite solutions. These solutions are summarized in Table 4.1 and plotted in Figure 4.1.

Table 4.1: The number of solutions to the LCP (4.1) depends on the signs of a and b. There are nine possible cases.

	$a < 0$	$a = 0$	$a > 0$
$b < 0$	0	0	1
$b = 0$	1	∞	1
$b > 0$	2	1	1

By eliminating y, we can rewrite (4.1) as

$$ax + b \geq 0, \tag{4.2a}$$

$$x \geq 0, \tag{4.2b}$$

$$x(ax + b) = 0. \tag{4.2c}$$

Now consider the following constrained minimization problem

$$x^* = \arg\min_{x \geq 0} x \left(\frac{1}{2} ax + b \right),$$

which is a quadratic program (QP). The Lagrangian of this problem is

$$\mathcal{L}(x, y) = x \left(\frac{1}{2} ax + b \right) - yx,$$

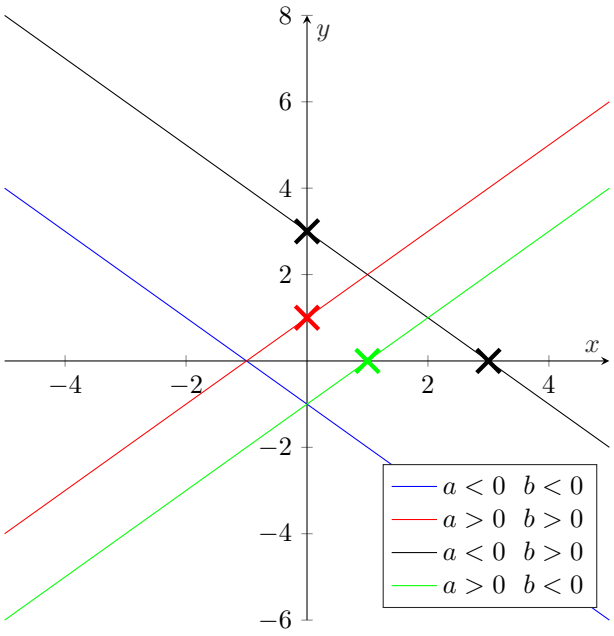

Figure 4.1: Examples of one-dimensional complementarity problems corresponding to (4.1). Solutions are marked with a cross. If a is greater than zero, then there exists a single solution. If a is less than zero, then two solutions exist if b is greater than zero, or no solutions exist if b is less than zero. If $a = 0$ and $b \geq 0$, then there is a single solution at $(0, b)$, unless b is also zero, in which case $(x, 0)$ is a solution for all $x \geq 0$.

where y is a Lagrange multiplier needed to enforce the constraint $x \geq 0$. The first-order optimality (KKT) conditions [75] are

$$\nabla_x \mathcal{L}(x, y) = ax + b - y = 0, \tag{4.3a}$$

$$x \geq 0, \tag{4.3b}$$

$$y \geq 0, \tag{4.3c}$$

$$xy = 0, \tag{4.3d}$$

where the last condition (4.3d) is the complementary slackness condition. From (4.3a) we have that $y = ax + b$, which allows us to rewrite (4.3) as

$$ax + b \geq 0,$$

$$x \geq 0,$$

$$x(ax + b) = 0,$$

which is identical to (4.2). Thus we have turned a constrained minimization problem into an LCP. The necessary optimality conditions for any QP lead to an LCP [57]. We will see this in Section 4.3, when we transform the constrained variational Stokes equations into a sequence of LCPs.

A general LCP in $\mathbb{R}^n$ is the following: Given $\mathbf{q} \in \mathbb{R}^n$ and $\mathbf{M} \in \mathbb{R}^{n \times n}$, find $\mathbf{w} \in \mathbb{R}^n$ and $\mathbf{z} \in \mathbb{R}^n$ such that

$$\mathbf{w} - \mathbf{M}\mathbf{z} = \mathbf{q}, \qquad \mathbf{w}, \mathbf{z} \geq \mathbf{0}, \qquad \mathbf{w} \cdot \mathbf{z} = 0. \tag{4.4}$$

Often, such a complementarity problem will be written in the compact notation

$$0 \leq \mathbf{q} + \mathbf{M}\mathbf{z} \perp \mathbf{z} \geq 0.$$

This notation is shorthand for saying that every component of $\mathbf{z}$ must be positive, every component of $\mathbf{w} = \mathbf{q} + \mathbf{M}\mathbf{z}$ must be positive, and the inner product $\mathbf{w} \cdot \mathbf{z}$ must be zero. It follows from these conditions that at least one of z_i or w_i must be zero for each $i = 1, \ldots, N$.

More generally, there exist *nonlinear* complementarity problems (NCPs) in $\mathbb{R}^n$: Given a mapping $F : \mathbb{R}^n \to \mathbb{R}^n$, find $\mathbf{z} \in \mathbb{R}^n$ such that

$$\mathbf{z} \geq \mathbf{0}, \qquad F(\mathbf{z}) \geq \mathbf{0}, \qquad \mathbf{z} \cdot F(\mathbf{z}) = 0. \tag{4.5}$$

In shorthand notation we write

$$0 \leq F(\mathbf{z}) \perp \mathbf{z} \geq 0.$$

4.2 Solution Procedures

Solution procedures for complementarity problem is an active field of research. Depending on the properties of the matrix $\mathbf{M}$, an LCP may be easy to solve with a greedy type algorithm, it may be NP-hard, or it may be somewhere in between. It has been shown that an LCP with general integer data is NP-hard [57], and the only algorithm that is guaranteed to find a solution is an exhaustive search.

Besides exhaustive search, there are several possible avenues to attempt to solve an LCP. These include pivot methods, splitting methods, and Newton-type methods. We focus on a Newton-type method known as the Fischer-Newton method that uses Newton's method to solve the Fischer-Burmeister function. To solve NCPs, they can be linearized with Newton's method to create a sequence of LCPs whose solutions converge to the solution of the NCP.

4.2.1 Linear Complementarity Problems

Any LCP can be converted into a root finding problem by defining the the Fischer-Burmeister function $\phi : \mathbb{R}^n \to \mathbb{R}^n$ [24]

$$\boldsymbol{\phi}(\mathbf{z}) = \boldsymbol{\phi}(\mathbf{z}, \mathbf{w}(\mathbf{z})) = \begin{pmatrix} \sqrt{z_1^2 + w_1^2} - z_1 - w_1 \\ \vdots \\ \sqrt{z_n^2 + w_n^2} - z_n - w_n \end{pmatrix}. \tag{4.6}$$

Equation (4.6) satisfies the property

$$\phi_i(\mathbf{z}, \mathbf{w}) = 0 \qquad \Longleftrightarrow \qquad 0 \leq z_i \perp w_i \geq 0.$$

This can be verified by looking at at each possible case:

	$z_i < 0$	$z_i = 0$	$z_i > 0$
$w_i < 0$	$\phi_i > 0$	$\phi_i > 0$	$\phi_i > 0$
$w_i = 0$	$\phi_i > 0$	$\phi_i = 0$	$\phi_i = 0$
$w_i > 0$	$\phi_i > 0$	$\phi_i = 0$	$\phi_i < 0$

The challenge of finding the root of $\boldsymbol{\phi}$ is that it is not differentiable if $z_i = w_i = 0$.

Nonetheless, to solve $\boldsymbol{\phi} = \mathbf{0}$, we apply Newton's method. Starting from an initial guess $\mathbf{z}^0$, we choose a search direction by solving the linear system

$$\mathbf{J}\Delta\mathbf{z}^k = \phi(\mathbf{z}^k),$$

and then update $\mathbf{z}^k$ according to $\mathbf{z}^{k+1} = \mathbf{z}^k + \tau\Delta\mathbf{z}^k$, where τ is some step size is determined in the usual ways, for example using backtracking. Ordinarily $\mathbf{J}$ would be the Jacobian of $\boldsymbol{\phi}$, however, since $\boldsymbol{\phi}$ not differentiable everywhere, the classical Jacobian may not work.

Instead of using the classical Jacobian, we make use of the generalized Jacobian. Before defining the generalized Jacobian, we must first define the B-subdifferential.

Definition 4.2.1. *Consider a Lipshitz continuous function $F : \mathbb{R}^n \to \mathbb{R}^m$ and assume that F is continuously differentiable on the subdomain $\mathcal{D} \subset \mathbb{R}^n$. The B-subdifferential of F at a point $\mathbf{x}$ is the set*

$$\partial_B F(\mathbf{x}) := \left\{ \lim_{\mathbf{y} \to \mathbf{x}} \nabla F \mid \forall \mathbf{y} \in \mathcal{D} \right\}.$$

In other words, $\partial_B F(\mathbf{x})$ is the set of limiting values of the gradient of F at $\mathbf{x}$. The generalized Jacobian of F at $\mathbf{x}$, $\partial F(\mathbf{x})$, is the convex-hull of $\partial_B F(\mathbf{x})$. If F is continuously differentiable at $\mathbf{x}$, then this limit is unique and the generalized Jacobian coincides with the classical Jacobian.

As a simple example, consider the function $f(x) = |x|$. This function is continuously differentiable everywhere, except $x = 0$. The B-subdifferential of $f(0)$ is the set of limiting values of the continuous derivative at $x = 0$. This is $+1$ if we are approaching from the right and -1 if we are approaching from the left. Therefore the B-subdifferential of $f(0)$ is the set $\{-1, 1\}$. The generalized derivative of f at $x = 0$ is the convex hull of this B-subdifferential, namely $\{m \mid |m| \leq 1\}$.

The two-dimensional analogue of the absolute value is the function $e : \mathbb{R}^2 \to \mathbb{R}$ given by $e(\mathbf{z}) = ||\mathbf{z}|| = \sqrt{\mathbf{z}^T\mathbf{z}}$. This function has the gradient $\nabla e = \mathbf{z}^T/||\mathbf{z}||$ but is not continuously differentiable at the origin. Letting $x = r\cos\theta$ and $y = r\sin\theta$, $\theta \in [0, 2\pi)$, the B-subdifferential at the origin is

$$\partial_B e(\mathbf{0}) = \left\{ \lim_{r\to 0} (r\cos\theta/r,\ r\sin\theta/r) \right\} = \{\cos\theta,\ \sin\theta\} = \{\mathbf{v} \mid ||\mathbf{v}|| = 1\}.$$

Then, the generalized derivative of $e(0)$ is the convex hull of this set,

$$\partial e(0) = \{\mathbf{v} \mid ||\mathbf{v}|| \leq 1\}.$$

We can apply this result to determine the generalized Jacobian of the Fischer-Burmeister function (4.6). Letting $\mathbf{y} = [z_i, w_i]$, each component of the range of ϕ can be written as

$$\phi_i(\mathbf{y}) = e(\mathbf{y}) - g(\mathbf{y}),$$

where $g(\mathbf{y}) = [1, 1]^T[z_i, w_i] = [1, 1]^T\mathbf{y}$. For $\mathbf{y} \neq \mathbf{0}$, the gradient of ϕ_i is

$$\nabla\phi_i = \frac{\mathbf{y}^T}{||\mathbf{y}||} - [1, 1]^T.$$

For $\mathbf{y} = 0$, the generalized gradient is

$$\partial\phi_i = \{\mathbf{v} - [1, 1] \mid ||\mathbf{v}|| \leq 1\}.$$

Combining this result with the chain rule, and recalling that $\mathbf{w} = \mathbf{q} + \mathbf{M}\mathbf{z}$, the generalized Jacobian of $\phi(\mathbf{z}, \mathbf{w})$ is

$$\mathbf{J} = \frac{\partial\phi_i}{\partial z_j} = \begin{cases} \{\mathbf{v} - [1, 1] \mid ||\mathbf{v}|| \leq 1\}, & z_i = w_i = 0, \\ \dfrac{z_i}{(z_i + w_i)^{1/2}} (z_k\delta_{kj} + 2w_k M_{kj}) - M_{ij} - \delta_{ij}, & \text{otherwise.} \end{cases}$$

At $z_i = w_i = 0$ the generalized Jacobian is a set of vectors. A descent direction at this point is chosen to be a particular element from this set. There are several possible choices for the Fischer-Burmeister function:

- **random**: the gradient is picked randomly from the set of possible directions
- **zero**: the gradient is taken to always be zero
- **perturbation**: the gradient can be approximated by perturbing z_i to some small value ϵ
- **approximation**: the gradient is approximated using finite difference

We will use a perturbation to compute the gradient at $z_i = w_i = 0$.

4.2.2 Nonlinear Complementarity Problems

To solve the NCP (4.5), a common procedure is to generate a sequence of solutions $\{\mathbf{z}^k\}$, such that $\mathbf{z}^{k+1}$ is a solution to the LCP

$$0 \leq \mathbf{q}^k + \mathbf{M}^k \mathbf{z}^{k+1} \perp \mathbf{z}^{k+1} \geq 0,$$

where $\mathbf{q}^k$ and $\mathbf{M}^k$ approximate $F(\mathbf{z})$ near $\mathbf{z}^k$. There are multiple choices for $\mathbf{q}^k$ and $\mathbf{M}^k$ including Newton's method where $\mathbf{M}^k = \nabla F(\mathbf{z}^k)$ and $\mathbf{q}^k = F(\mathbf{z}^k) - \mathbf{M}^k \mathbf{z}^k$. As with Newton's method for nonlinear equations, provided we are in the neighborhood of the solution $\mathbf{z}^*$ and F is sufficiently differentiable, this method's iterates converge quadratically to the solution of the NCP.

4.3 Variational Stokes Equations

In the algorithm outlined in Chapter 3, at each time step we solve the Stokes equations, and then advance the rigid bodies according to the ODEs (3.12). If two bodies are sufficiently close and Δt is sufficiently large, then they may overlap after advancing in time. This is especially the case in concentrated suspensions, where it is certain that there will be rigid bodies in close proximity to one another. Spatial and temporal adaptivity [41] help, but as the distance between bodies approaches zero, the required spatial refinement and time step size become computationally infeasible.

As an alternative to adaptivity, a repulsion force can be used. This can be a force based on a Morse or Lennard-Jones potential [25, 49] or a spring based force [37, 85], however both these methods require tuning parameters and add stiffness to the system. Therefore, small time step sizes

are required when rigid bodies are sufficiently close. In addition, these methods do not explicitly guarantee that overlap is avoided.

A different approach [51] uses a repulsion force that is free of tuning parameters and explicitly guarantees that overlap between bodies is avoided. This is done by imposing a constraint on the variational form of the Stokes equations. In Section 4.4 we define a metric $\mathbf{V}$ to measure overlap between bodies. This metric will be defined so that $\mathbf{V} < \mathbf{0}$ means that there is overlap between bodies. We will thus constrain the solution to the Stokes equations $\mathbf{u}$ to be a solution such that, after advancing the rigid bodies in time, $\mathbf{V} \geq \mathbf{0}$.

Without defining $\mathbf{V}$ exactly, we can define a few of its properties. First, $\mathbf{V}$ must be a function of $\mathbf{u}$, otherwise constraining $\mathbf{u}$ by $\mathbf{V}$ would not be possible. Second, we let $\mathbf{V}$ a vector, with each component measuring the overlap between a pair of bodies. Thus $\mathbf{V} \in \mathbb{R}^m$, where $m = \binom{n_p}{2}$.

The Stokes equations are the Euler-Lagrange equations of the constrained minimization problem,

$$\begin{aligned} &\min_{\mathbf{u}} \int_V \left(\frac{1}{2} \nabla \mathbf{u} : \nabla \mathbf{u} \right) \, \mathrm{d}V, \\ &\text{such that } \nabla \cdot \mathbf{u} = 0. \end{aligned}$$

By applying the KKT conditions, the pressure p enters as a Lagrange multiplier to enforce the incompressibility constraint. As mentioned above, in addition to the incompressibility constraint, we also wish to enforce $\mathbf{V} \geq \mathbf{0}$, and this leads to the constrained minimization problem

$$\min_{\mathbf{u}} \int_V \left(\frac{1}{2} \nabla \mathbf{u} : \nabla \mathbf{u} \right) \, \mathrm{d}V, \tag{4.8a}$$

$$\text{such that } \nabla \cdot \mathbf{u} = 0, \quad \mathbf{V} \geq 0. \tag{4.8b}$$

The Lagrangian of (4.8) is

$$\mathcal{L}(\mathbf{u}, p, \boldsymbol{\lambda}) = \int_V \left(\frac{1}{2} \nabla \mathbf{u} : \nabla \mathbf{u} - p \nabla \cdot \mathbf{u} \right) \, \mathrm{d}V + \boldsymbol{\lambda} \cdot \mathbf{V}(\mathbf{u}),$$

where $\boldsymbol{\lambda} \in \mathbb{R}^m$ is a Lagrange multiplier that enforces the no overlap constraint.

The first-order optimality conditions of the Lagrangian yield the forced Stokes equations,

$$-\Delta \mathbf{u} + \nabla p = \mathbf{F}_r(\mathbf{x}), \tag{4.9a}$$

$$\nabla \cdot \mathbf{u} = 0, \tag{4.9b}$$

$$\mathbf{V}(\mathbf{u}) \geq \mathbf{0}, \tag{4.9c}$$

$$\boldsymbol{\lambda} \geq \mathbf{0}, \tag{4.9d}$$

$$\boldsymbol{\lambda} \cdot \mathbf{V}(\mathbf{u}) = 0, \tag{4.9e}$$

where the repulsion force $\mathbf{F}_r$ is

$$\mathbf{F}_r(\mathbf{x}) = \int_S \mathrm{d}_{\mathbf{u}} \mathbf{V}^T \boldsymbol{\lambda} \delta(\mathbf{x} - \mathbf{X}(s, t_0)) \, \mathrm{d}S. \tag{4.10}$$

Note that $\mathrm{d}_{\mathbf{u}}\mathbf{V}$ is nonzero only on the boundaries of the rigid bodies, since changing the velocity at any other point inside the fluid domain does not change the amount of overlap between bodies. This is important, as it lets us reuse the boundary integral formulation developed in the previous chapters without introducing any volume integrals. The problem (4.9) is an NCP since the relationship between $\mathbf{V}$ and $\mathbf{u}$ is nonlinear.

4.4 Space-Time Interference Volumes

One possible choice for $\mathbf{V}$ is a signed distance function [81], where we measure the closest distance between each rigid body pair. If two bodies are overlapping then this distance is negative. This choice, though simple, has the disadvantage that if too large a time step is taken, contact may be missed even though the configuration is contact-free at the end of the time step. Space-time interference volumes [30, 51] mitigate this issue by computing the volume swept out in the space-time plane during each time step. This value will be negative if there is contact at any point during the time step, even if the final configuration is contact-free.

Let $\mathbf{x}(s, \tau)$ be a paramaterization of the boundary of the domain (all rigid bodies and walls) at time τ between an initial contact-free time t_0 and $t_1 = t_0 + \Delta t$. The collection of points $\mathbf{x}(s, \tau)$ define a moving boundary $S(\tau)$. For each point $\mathbf{x}(s, \tau)$ let $\tau_I(s)$, $t_0 \leq \tau_I \leq t_1$ be the intersection time. That is, $\tau_I(s)$ is the first instance when $\mathbf{x}(s, \tau)$ comes into contact with a different point on $S(\tau_I)$. Then, the space-time volume for the time interval $[t_0, t_1]$ is

$$V^C(S, t_1) = -\int_{S(t_0)} \int_{\tau_I(s)}^{t_1} \sqrt{\epsilon^2 + (\mathbf{u}(s, \tau) \cdot \mathbf{n}(s, \tau))^2} \, \mathrm{d}\tau \, \mathrm{d}s, \tag{4.11}$$

where $\mathbf{n}(s,\tau)$ is the normal to $S(\tau)$ at $\mathbf{x}(s,\tau)$ and $\mathbf{u}(s,\tau)$ is its velocity. The small constant ϵ is used to smooth the expression and the time integration is over the entire history of particle overlap.

For a fixed τ, the set of points such that $\tau_I(s) \leq \tau$ defines a set of boundary segments. Let $s_1(\tau)$ and $s_2(\tau)$ be the endpoints of one such segment at time τ. Using this notation, $\tau(s) \leq \tau$ is equivalent to $s_1(\tau) \leq s \leq s_2(\tau)$. We can exchange the order of integration and rewrite (4.11) as

$$\Delta V^C = -\int_{t_0}^{t_1} \int_{s_1(\tau)}^{s_2(\tau)} \sqrt{\epsilon^2 + (\mathbf{u}(s,\tau)\cdot\mathbf{n}(s,\tau))^2}\, \mathrm{d}s\, \mathrm{d}\tau.$$

To first-order in Δt this is equivalent to

$$\Delta V^C = -\int_{s_1(t_0)}^{s_2(t_0)} \sqrt{\epsilon^2 + (\mathbf{u}(s,t_0)\cdot\mathbf{n}(s,t_0))^2}\, \mathrm{d}s\Delta t,$$

and the rate of change of V^C with respect to time is $\Delta V^C/\Delta t$. This lets us define

$$V(\mathbf{u},t_0) = -\int_{s_1(t_0)}^{s_1(t_0)} \sqrt{\epsilon^2 + (\mathbf{u}(s,t_0)\cdot\mathbf{n}(s,t))^2}\, \mathrm{d}s + \epsilon. \tag{4.12}$$

We will use (4.12) as a constraint in the next section. Note that we have added ϵ to make sure $V(\mathbf{u},t_0)$ can be zero. The variation of $V(\mathbf{u},t_0)$ with respect to $\mathbf{u}$ is

$$\mathrm{d}_{\mathbf{u}}V[\delta\mathbf{u}] = \frac{\mathrm{d}}{\mathrm{d}h}V(\mathbf{u}+h\delta\mathbf{u},t_0)\Big|_{h=0} = -\int_{s_1(t_0)}^{s_2(t_0)} \frac{(\mathbf{u}\cdot\mathbf{n})(\mathbf{n}\cdot\delta\mathbf{u})}{\sqrt{\epsilon^2+(\mathbf{u}\cdot\mathbf{n})^2}}\, \mathrm{d}s. \tag{4.13}$$

Instead of just preventing overlaps between rigid bodies, we can control the minimum distance between them. Defining $d_m \geq 0$, we modify the computation of τ_I to be the time of contact of the displaced surfaces $S(\tau) + d_m\mathbf{n}(\tau)$. Keeping bodies sufficiently separated means that we can limit potentially expensive near singular integration and control the stiffness of the problem.

4.5 The Boundary Integral Formulation as a NCP

At each time step the cannonical equations (2.31) and (2.32) are solved. We write these equations in the compact form

$$\mathbf{A}\boldsymbol{\Psi} = \mathbf{b},$$

where $\boldsymbol{\Psi}$ is a vector consisting of the density function $\boldsymbol{\eta}$ as well as the translational and rotational velocities of all the bodies and $\mathbf{b}$ is a vector consisting of the velocities on the surfaces of the bodies, as well as the net force and torque on each body. Modifying the right hand allows us to introduce contact forces between bodies,

$$\mathbf{A}\boldsymbol{\Psi} = \mathbf{b} + \mathbf{G}\hat{\mathbf{f}}, \tag{4.14}$$

where $\hat{\mathbf{f}}$ is a vector containing the forces and torques on each rigid body, and $\mathbf{G} \in \mathbb{R}^{2N+3n} \times \mathbb{R}^{3n}$ maps the contact forces to their corresponding velocities. Defining $\boldsymbol{\lambda} \in \mathbb{R}^m$ to be a scaling factor associated with each possible collision, and $\hat{\mathbf{F}}_c \in \mathbb{R}^{3n_p} \times \mathbb{R}^m$ to be the repulsion forces associated with each collision region, we have

$$\hat{\mathbf{f}} = \hat{\mathbf{F}}_c \cdot \boldsymbol{\lambda} = \begin{pmatrix} \hat{\mathbf{f}}_1^1 & \dots & \hat{\mathbf{f}}_1^m \\ \vdots & \ddots & \vdots \\ \hat{\mathbf{f}}_{n_p}^1 & \dots & \hat{\mathbf{f}}_{n_p}^m \end{pmatrix} \begin{pmatrix} \lambda_1 \\ \vdots \\ \lambda_m \end{pmatrix} = \begin{pmatrix} \lambda_1 \hat{\mathbf{f}}_1^1 + \ldots + \lambda_m \hat{\mathbf{f}}_1^m \\ \vdots \\ \lambda_1 \hat{\mathbf{f}}_{n_p}^1 + \ldots + \lambda_m \hat{\mathbf{f}}_{n_p}^m \end{pmatrix},$$

where $\hat{\mathbf{f}}_i^q$ are the forces and torques on rigid body i due to contact region q. Equation (4.10) tells us that $\hat{\mathbf{f}}_i$ is related to $\mathrm{d}_{\mathbf{u}}\mathbf{V}$, specifically $\hat{\mathbf{f}}_i^q = (\mathbf{F}_i^q, L_i^q)$ where,

$$\mathbf{F}_i^q = \int_{\Gamma_i^q} \mathrm{d}_{\mathbf{u}}\mathbf{V} \, \mathrm{d}S, \qquad L_i^q = \int_{\Gamma_i^q} (\mathrm{d}_{\mathbf{u}}\mathbf{V}) \cdot (\mathbf{x} - \mathbf{c}_q^p)^{\perp} \, \mathrm{d}S, \tag{4.15}$$

and Γ_i^q is the surface of part of rigid body i that belongs to contact region q.

We require $\boldsymbol{\lambda} \geq 0$ so that overlapping bodies repel one another. As discussed in the introduction to this section, we also require that if there is no contact, then the magnitude of the repulsion force is zero. Thus we have the complementarity problem

$$0 \leq \mathbf{V}(\boldsymbol{\Psi}) \perp \boldsymbol{\lambda} \geq 0. \tag{4.16}$$

Performing a first-order linearization

$$0 \leq \mathbf{V}(\Psi^k) + \mathrm{d}_{\boldsymbol{\Psi}}\mathbf{V}(\boldsymbol{\Psi}^k) \cdot \Delta\boldsymbol{\Psi} \perp \boldsymbol{\lambda}^{k+1} \geq 0.$$

From (4.14), we know that

$$\Delta\boldsymbol{\Psi} = \mathbf{A}^{-1} \cdot \mathbf{G} \cdot \hat{\mathbf{F}}_c \cdot \boldsymbol{\lambda}^{k+1}.$$

Thus we have the LCP

$$0 \leq \mathbf{V}(\boldsymbol{\Psi}^k) + \mathbf{B}\boldsymbol{\lambda}^{k+1} \perp \boldsymbol{\lambda}^{k+1} \geq 0, \tag{4.17}$$

where $\mathbf{B} = \mathrm{d}_{\boldsymbol{\Psi}}\mathbf{V} \cdot \mathbf{A}^{-1} \cdot \mathbf{G} \cdot \hat{\mathbf{F}}_c$. Algorithm 1 presents the algorithm that incorporates repulsion forces to advance to a contact-free time step. The matrix $\mathbf{B}$ is an $m \times m$ matrix, where the entry B_{ij} is the change induced by the i^{th} contact volume by the j^{th} contact force. This matrix is sparse and the largest entries will be along the diagonal, since a repulsion force will affect mainly its corresponding volume. Figure 4.2 provides a sketch of this algorithm applied to circles in an extensional flow.

Algorithm 1: Algorithm to advance simulation to a contact-free configuration.

Data: initial positions $\hat{\mathbf{q}}^0$

Result: contact-free configuration $\hat{\mathbf{q}}^1$

Solve $\mathbf{A\Psi} = \mathbf{b}$ for $\hat{\mathbf{u}}$;

$\hat{\mathbf{q}}^1 \leftarrow \hat{\mathbf{q}}^0 + \Delta t \hat{\mathbf{u}}$;

Compute $\mathbf{V}(\hat{\mathbf{q}}^1)$;

/* NCP iteration */

while $\mathbf{V} < \mathbf{0}$ **do**

  Compute $\mathrm{d}_{\mathbf{u}}\mathbf{V}$, $\hat{\mathbf{F}}$;

  $\mathbf{B} \leftarrow \mathrm{d}_{\mathbf{\Psi}}\mathbf{V} \cdot \mathbf{A}^{-1} \cdot \mathbf{G} \cdot \hat{\mathbf{F}}$; // Construct B

  Solve LCP: $0 \leq \mathbf{V} + \mathbf{B}\boldsymbol{\lambda} \perp \boldsymbol{\lambda} \geq 0$ for $\boldsymbol{\lambda}$;

  $\mathbf{b} \leftarrow \mathbf{b} + \mathbf{G}\hat{\mathbf{F}}_c\boldsymbol{\lambda}$;

  Solve $\mathbf{A\Psi} = \mathbf{b}$ for $\mathbf{\Psi}$; // Stokes solve

  $\hat{\mathbf{q}}^1 \leftarrow \hat{\mathbf{q}}^0 + \Delta t \hat{\mathbf{u}}$;

  Compute $\mathbf{V}(\hat{\mathbf{q}}^1)$;

Each time step requires the solution to a sequence of LCPs. Experience has shown that the solution often converges to the solution of the NCP after only one or two iterations. However, in certain situations the sequence of LCPs can be very lengthy. Since each LCP solve corresponds to solving the Stokes equations with a different contact force, limiting the number of required LCP solves is critical to keeping the computational cost reasonable. An approach that we have used is a heuristic adaptive time stepping routine. If the number of LCP iterations goes above a critical value, we halve the time step size. Conversely, if the number of LCP iterations is small (less than five for example) we increase the time step size slightly.

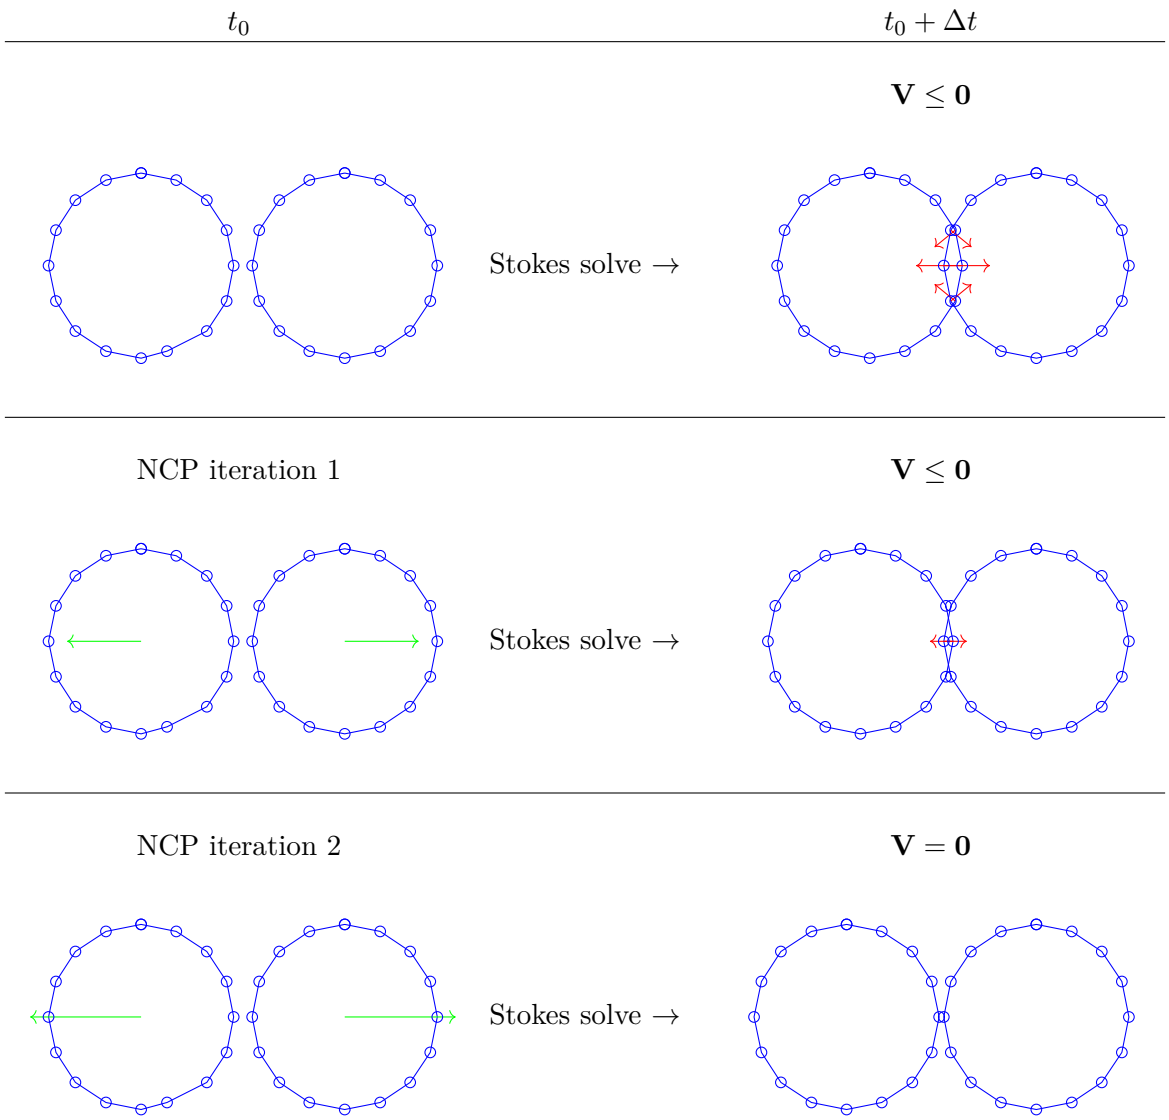

Figure 4.2: Consider two circular bodies in the extensional flow $(-x, y)$. Starting from an initial contact-free configuration, the Stokes equations are solved. After time stepping, the circles are pushed into one another and overlap occurs. A linear interpolation of the bodies' tracjectories are used to estimate $\mathbf{V}$ and $\mathrm{d}_{\mathbf{u}}\mathbf{V}$. From $\mathrm{d}_{\mathbf{u}}\mathbf{V}$ (red arrows), we compute the direction of the net force and torque on each body. Solving the LCP (4.17) gives the magnitude $\boldsymbol{\lambda}^0$. This repulsion force updates the right hand side $\mathbf{b}$ and the Stokes equations are solved again. In this example, this repulsion force (green arrow) is not enough to prevent overlap, and the LCP (4.17) is solved again for $\boldsymbol{\lambda}^1$. This new repulsion force updates $\mathbf{b}$ and the Stokes equations are solved again. At this point the circles do not overlap and the configuration is accepted and we advance to the next time step.

CHAPTER 5

RESULTS

We use our new time stepping method to simulate bounded and unbounded suspensions of two-dimensional rigid bodies in a viscous fluid. The main parameters are the minimum separation distance d_m, the number of discretization points of each rigid body, N_p, and each solid wall, N_w, and the initial time step size Δt. We perform convergence studies and investigate the effect of the STIV algorithm on the reversibility of the flow. To further demonstrate the consequence of STIV, we include plots of streamlines that cross whenever the collision detection algorithm is applied. The particular experiments we perform are now summarized.

- **Shear Flow**: We consider the standard problem of two identical rigid circles in the shear flow $\mathbf{u} = (y, 0)$ with the left body slightly elevated from the right body. We report similar results to those presented in [51], but we are able to take smaller initial displacements and minimum separation distances. The contact algorithm breaks the reversibility of the flow, and this effect is illustrated and quantified.

- **Taylor-Green Flow**: We simulate a concentrated suspension of 48 rigid ellipses in an unbounded Taylor-Green flow. At the prescribed separation distance, our new time stepping method is able to stably reach the time horizon, while the locally semi-implicit time integrator proposed by Lu et al. [51] results in the STIV algorithm stalling, even with $\Delta t = 10^{-8}$.

- **Porous Monolayer Injection**: We consider a suspension of confined rigid circular bodies motivated by an experiment by MacMinn et al. [53]. The geometry is an annulus with an inflow at the inner boundary and an outflow at the outer boundary. We again examine the effect of contact force on the reversibility of the flow, and compute the shear strain rate and make qualitative comparisons to results for deformable bodies [53].

- **Taylor-Couette Flow**: With the ability to do high area fraction suspensions without imposing a large non-physical minimum separation distance, we simulate rigid bodies of varying aspect ratios inside a Taylor-Couette device. We examine the effect of the rigid body shape and area fraction on the effective viscosity and the alignment angles.

5.1 Shear Flow

We consider two rigid circular bodies in the shear flow $\mathbf{u}(\mathbf{x}) = (y, 0)$. One body is centered at the origin, while the other body is placed to the left and above of the origin. With this initial condition, the particles come together, interact, and then separate. Both bodies are discretized with $N = 32$ points and the arc length spacing $h = 2\pi/32 \approx 0.196$. This experiment was also performed by Lu et al. [51], and we compare the two time stepping methods.

We start by considering the time step size $\Delta t = 0.4$ and minimum separation distance $\delta = 0$ (no contact algorithm). Our new globally implicit method successfully reaches the time horizon without requiring a repulsion force. However, with the same Δt, the local explicit time stepping results in a collision between the bodies, so the collision algorithm is required to reach the time horizon. Alternatively, the time step size can be reduced, but, as we will see, for sufficiently dense suspensions, even an excessively small time step size results in collisions. Next, in Figure 5.1, we investigate the effect of the minimum separation distance on the position of the rigid bodies. The top plot shows the trajectory of the left body as it approaches, interacts, and finally separates from the body centered at the origin. In this simulation, we use our new globally implicit time integrator, but the STIV contact algorithm is not applied. The bottom left plot shows the trajectory of the particle when the contact algorithm is applied with varying levels of separation. Notice that the trajectories are identical until near $x = 0$ when the particle separation first falls below the minimum separation distance. Finally, in the bottom right plot, the final vertical displacement of the body initially on the left is plotted. These results are computed for the locally implicit time stepping method [51], and the general trend of the trajectories are similar.

We next investigate the effect of the collision algorithm on the time reversibility of the flow. We reverse the shear direction at $t = 10$ and measure the error between the body's center at $t = 0$ and $t = 20$. We expect an error that is the sum of a first-order error caused by time stepping, and a fixed constant caused by the minimum separation distance. The results for various values of d_m are reported in Table 5.1. The contact algorithm is not applied when $d_m = 0$ and $d_m = h$, and we observe the expected first-order convergence. When $d_m \geq 2h$, the bodies are deflected onto contact-free streamlines when their proximity reaches the minimum separation distance. After the flow is reversed, the bodies again pass one another, but they are now on contact-free streamlines, so the initial deflection is not reversed. For these larger values of d_m, we see in Table 5.1 that the

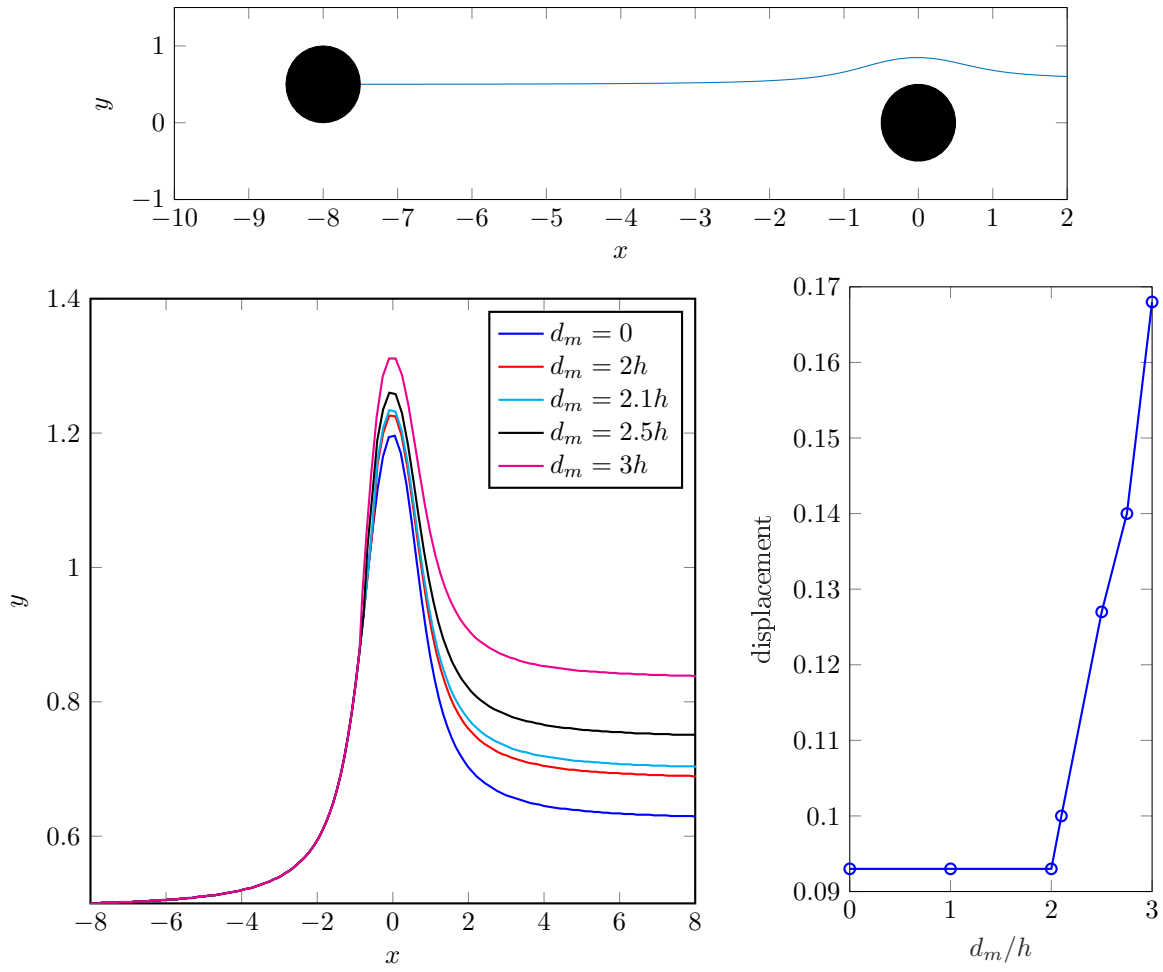

Figure 5.1: Shear experiment. Top: The initial setup and trajectory of the left body. Bottom left: The left body's trajectory for varying minimum separation distances. Notice how the trajectories are identical until shortly before $x = 0$ when the contact algorithm is first applied. Bottom right: The final vertical displacement of the left particle for varying minimum separation distances.

error eventually plateaus as Δt is decreased indiciating that the error due to the contact algorithm dominates.

The break in reversibility is further demonstrated by examining individual streamlines. In Figure 5.2, we compute the streamline of the left body for three different initial placements. We set $d_m = 3h$ for all the streamlines. With this threshold, only the bottom-most streamline falls below d_m. Therefore, as the bodies approach, the streamlines behave as expected—they do not cross.

Table 5.1: A study of time reversibility of the shear flow example. At $t = 10$, the flow direction is reversed and we calculate the relative error in the initial and final positions. When the collision constraint is active and force is needed to keep the bodies apart the contact algorithm dominates the error in the reversibility.

	Δt				
d_m	4×10^{-2}	2×10^{-2}	1×10^{-2}	5×10^{-3}	2.5×10^{-3}
0	1.35×10^{-1}	7.32×10^{-2}	3.74×10^{-2}	2.00×10^{-2}	1.01×10^{-2}
h	1.35×10^{-1}	7.32×10^{-2}	3.74×10^{-2}	2.00×10^{-2}	1.01×10^{-2}
$2h$	1.88×10^{-1}	1.41×10^{-1}	1.17×10^{-1}	1.08×10^{-1}	1.02×10^{-1}
$2.25h$	2.55×10^{-1}	2.08×10^{-1}	1.87×10^{-1}	1.78×10^{-1}	1.73×10^{-1}
$2.50h$	3.05×10^{-1}	2.69×10^{-1}	2.52×10^{-1}	2.45×10^{-1}	2.40×10^{-1}
$2.75h$	3.64×10^{-1}	3.31×10^{-1}	3.13×10^{-1}	3.07×10^{-1}	3.03×10^{-1}
$3.00h$	4.12×10^{-1}	3.88×10^{-1}	3.72×10^{-1}	3.67×10^{-1}	3.63×10^{-1}

However, when the contact algorithm is applied to the blue streamline, the streamlines cross.

5.2 Taylor-Green

For planar flows, we can separate suspensions into dilute and concentrated regimes by comparing the number of bodies per unit area, ν, to the average body length ℓ. If $\nu < 1/\ell^2$, the suspension is dilute, otherwise it is concentrated (in 2D planar suspensions, unlike 3D suspensions, there is no semi-dilute regime). We consider the suspension of 75 rigid bodies in the Taylor-Green flow $\mathbf{u}^\infty = (\cos(x)\sin(y), -\sin(x)\cos(y))$. The number of bodies per unit area is $\nu \approx 3.1$ which is greater than $1/\ell^2 = 1.1$. Therefore, this suspension is well within the concentrated regime.

We discretize the bodies with $N = 32$ points and select the minimum separation distance $d_m = 0.05h$. Snapshots of the simulation are shown in Figure 5.3. In this concentrated suspension, the bodies frequently come into contact. If the interactions between these nearly touching bodies are treated explicitly, this leads to stiffness. Our time stepper controls this stiffness by treating these interactions implicitly, and the simulation successfully reaches the time horizon. We performed the same simulation, but with the locally implicit time stepping method [51]. Because of the near-contact, smaller time step sizes must be taken. We took time step sizes as small as 10^{-8}, and the method was not able to successfully reach the time horizon. This exact behavior has also been observed for vesicle suspensions [63]. In the bottom right plot of Figure 5.3, we show the trajectory

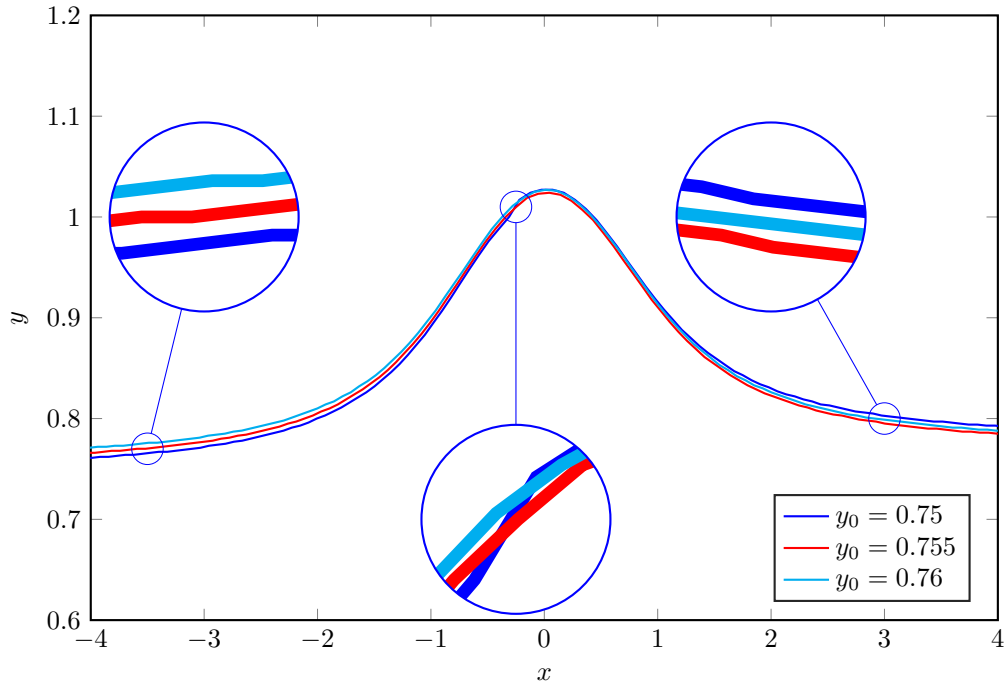

Figure 5.2: The contact algorithm causes streamlines to cross. Keeping the minimum separation fixed at $d_m = 3h$, we vary the starting y location of the left body. The teal and red streamlines do not require a repulsion force to enforce the minimum separation between the bodies, but the blue streamline does. Once the contact algorithm is applied, the blue streamline crosses the other streamlines (middle inset). This crossing of the streamlines breaks the reversibility of the simulation.

of one body for different time step sizes. The dots denote locations where the contact algorithm is applied. For this very complex flow, the trajectories are in good agreement with different time step sizes.

5.3 Fluid Driven Deformation

A recent experiment considers a dense monolayer packing of soft deformable bodies [53]. Motivated by this experiment, we perform numerical simulations of rigid bodies in a similar device. We pack rigid bodies in a Couette device, but with a very small inner boundary. The boundary conditions are an inflow or outflow of rate Q at the inner boundary with an outflow or inflow at the outer cylinder. This boundary condition corresponds to injection and suction of fluid from the center of the experimental microfluidic device. In the experimental setting, the soft bodies are able

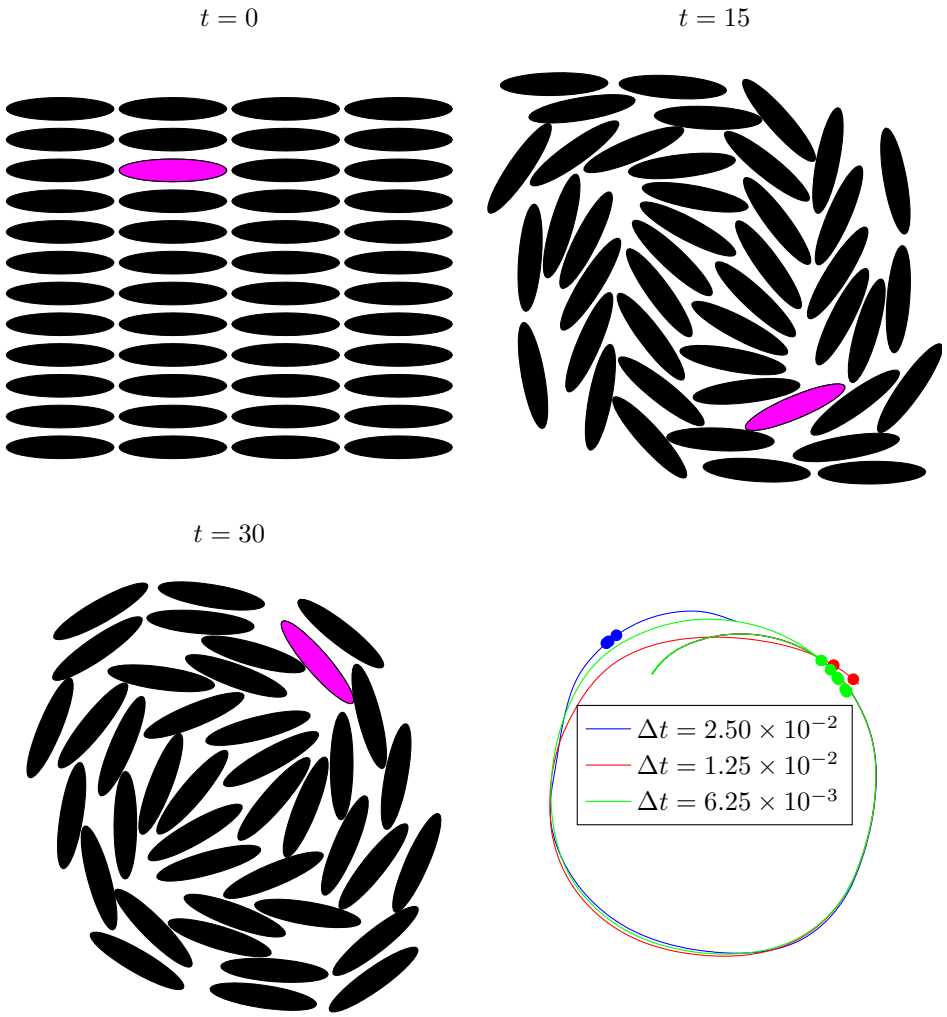

Figure 5.3: Snapshots of a dense suspension in an unbounded Taylor-Green flow. The number of bodies per unit area ν is approximately 3.1. This is greater than $1/\ell^2 = 1.1$, which puts the simulation well within the concentrated regime. Bodies are discretized with 32 points and the minimum separation is $\delta = 0.05h$. The bottom right plot shows the trajectory of the center of the colored body for different step sizes. Each line in that plot is marked where a contact force is applied to the colored body to enforce the minimum separation.

to reach the outer boundary, and the resulting boundary condition would not be uniform at the outer wall. So that we can apply the much simpler uniform inflow or outflow at the outer boundary, we force the rigid bodies to remain well-separated from the outer wall. We accomplish this by placing a ring of *fixed* rigid bodies halfway between the inner and outer cylinders (Figure 5.4). The

spacing between these fixed bodies is sufficiently small that the mobile bodies are not able to pass. Since the outer boundary is well-separated from the fixed bodies, the outer boundary condition is justifiably approximated with a uniform flow.

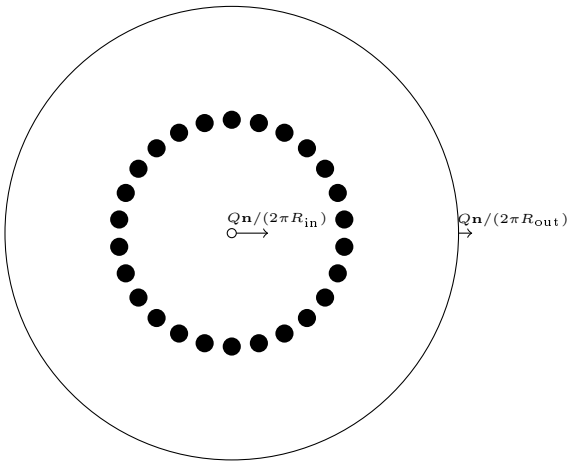

Figure 5.4: The geometry used in our numerical experiment that is motivated by the experimental setup of MacMinn et al. [53]. The fixed solid bodies are shaded in black.

We start by examining the effect of the contact algorithm on the reversibility of the flow. We again reverse the flow at time T and run the simulation until time $2T$. The rigid bodies are in contact for much longer than the shear example in Section 5.1, so maintaining reversibility is much more challenging. Figure 5.5 shows several snapshots of the simulation, and the bottom right plot superimposes the initial and final configurations. We observe only a minor violation of reversibility, and it is largest for bodies that were initially near the fixed bodies—the contact algorithm is applied to these bodies most frequently.

In [53], the shear strain rate is measured to better characterize the flow. In Figure 5.6, we plot the shear strain rate for the simulation in Figure 5.5. A qualitative comparison of the numerical and experimental results are in good agreement. In particular, the petal-like patterns in Figure 5.6 are also observed in the experimental results.

5.4 Taylor-Couette

In many industrial applications, for example pulp and paper manufacturing, suspensions of rigid elongated fibers are encountered. Motivated by these suspensions we investigate rheological

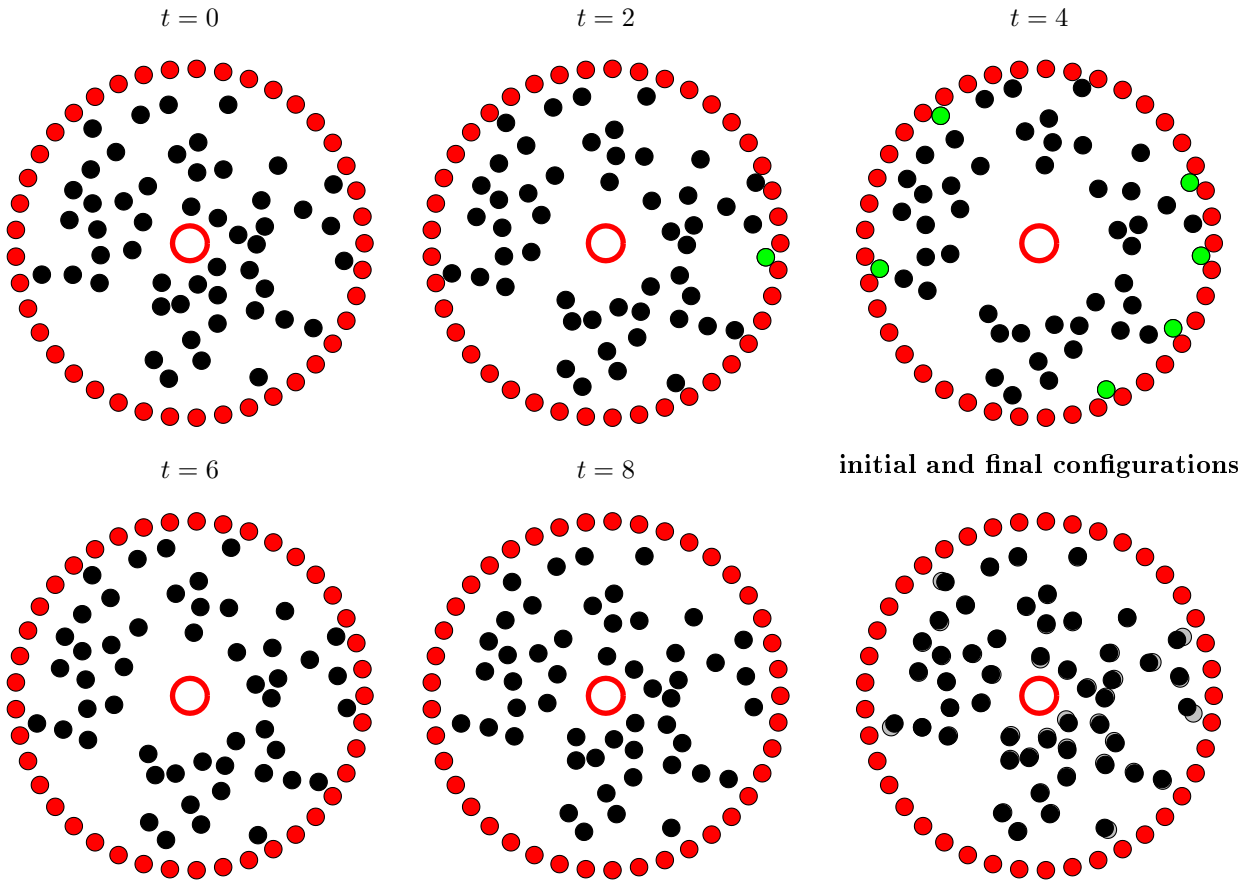

Figure 5.5: Snapshots of a rigid body suspension motivated by an experiment for deformable bodies [53]. Fluid is injected at a constant rate starting at $t = 0$. At $t = 4$ the flow direction is reversed. Fixed bodies are colored in red, while bodies subject to a repulsion force are colored in green. The initial configuration has been superimposed on the final configuration at $t = 8$ to show the effect of the repulsion forces on reversibility.

and statistical properties of confined suspensions. We consider suspensions of varying area fraction and body aspect ratio; specifically we consider 5, 10, and 15 percent area fractions and elliptical bodies of aspect ratio, λ of 1, 3, and 6. In all the examples, $\nu < 1/\ell^2$, so all the suspensions are in the dilute regime. The bodies initial locations are random, but non-overlapping (Figure 5.7). The flow is driven by rotating the outer cylinder at a constant angular velocity while the inner cylinder remains fixed.

Before measuring any rheological properties, the outer wall completes one full revolution so that the bodies are well-mixed and approaching a statistical equilibrium. We start by considering the

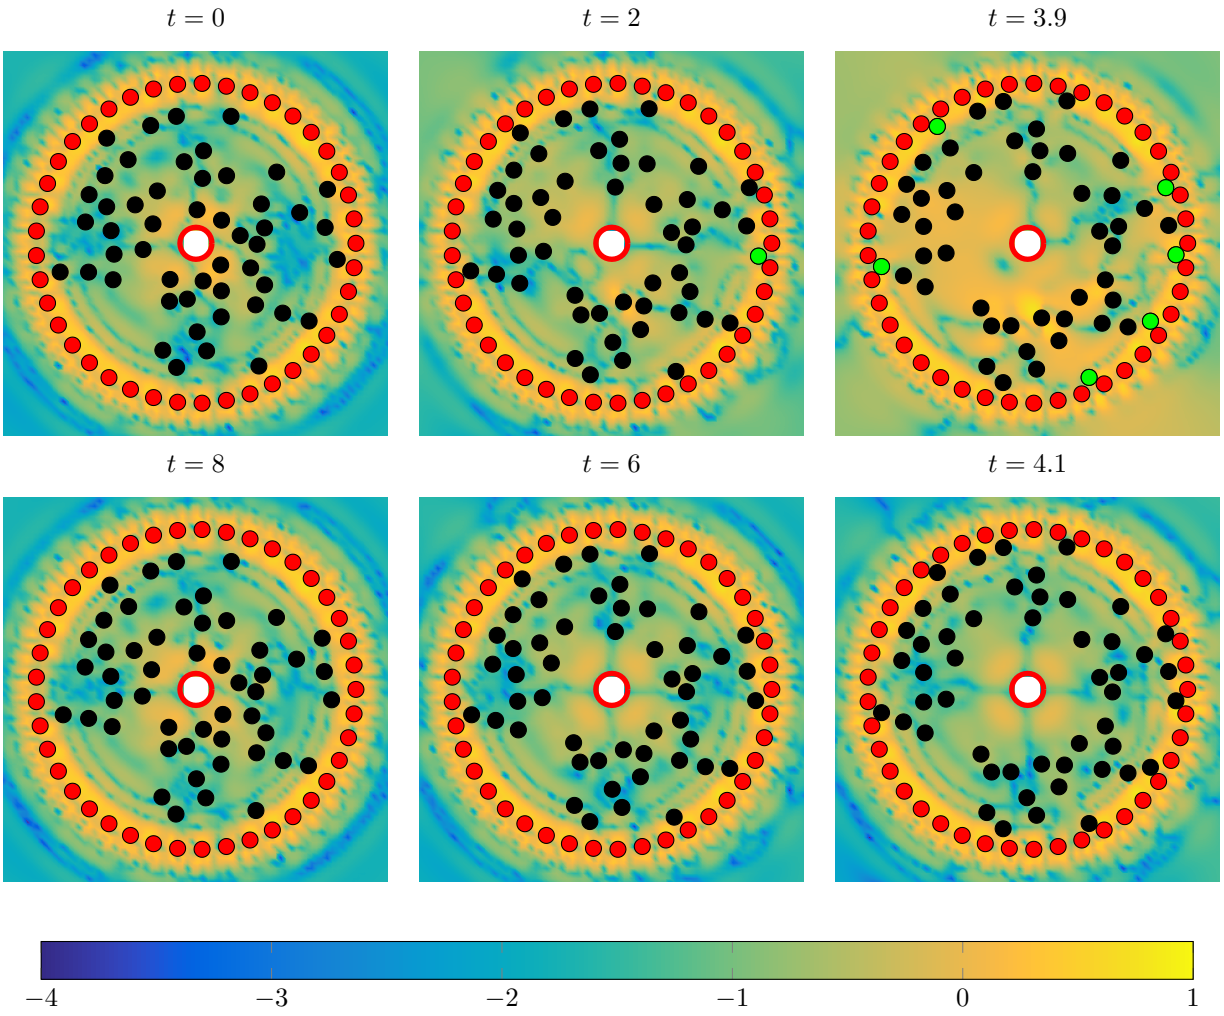

Figure 5.6: The shear strain rate $\log_{10}(|\sigma_{xy}|)$ of the suspension in Figure 5.5. The formation of the petal-like patterns is also observed experimentally for a suspension of deformable bodies [53].

alignment of the bodies. The alignment is particularly insightful since many industrial processes involve fibers suspended in a flow, and the alignment affects the material properties [45]. One way to measure the alignment is the order parameter, S, defined as,

$$S = \left\langle \frac{d\cos^2\tilde{\theta} - 1}{d - 1} \right\rangle,$$

where d is the dimension of the problem (2 in our case), $\tilde{\theta}$ is the deviation from the expected angle, and $\langle\cdot\rangle$ averages over all bodies. If $S = 1$, all bodies are perfectly aligned with the shear direction,

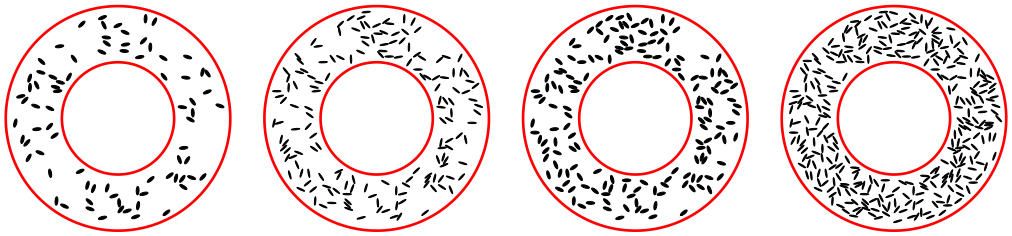

Figure 5.7: Four initial configurations for Taylor-Couette flow with varying volume fraction ϕ and aspect ratio λ. From left to right: 1) $\phi = 5\%$, $\lambda = 3$, 2) $\phi = 5\%$, $\lambda = 6$, 3) $\phi = 10\%$, $\lambda = 3$, 4) $\phi = 10\%$, $\lambda = 6$.

$S = 0$ corresponds to a random suspension (no alignment), and $S = -1$ means that all bodies are perfectly aligned perpendicular to the shear direction. In our geometric setup, a body centered at (x, y) has an expected angle of $\tilde{\theta} \tan^{-1}(y/x) + \pi/2$, and the average alignment of the bodies will be in the direction of the shear, which is also perpendicular to the radial direction.

Since the initial condition is random, the initial configurations in Figure 5.7 have an order parameter $S \approx 0$. As the outer cylinder rotates, we see in Figure 5.8 that S increases rapidly. The area fractions ϕ we consider have a minor effect on S; however, the aspect ratio has a large effect. In particular, suspensions with slender bodies align much better with the flow.

This matches the known dynamics of a single body in an unbounded shear flow, where the body will align with the shear direction on average. Bodies with a high aspect ratio rotate quickly when then they are perpendicular to the shear direction and spend more time nearly aligned with the shear direction. We compare our results to the time averaged order parameter of a single elliptical body in an unbounded shear flow. If the shear rate is $\dot{\gamma}$, a single elliptical body rotates with period $\tau = \pi/(2|\dot{\gamma}|)(\lambda + \lambda^{-1})$ [35] according to

$$\varphi(t) \;=\; \tan^{-1}\left(\frac{1}{\lambda}\tan\left(\frac{\lambda\dot{\gamma}t}{\lambda^2+1}\right)\right).$$

The time average order parameter is then,

$$\langle S\rangle \;=\; \frac{1}{\tau}\int_0^\tau \left(2\cos^2(\varphi(t)) - 1\right)\,\mathrm{d}t \;=\; \frac{\lambda-1}{\lambda+1}.$$

Independent of the shear rate, for $\lambda = 3$ the theoretical $\langle S\rangle$ is $1/2$ and for $\lambda = 6$ it is $5/7$. Table 5.2 shows the time and space averaged order parameter for the Couette apparatus. We see that in all cases our computed time averaged order parameter is higher than the theoretical single fiber case.

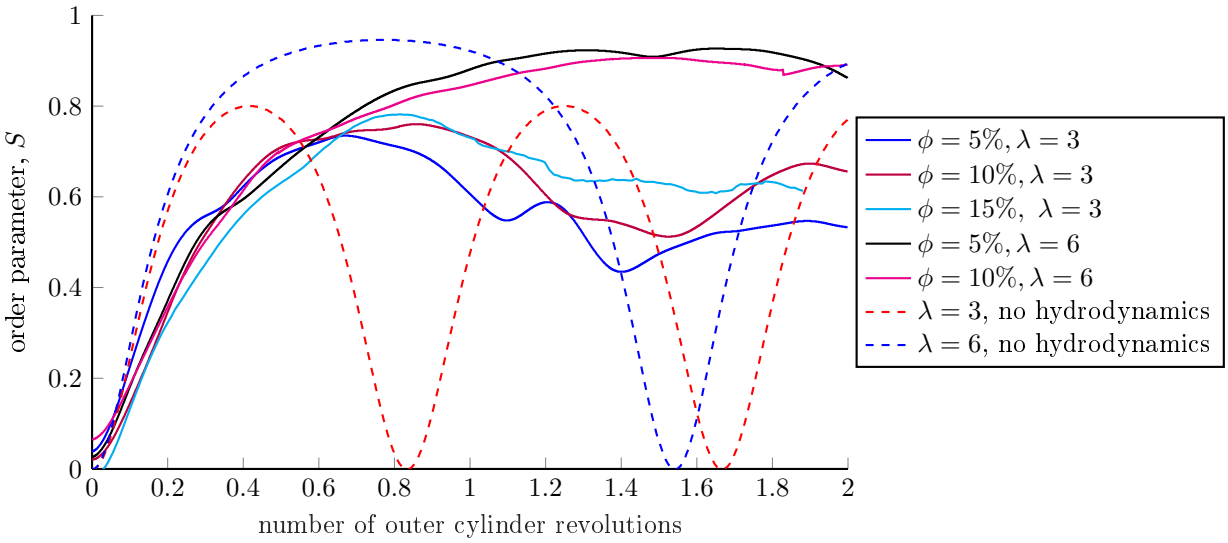

Figure 5.8: The order parameter of different fiber concentrations and aspect ratios. We see that the 6:1 fibers align better. The 6:1 fibers rotate through the angle perpendicular to the shear direction more quickly than the 3:1 fibers and thus spend more time approximately aligned with the shear direction. The dashed lines represent the order parameter for a suspension in an unbounded shear flow with bodies that do not interact hydrodynamically. The red line shows $\lambda = 3$, while the blue line shows $\lambda = 6$.

This could be due to the hydrodynamic interactions between the bodies, or the effect of the solid walls.

In the absence of solid walls and hydrodynamic interactions between bodies, a suspension will align and disalign. The period of the order parameter in this case is the same as the rotational period for a single fiber. In Figure 5.8 the theoretical order parameter is shown for a suspension of non-hydrodynamically interacting fibers in an unbounded shear flow. Hydrodynamic interactions prevent the suspension from disaligning completely.

Another quantity of interest in rheology is the effective viscosity of a suspension. The shear viscosity μ relates the bulk shear stress σ_{xy} of a Newtonian fluid to the bulk shear rate $\dot{\gamma}$,

$$\sigma_{xy} = \mu\dot{\gamma}.$$

Adding bodies increases the bulk shear stress of a suspension. The proportionality constant relating the increased σ_{xy} to the shear rate is the *apparent viscosity*, and the ratio between the apparent viscosity and the bulk viscosity is the effective viscosity μ_{eff}. Experimentally, the bulk shear stress

Table 5.2: The time averaged order parameter during the second revolution of the Couette apparatus. The higher aspect ratio fibers align better on average. The alignment is in all cases higher than predicted for a single Jeffery orbit.

area fraction, ϕ	aspect ratio, λ	computed $\langle S \rangle$	theoretical $\langle S \rangle$ (single fiber)
5%	3	0.52	0.50
10%	3	0.60	0.50
15%	3	0.65	0.50
5%	6	0.91	0.71
10%	6	0.89	0.71

is often computed by measuring the torque on the inner cylinder [40]. Numerically, this is simply the strength of the rotlet corresponding to the inner cylinder. By computing the ratio of the torque on the inner cylinder with bodies to the torque without bodies we determine the effective viscosity of a suspension. Figure 5.9 shows the effective viscosity increases with ϕ, but is generally lower for bodies with aspect ratio $\lambda = 6$. This is because higher aspect ratio bodies align themselves better, and thus contribute less to the bulk shear stress. The spikes in 5.9 occur when a repulsion force is added to the system. Similar spikes were observed in Lu *et al.* [51]. To smooth the results we use a multiscale local polynomial transform to smooth the data shown in Figure 5.9.

Finally, instead of computing the instantaneous effective viscosity, experimenters are interested in the time averaged effective viscosity of a suspension. In Table 5.3, we report the average instantaneous effective viscosity over the second revolution of the outer cylinder

Table 5.3: Time averaged effective viscosity for various area fractions and aspect ratios. The time average is done between the first and second revolutions of the outer cylinder. As ϕ increases the effective viscosity increases as expected. In general higher aspect ratio bodies increase the viscosity less than lower aspect bodies.

λ	5% area fraction	10% area fraction	15% area fraction
1	1.12	1.22	1.42
3	1.10	1.23	1.36
6	1.08	1.18	

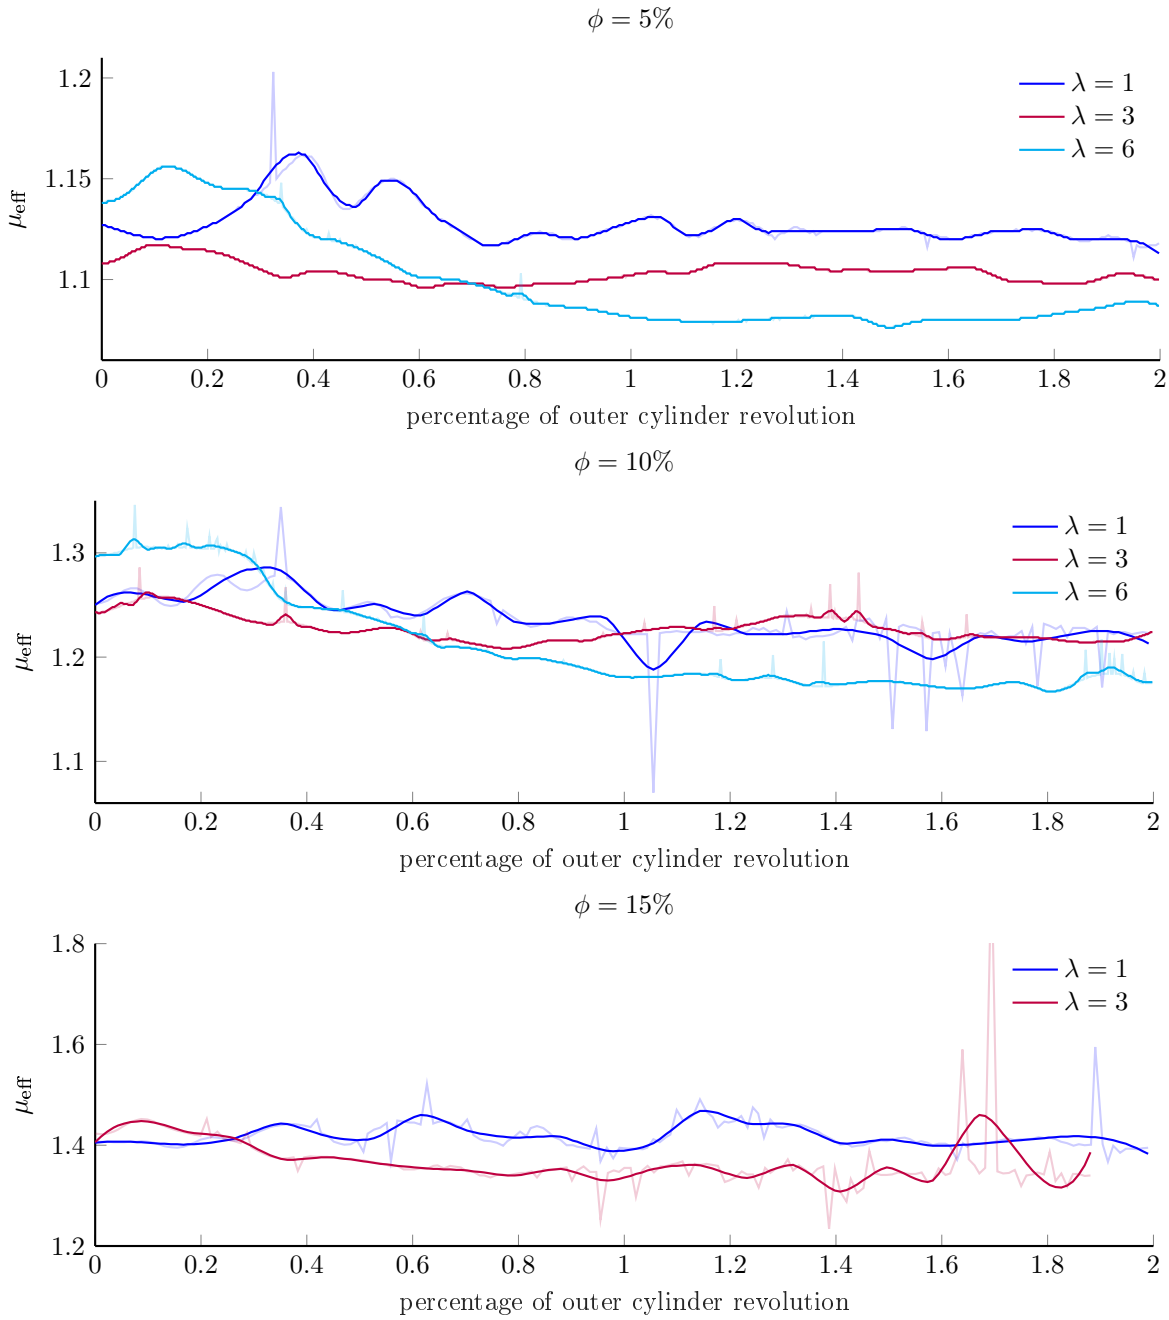

Figure 5.9: Instantaneous bulk effective viscosity for various volume fractions and body aspect ratios. The inner cylinder is fixed, while the outer one rotates at a constant angular velocity. The transparent lines represent the raw data, while the solid lines have been smoothed using a multiscale local polynomial transform.

CHAPTER 6

CONCLUSION

We have developed numerical methods to simulate the motion of two-dimensional rigid particles in a Stokesian fluid. The method makes no assumptions on the particle shape, nor on the particle concentration. Moreover, it works in bounded and unbounded flows. One challenge when simulating suspensions is that numerical errors can cause bodies to overlap. To avoid overlap we introduce a modification of the repulsion force in [51]. This repulsion force ensures separation between particles, and is completely free of any tuning parameters. In the original contact algorithm, all inter-body hydrodynamic interactions are discretized explicitly by lagging them from the previous time step. This results in a block-diagonal system to solve at each time step. Unfortunately, discretizing the inter-body interactions explicitly necessitates a small time size, particularly for concentrated suspensions. The approach taken by Lu *et al.* [51] was to maintain a sufficiently large minimum separation distance to control the stiffness. Instead, we have taken the approach of using a globally-implicit time stepper to simulate concentrated two-dimensional suspensions without requiring a large minimum separation distance nor an excessively small time step size. The disadvantage of the new globally-implicit time stepper is that a full dense linear system must be solved at every time step since all the bodies are coupled. In certain simulations, the additional cost of performing this dense matrix solve every time step is more than offset by the ability to take larger time steps. We use our stable algorithm to investigate the rheological properties of various suspensions and study the effect of the contact force on the time reversibility of the simulation. We investigate the effect of the concentration and aspect ratio of rigid fibers on the alignment angle and effective viscosity of a suspension confined inside a Couette apparatus.

6.1 Limitations

Physically, in order to use a boundary integral representation, the method requires that the steady Stokes equations are valid. Therefore, it is assumed that both the Reynolds number and the Strouhal number are very small. This means that the length scale is small, the velocity is slow,

the viscosity of the solvent is large, and the time scale is large. In addition to these assumptions, it is also necessary that the solvent is Newtonian. If any of these assumptions are not valid, then boundary integral equations become much more difficult or impossible to implement. It is possible to still use an integral equation formulation, however the formulation often includes computationally expensive volume integrals.

Another limitation is that the suspended bodies are assumed to be completely rigid. The formulation for non-rigid bodies is slightly different, and BIEs have been used to simulate different particle suspensions, including vesicles [63,64,68], drops [73], and flexible fibers [77]. The repulsion force described in Chapter 4 does not require the bodies to be rigid, and in the fact the original paper by Lu *et al.* [51] demonstrated the robustness of the method for vesicles suspensions.

6.2 Future Work

There are many avenues for future work. Mathematically, since the linear complementarity problems that need to be solved are not positive definite, there is no guarantee that solutions exist or are unique. This may lead to a non-convergent sequence of solutions to the nonlinear complementarity problem, and is believed to cause problems in simulations involving solid walls. We attempt to circumvent this issue by reducing the time step size if the number of NCP iterations is too large. However, this is a heuristic and contradicts our goal to avoid heuristics and tuning parameters. The robustness of the contact algorithm can be improved by using a different measure for contact. In [81], instead of the STIV, the signed distance between bodies is used to measure overlap. This method results in a sequence of symmetric, positive definite LCPs, where solutions are guaranteed to exist and be unique. This choice of metric, however, only detects contact at the end of each time step. Unlike STIVs, if contact occurs at an intermediate time (i.e. one body passes completely through another) this contact will not be detected.

6.2.1 Three-dimensional Simulations

The method as it stands has only been implemented for two dimensional (planar) flows. For a fixed number of unknowns, two-dimensional simulations allow us to simulate far more rigid bodies. That said, three-dimensional simulations are more realistic and would allow us to properly investigate certain physical features that only show up in three dimensions. The Stokes paradox,

that in two-dimensions prevents us from simulating particles undergoing a net force in an unbounded domain, does not exist in three dimensions where the fundamental solution decays as r^{-1}. This would allow us to simulate other important problems such as sedimentation, without introducing a bounding wall.

Three-dimensional computations involving boundary integrals to simulate rigid body suspensions are well- developed [17, 55]. After replacing the two-dimensional kernels with their three-dimensional counterparts, the formulation described in Chapter 2 is identical. The three-dimensional double-layer kernel is singular, so different quadrature techniques must be used to evaluate the intra-body interactions [12]. The space-time interference volume required to compute the repulsion force as discussed in Chapter 4 was originally developed in three dimensions [30]. In addition, the near singular integration technique described in Section 3.3 has been implemented in three dimensions [84] and other methods to compute the interactions between nearly touching bodies in three dimensions exist [3, 72].

6.2.2 Periodic Boundary Conditions

Although we have used our method to compute the viscosity of a suspension inside a Couette apparatus, in reality wall effects play a large role in the viscosity of a confined suspension. One remedy is to enlarge the size of the computational domain and compute the viscosity based on a representative region far from the walls. However, the discretization of the solid walls and the extra particles not in the representative region increases the size of the resulting linear system. To more accurately and efficiently compute quantities like the effective viscosity, periodic boundary conditions should be used. Single and double periodic boundary conditions been implemented in a boundary integral setting [2, 67]. Periodic fast summation methods like the Ewald summation [21] and the periodic FMM [67, 82] are used to accelerate the linear solve.

For unbounded shear flow however, the background flow is periodic in the shear direction but grows in the direction normal to the shear plane. The Lees-Edward boundary conditions [46] have been used in molecular dynamics simulations to model a sheared system without introducing any solid walls. These boundary conditions are illustrated in Figure 6.1.

As with regular periodic simulations, a master cell of dimensions $H \times H$ is replicated vertically and horizontally. The cells above and below the master cell translate with constant velocity $V = \dot{\gamma}\bar{y}$, where $\dot{\gamma}$ is the shear rate and $\bar{y}$ is the y component of the center of the cell. In the x direction

the periodicity is enforced as usual. That is, if the body leaves the master cell from one side it reappears on the other (green body in Figure 6.1). In the vertical direction however, we need to take into account the relative movement of the cells.

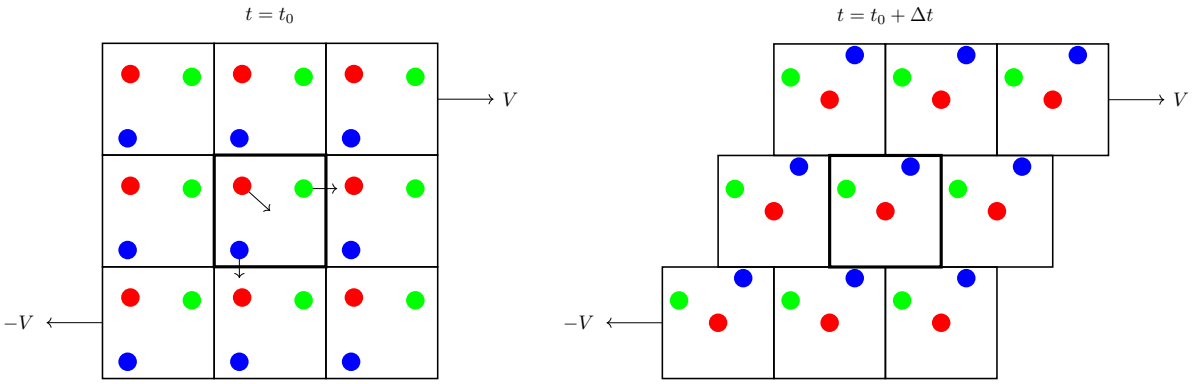

Figure 6.1: Sketch of the Lees-Edwards boundary condition for unbounded shear flow. A master cell is replicated vertically and horizontally. The cells above and below the master cell translate relative to it with velocity $V = \dot{\gamma}\bar{y}$, where $\dot{\gamma}$ is the shear rate and $\bar{y}$ is the y component of the center of the cell.

Consider the blue body in Figure 6.1 and Figure 6.2. By time $t + \Delta t$, the cell directly above the master cell at $t = t_0$ has shifted by $V\Delta t$ units. Therefore instead of the blue body exiting from the bottom and reentering at the top with the same x component, it must reenter from its horizontal position in the shifted cell. In addition, the velocity of the blue body changes to match the cell it is entering from. Instead of moving at velocity (u, v) it will now be moving with velocity $(u + V, v)$.

Thus the Lee-Edwards boundary conditions for a body attempting to move to (x, y) with velocity (u, v) can be summarized as follows:

- Body exits from the left edge of master cell:

$$(x, y) \to (x - H, y), \qquad (u, v) \to (u, v).$$

- Body exits from right edge of master cell:

$$(x, y) \to (x + H, y) \qquad (u, v) \to (u, v).$$

- Body exits from bottom of master cell:

$$(x, y) \to (x + V\Delta t, y + H), \qquad (u, v) \to (u + V, v).$$

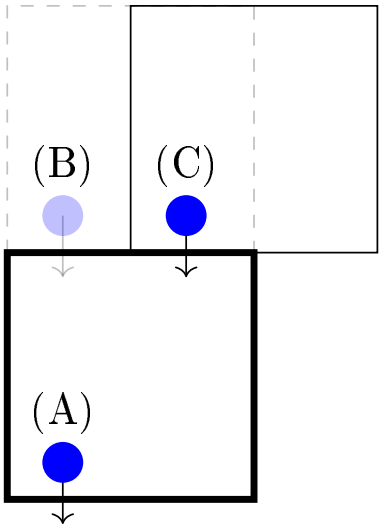

Figure 6.2: Sketch of just the blue body in Figure 6.1. The blue body exits the master cell from the bottom (A). Because the cell on top has shifted by $V\Delta t$, instead of reentering from the top (B) as in normal periodic boundary conditions, the body reenters from its horizontal location in the shifted cell (C).

- Body exits from top of master cell:

$$(x, y) \to (x - V\Delta t, y - H), \qquad (u, v) \to (u - V, v).$$

These boundary conditions have been used to simulate rigid body suspensions using the finite element method [33], the lattice Boltzmann method [50], and Stokesian dynamics [70], however they not been investigated using BIEs. To implement a repulsion force using the methods described in Chapter 4 would require a periodized computation of the STIV, something that is not yet implemented.

APPENDIX A

INDEX NOTATION

Index notation is a tool to facilitate concise manipulation of tensor operations. Details on the definitions and properties of tensors can be found in many engineering and physics textbooks. A vector in $\mathbb{R}^d$ is first-order tensor and can be represented in index notation as

$$\mathbf{a} = a_i,$$

where i ranges from 1 to d. In this dissertation, we will always take $d = 2$. A second-order tensor can be represented as

$$\mathbf{A} = A_{ij},$$

where now both i and j range from 1 to d. Higher order tensors can be defined in a similar manner.

A.1 Einstein Summation Convention

When manipulating tensor expressions, it is often necessary to perform summation operations across indices. For example the dot product between two vectors in is defined as

$$\mathbf{a} \cdot \mathbf{b} = a_1 b_1 + a_2 b_2 = \sum_{i=1}^{2} a_i b_i.$$

Einstein simplified this notation by adopting the convention that any repeated index in an expression always implies a summation. The dot product can then be written more concisely as

$$\mathbf{a} \cdot \mathbf{b} = a_i b_i.$$

This notation can also be used to define matrix-vector products and the divergence operator

$$\mathbf{A}\mathbf{x} = \mathbf{b} \Leftrightarrow A_{ij} x_i = b_j, \qquad \nabla \cdot \mathbf{a} = \frac{\partial}{x_i} a_i = \partial_i a_i.$$

Often we will use the shorthand ∂_i to represent $\partial / \partial x_i$. The gradient can be defined then as

$$\nabla \mathbf{a} = \partial_j a_i.$$

A.2 Special Tensors

There are a couple of special tensors that come up often in index notation. The first is the Kronecker-delta tensor, a second-order tensor defined as

$$\delta_{ij} = \begin{cases} 0 & i \neq j, \\ 1 & i = j. \end{cases}$$

This tensor is equivalent to the identity matrix in linear algebra.

The Levi-Civita tensor is third-order tensor defined by

$$\epsilon_{ijk} = \begin{cases} 1 & \text{if } (i,j,k) \text{ is } (1,2,3),\ (2,3,1),\ (3,1,2). \\ -1 & \text{if } (i,j,k) \text{ is } (3,2,1),\ (1,3,2),\ (2,1,3). \\ 0 & \text{otherwise.} \end{cases}$$

The Levi-Civita tensor appears in the cross product and curl:

$$\mathbf{a} \times \mathbf{b} = \epsilon_{ijk} a_j b_k, \qquad \nabla \times \mathbf{a} = \epsilon_{ijk} \partial_j a_k.$$

REFERENCES

[1] Suresh G. Advani and Charles L. Tucker III. The Use of Tensors to Describe and Predict Fiber Orientation in Short Fiber Composites. *Journal of Rheology*, 31(8):751–784, 1987.

[2] Ludvig af Klinteberg and Anna-Karin Tornberg. Fast Ewald summation for Stokesian particle suspensions. *International Journal for Numerical Methods in Fluids*, 76(10):669–698, 2014.

[3] Ludvig af Klinteberg and Anna-Karin Tornberg. A fast integral equation method for solid particles in viscous flow using quadrature by expansion. *Journal of Computational Physics.*, 326:420–445, 2016.

[4] M. Anitescu and F. A. Potra. Formulating dynamic multi-rigid-body contact problems with friction as solvable linear complementarity problems. *Nonlinear Dynamics*, 14(3):231–247, 1997.

[5] G. Ausias, X.J. Fan, and R.I. Tanner. Direct simulation for concentrated fibre suspensions in transient and steady state shear flows. *Journal of Non-Newtonian Fluid Mechanics*, 135(1):46–57, 2006.

[6] Alex Barnett, Bowei Wu, and Shravan Veerapaneni. Spectrally Accurate Quadratures for Evaluation of Layer Potentials Close to the Boundary for the 2D Stokes and Laplace Equations. *SIAM Journal on Scientific Computing*, 37(4):B519–B542, 2015.

[7] G. K. Batchelor. Slender-body theory for particles of arbitrary cross-section in Stokes flow. *Journal of Fluid Mechanics*, 44:419, 1970.

[8] G. K. Batchelor. The stress generated in a non-dilute suspension of elongated particles by pure straining motion. *Journal of Fluid Mechanics*, 46(4):813–829, 1971.

[9] Jan Bender, Kenny Erleben, and Jeff Trinkle. Interactive simulation of rigid body dynamics in computer graphics. *Computer Graphics Forum*, 33(1):246–270, 2012.

[10] Stephen C. Billups and Katta G. Murty. Complementarity problems. *Journal of Computational and Applied Mathematics*, 124(1):303 – 318, 2000.

[11] J. F. Brady and G. Bossis. Stokesian Dynamics. *Annual Review of Fluid Mechanics*, 20(1):111–157, 1988.

[12] James Bremer and Zydrunas Gimbutas. A Nyström method for weakly singular integral operators on surfaces. *Journal of Computational Physics*, 231(14):4885 – 4903, 2012.

[13] Lukas Bystricky, Sachin Shanbhag, and Bryan D. Quaife. Stable and contact-free time stepping for dense rigid particle suspensions. Preprint available at `https://arxiv.org/pdf/1804.11012.pdf`.

[14] S. L. Campbell, I. C. F. Ipsen, C. T. Kelley, C. D. Meyer, and Z. Q. Xue. Convergence estimates for solution of integral equations with GMRES. *Journal of Integral Equations and Applications*, 8(1):19–34, 1996.

[15] J. Carlson, J. A. Carlson, A. Jaffe, and A Wiles. *The Millennium Prize Problems*. Amsns AMS non-series Title Series. American Mathematical Society, 2006.

[16] I. L. Claeys and J. F. Brady. Suspensions of prolate spheroids in stokes flow. 1. dynamics of a finite number of particles in an unbounded fluid. *J. Fluid Mech.*, 68:411, 1993.

[17] Eduardo Corona, Leslie Greengard, Manas Rachh, and Shravan Veerapaneni. An integral equation formulation for rigid bodies in Stokes flow in three dimensions. *Journal of Computational Physics*, 332:504–519, 2017.

[18] Hua-Shu Dou, Boo Cheong Khoo, Nhan Phan-Thien, Khoon Seng Yeo, and Rong Zheng. Simulations of fibre orientation in dilute suspensions with front moving in the filling process of a rectangular channel using level-set method. *Rheologica Acta*, 46(4):427–447, 2007.

[19] Albert Einstein. *Eine neue Bestimmung der Moleküldimensionen*. PhD thesis, University of Zurich, 1906.

[20] Kenny Erleben. Numerical methods for linear complementarity problems in physics-based animation. In *ACM SIGGRAPH 2013 Courses*, SIGGRAPH '13, pages 8:1–8:42, New York, NY, USA, 2013. ACM.

[21] Peter P. Ewald. Die berechnung optischer und elektrostatischer gitterpotentiale. *Annalen der Physik*, 369(3):253–287, 1921.

[22] Xijun Fan, N. Phan-Thien, and Rong Zheng. A direct simulation of fibre suspensions. *Journal of Non-Newtonian Fluid Mechanics*, 74(1):113–135, 1998.

[23] M. C. Ferris and J. S. Pang. Engineering and economic applications of complementarity problems. *SIAM Review*, 39:669–713, 1995.

[24] A. Fischer. A special newton-type optimization method. *Optimization*, 24(3-4):269–284, 1992.

[25] Daniel Flormann, Othmane Aouane, Lars Kaestner, Christian Ruloff, Chaouqi Misbah, Thomas Podgorski, and Christian Wagner. The buckling instability of aggregating red blood cells. *Scientific Reports*, 7:7928, 2017.

[26] Fransisco Folgar and Charles L. Tucker III. Orientation behavior of fibers in concentrated suspensions. *Journal of Reinforced Plastics and Composites*, 3(2):98–119, 1984.

[27] Anne Greenbaum, Leslie Greengard, and Anita Mayo. On the numerical solution of the biharmonic equation in the plane. *Physica D: Nonlinear Phenomena*, 60(1):216–225, 1992.

[28] L. Greengard and V. Rokhlin. A Fast Algorithm for Particle Simulations. *Journal of Computational Physics*, 135(2):280–292, 1987.

[29] D. H. Griffel. *Applied Functional Analysis*. Dover Books on Mathematics. Dover Publications, 2002.

[30] David Harmon, Daniele Panozzo, Olga Sorkine, and Denis Zorin. Interference-aware Geometric Modeling. *ACM Transactions on Graphics*, 30(6):137:1–137:10, 2011.

[31] Johan Helsing and Rikard Ojala. On the evaluation of layer potentials close to their sources. *Journal of Computational Physics*, 227:2899–2921, 2008.

[32] George C. Hsiao. Boundary element methods—an overview. *Applied Numerical Mathematics*, 56(10):1356 – 1369, 2006.

[33] Wook Ryol Hwang and Martien A. Hulsen. Direct numerical simulations of hard particle suspensions in planar elongational flow. *Journal of Non-Newtonian Fluid Mechanics*, 136(2):167 – 178, 2006.

[34] David Abram Jack. *Advanced Analysis of Short-Fiber Polymer Composite Material Behavior*. PhD thesis, University of Missouri, 2006.

[35] G. B. Jeffery. The Motion of Ellipsoidal Particles Immersed in a Viscous Fluid. *Proceedings of the Royal Society of London Series A*, 102:161–179, 1922.

[36] C. G. Joung, N. Phan-Thien, and X.J. Fan. Direct simulation of flexible fibers. *Journal of Non-Newtonian Fluid Mechanics*, 99(1):1–36, 2001.

[37] Gökberk Kabacaoğlu, Bryan Quaife, and George Biros. Low-resolution simulations of vesicle suspensions in 2D. *Journal of Computational Physics*, 357:43–77, 2018.

[38] Seppo J. Karrila and Sangtae Kim. Integral Equations of the Second Kind for Stokes Flow: Direct Solutions for Physical Variables and Removal of Inherent Accuracy Limitations. *Chemical Engineering Communications*, 82(1):123–161, 1989.

[39] Sangtae Kim and Seppo J. Karrila. *Microhydrodynamics: Principles and Selected Applications*. Butterworth-Heinemann, Stoneham, MA, 1991.

[40] Erin Koos, Esperanza Linares Guerrero, Melany L. Hunt, and Christopher E. Brennen. Rheological measurements of large particles in high shear rate flows. *Physics of Fluids*, 24:013302, 2012.

[41] M. C. A. Kropinski. Integral equation methods for particle simulations in creeping flows. *Computers & Mathematics with Applications*, 38(5):67–87, 1999.

[42] Anthony J. C. Ladd. Numerical simulations of particulate suspensions via a discretized Boltzmann equation. Part 1. Theoretical foundation. *Journal of Fluid Mechanics*, 271:285–309, 1994.

[43] Anthony J. C. Ladd. Numerical simulations of particulate suspensions via a discretized Boltzmann equation. Part 2. Numerical results. *Journal of Fluid Mechanics*, 271:311–339, 1994.

[44] O. Ladyzhenskaia and R. Silverman. *The Mathematical Theory of Viscous Incompressible Flow*. Martino Fine Books, 2014.

[45] Ronald G. Larson. *The Structure and Rheology of Complex Fluids*. Oxford University Press, New York, NY, 1999.

[46] A. W. Lees and S. F. Edwards. The computer study of transport processes under extreme conditions. *Journal of Physics C: Solid State Physics*, 5(15):1921, 1972.

[47] Stefan B. Lindström and Tetsu Uesaka. Simulation of the motion of flexible fibers in viscous fluid flow. *Physics of Fluids*, 19(11):113307, 2007.

[48] Stefan B. Lindström and Tetsu Uesaka. Simulation of semidilute suspensions of non-Brownian fibers in shear flow. *The Journal of Chemical Physics*, 128(2):024901, 2008.

[49] Yaling Liu and Wing Kam Liu. Rheology of red blood cell aggregation by computer simulation. *Journal of Computational Physics*, 220(1):139–154, 2006.

[50] Eric Lorenz, Alfons Hoekstra, and Alfonso Caiazzo. Lees-Edwards boundary conditions for lattice Boltzmann suspension simulations. *Physical review. E, Statistical, nonlinear, and soft matter physics*, 79:036706, 04 2009.

[51] Libin Lu, Abtin Rahimian, and Denis Zorin. Contact-aware simulations of particulate Stokesian suspensions. *Journal of Computational Physics*, 347:160–182, 2017.

[52] Enkeleida Lushi and Petia M. Vlahovska. Periodic and chaotic orbits of plane-confined microrotors in creeping flows. *Journal of Nonlinear Science*, 25(5):1111–1123, Oct 2015.

[53] Christopher W. MacMinn, Eric R. Dufresne, and John S. Wettlaufer. Fluid-driven deformation of a soft granular material. *Physical Review X*, 5:011020, 2015.

[54] Giulio Maier. A quadratic programming approach for certain classes of non linear structural problems. *Meccanica*, 3(2):121–130, 1968.

[55] Andrea Albeto Mammoli. The treatment of lubrication forces in boundary integral equations. *Proceedings of the Royal Society A*, 462:855–881, 2006.

[56] Rajat Mittal and Gianluca Iaccarino. Immersed Boundary Methods. *Annual Review of Fluid Mechanics*, 37(1):239–261, 2005.

[57] K. G. Murty. *Linear Complementarity, Linear and Non Linear Programming*. Sigma series in applied mathematics. Heldermann Verlag, 1988.

[58] Igor V. Pivkin, Bruce Caswell, and George Em Karniadakisa. Dissipative Particle Dynamics. 27:85–110, 2010.

[59] Pit Polfer, Torsten Kraft, and Claas Bierwisch. Suspension modeling using smoothed particle hydrodynamics: Accuracy of the viscosity formulation and the suspended body dynamics. *Applied Mathematical Modelling*, 40:2606–1618, 2016.

[60] Henry Power. The completed double layer boundary integral equation method for two-dimensional Stokes flow. *IMA Journal of Applied Mathematics*, 51(2):123–145, 1993.

[61] Henry Power and Guillermo Miranda. Second Kind Integral Equation Formulation of Stokes Flows Past a Particle of Arbitrary Shape. *SIAM Journal on Applied Mathematics*, 47(4):689–698, 1987.

[62] C. Pozrikidis. *Boundary Integral and Singularity Methods for Linearized Viscous Flows*. Cambridge University Press, Cambridge, UK, 1992.

[63] Bryan Quaife and George Biros. High-volume fraction simulations of two-dimensional vesicle suspensions. *Journal of Computational Physics*, 274:245–267, 2014.

[64] Bryan Quaife and George Biros. High-order Adaptive Time Stepping for Vesicle Suspensions with Viscosity Contrast. *Procedia IUTAM*, 16:89–98, 2015.

[65] Bryan Quaife and George Biros. Adaptive time stepping for vesicle suspensions. *Journal of Computational Physics*, 306:478–499, 2017.

[66] Bryan Quaife and M Nicholas J. Moore. A boundary-integral framework to simulate viscous erosion of a porous medium. Preprint available at `https://arxiv.org/pdf/1804.06016.pdf`, 2018.

[67] Gary R. Marple, Alex Barnett, Adrianna Gillman, and Shravan Veerapaneni. A fast algorithm for simulating multiphase flows through periodic geometries of arbitrary shape. 38, 10 2015.

[68] Abtin Rahimian, Shravan Kumar Veerapaneni, and George Biros. Dynamic simulation of locally inextensible vesicles suspended in an arbitrary two-dimensional domain, a boundary integral method. *Journal of Computational Physics*, 229(18):6466–6484, 2010.

[69] Youcef Saad and Martin H. Schultz. GMRES: A generalized minimal residual algorithm for solving nonsymmetric linear systems. *SIAM Journal on Scientific and Statistical Computing*, 7(3):856–869, 1986.

[70] Akira Satoh, Roy W. Chantrell, Geoff N. Coverdale, and Shin ichi Kamiyama. Stokesian dynamics simulations of ferromagnetic colloidal dispersions in a simple shear flow. *Journal of Colloid and Interface Science*, 203(2):233 – 248, 1998.

[71] Christian F. Schmid, Leonard H. Switzer, and Daniel J. Klingenberg. Simulations of fiber flocculation: Effects of fiber properties and interfiber friction. *Journal of Rheology*, 44(4):781–809, 2000.

[72] Michael Siegel and Anna-Karin Tornberg. A local target specific quadrature by expansion method for evaluation of layer potentials in 3D. *Journal of Computational Physics*, 364:365–392, 2018.

[73] Chiara Sorgentone and Anna-Karin Tornberg. A highly accurate boundary integral equation method for surfactant-laden drops in 3D. *Journal of Computational Physics*, 360:167–191, 2018.

[74] D. E. Stewart and J. C. Trinkle. An implicit time-stepping scheme for rigid body dynamics with inelastic collisions and coulomb friction. *International Journal for Numerical Methods in Engineering*, 39(15):2673–2691, 1996.

[75] Rangarajan K. Sundaram. *A First Course in Optimization Theory*. Cambridge University Press, 1996.

[76] Anna-Karin Tornberg and Katarina Gustavsson. A numerical method for simulations of rigid fiber suspensions. *Journal of Computational Physics*, 215:172–196, 2006.

[77] Anna-Karin Tornberg and Michael J. Shelley. Simulating the dynamics and interactions of flexible fibers in Stokes flows. *Journal of Computational Physics*, 196(1):8–40, 2004.

[78] Lloyd N. Trefethen and David Bau. *Numerical Linear Algebra*. SIAM, 1997.

[79] Lloyd N. Trefethen and J. A. C. Weideman. The Exponentially Convergent Trapezoidal Rule. *SIAM Review*, 56(3):385–458, 2014.

[80] Satoru Yamamoto and Takaaki Matsuoka. Dynamic simulation of microstructure and rheology of fiber suspensions. *Polymer Engineering and Science*, 36:2396, 1996.

[81] Wen Yan, Eduardo Corona, Shravan Veerapaneni, and Michael Shelley. Resolving collisions in Stokes suspensions with an efficient and stable potential-free constrained optimization algorithm. 70th Annual Meeting of the APS Division of Fluid Dynamics, 2017.

[82] Wen Yan and Michael Shelley. Flexibly imposing periodicity in kernel independent FMM: A multipole-to-local operator approach. *Journal of Computational Physics*, 355:214 – 232, 2018.

[83] Kyongmin Yeo, Enkeleida Lushi, and Petia M. Vlahovska. Collective dynamics in a binary mixture of hydrodynamically coupled microrotors. *Phys. Rev. Lett.*, 114:188301, 2015.

[84] Lexing Ying, George Biros, and Denis Zorin. A high-order 3D boundary integral equation solver for elliptic PDEs in smooth domains. *Journal of Computational Physics*, 219(1):247–275, 2006.

[85] Shihani Zhao and Tsorng-Whay Pan. Numerical Simulation of Red Blood Cell Suspensions Behind a Moving Interface in a Capillary. *Numerical Mathematics Theory Methods and Applications*, 7(4):499–511, 2014.

BIOGRAPHICAL SKETCH

Lukas Bystricky received a B.Sc. in Mathematics from the University of British Columbia in 2012. After working for a year in Thundery Bay, Ontario he began his graduate studies in Scientific Computing at Florida State University.