

This is the accepted manuscript version of the following article:

L. af Klinteberg, A.-K. Tornberg, A fast integral equation method for solid particles in viscous flow using quadrature by expansion, *J. Comput. Phys.* (2016), <http://dx.doi.org/10.1016/j.jcp.2016.09.006>

©2016. This manuscript version is made available under the CC-BY-NC-ND 4.0 license <http://creativecommons.org/licenses/by-nc-nd/4.0/>

A fast integral equation method for solid particles in viscous flow using quadrature by expansion

Ludvig af Klinteberg^{a,*}, Anna-Karin Tornberg^a

^a *Department of Mathematics / Linné FLOW Centre / Swedish e-Science Research Centre, KTH Royal Institute of Technology, 100 44 Stockholm, Sweden*

Abstract

Boundary integral methods are advantageous when simulating viscous flow around rigid particles, due to the reduction in number of unknowns and straightforward handling of the geometry. In this work we present a fast and accurate framework for simulating spheroids in periodic Stokes flow, which is based on the completed double layer boundary integral formulation. The framework implements a new method known as quadrature by expansion (QBX), which uses surrogate local expansions of the layer potential to evaluate it to very high accuracy both on and off the particle surfaces. This quadrature method is accelerated through a newly developed precomputation scheme. The long range interactions are computed using the spectral Ewald (SE) fast summation method, which after integration with QBX allows the resulting system to be solved in $M \log M$ time, where M is the number of particles. This framework is suitable for simulations of large particle systems, and can be used for studying e.g. porous media models.

Keywords: Viscous flow, Stokes equations, Boundary integral methods, Quadrature by expansion, Fast Ewald summation

1. Introduction

Fluid flows involving microscopic, rigid particles are common both in nature and in industrial processes. The macroscopic properties of such systems are often determined by the particle interactions happening on the smallest scale of the flow. Examples of such systems can be found in sedimenting suspensions [17] and porous media modeling [13]. To fully understand the interactions between particles, and the effect those interactions have on the fluid flow, numerical simulation is a valuable tool. Stokes equations are often valid in this context, due to the small particle size and low fluid velocity. For problems governed by Stokes equations it is possible to use boundary integral methods, where the solution is represented as a layer potential from the boundaries of the domain (in this case the particle surfaces). Since the problem then is formulated as a boundary integral

*Corresponding author

Email address: ludvigak@kth.se (Ludvig af Klinteberg)

equation (BIE) on the union of the particle surfaces, the dimension of the domain which has to be discretized is reduced from \mathbb{R}^3 to \mathbb{R}^2 , significantly reducing the number of unknowns. This also removes the problem of combining a volume grid with moving boundaries, which can be challenging. However, boundary integral methods come with a set of challenges of their own, two of which can be considered major.

The first major challenge of boundary integral methods is that the resulting linear system after e.g. a Nyström discretization is dense, such that the cost of one left hand side evaluation is $\mathcal{O}(N^2)$, where N is the number of unknowns. This puts a severe limit on the size of the systems that can be considered, even if a rapidly converging iterative method is used. This can be overcome by evaluating the layer potential using a fast method, such as the fast multipole method (FMM) [16] or a fast Ewald summation method [14, 24] (for periodic problems). These methods reduce the cost of one left hand side evaluation to $\mathcal{O}(N)$ and $\mathcal{O}(N \log N)$, respectively, thereby making boundary integral methods suitable for large-scale computations.

The second major challenge of boundary integral methods is that one needs accurate quadrature of singular and nearly singular integrals when evaluating the layer potential. Developing such methods which are accurate, fast and work for arbitrary geometries is a topic of current research, particularly for the problem of nearly singular quadrature in three dimensions. In two dimensions there are efficient methods for nearly singular quadrature which gain a lot of their power from the complex variable formulation [18, 6]. In three dimensions the situation appears less resolved, though several different methods have been successfully used in practical applications [9, 10, 37, 29, 34, 38]. A common feature of many of these methods is however that they are highly target specific, meaning that the cost grows rapidly if there are many nearly singular integrals to be evaluated.

Quadrature by expansion (QBX) [20, 7] is a fairly new method for numerical integration of singular and nearly singular integrals. The method is built around evaluating layer potentials through local expansions, and comes equipped with a solid convergence theory [15]. It was originally presented for the Helmholtz kernel in two dimensions, but the principles of the method generalize directly both to three dimensions and other kernels. A promising feature of QBX is that it is possible to combine it with an FMM, thereby creating a fully $\mathcal{O}(N)$ method that is able to accurately evaluate the potential everywhere.

In this paper we extend the QBX framework to deal with the Stokes double layer potential in three dimensions, and combine it with the fast Ewald summation method presented in [2]. The result is a robust and scalable framework for computing Stokes flow around periodic systems of rigid, spheroidal particles. To limit the scope of the present paper, we restrict our attention to stationary particles. The method could easily be extended to dynamic problems — such as sedimentation — by coupling it to an ODE solver, as was done in [2].

The structure of this paper is as follows: In section 2 we state the necessary boundary integral formulation for Stokes flow around rigid particles. In section 3 we introduce QBX and discuss errors and parameter selection. In section 4 we show how QBX can be accelerated for our problem, leading to a computationally feasible method. In section 5

we briefly touch on the subject of Ewald summation, and how to combine it with QBX. Finally, in section 6 we present selected numerical results, which we draw both from validation tests and from an example application. We also include an appendix (A), which covers the computational details of the acceleration scheme presented in section 4.

2. Formulation

Our application of interest is that of rigid particles in a Newtonian fluid. The particles are assumed small enough that the Reynolds number is close to zero, such that we can model the flow as being Stokes flow, governed by the Stokes equations,

$$\begin{aligned} -\nabla P + \mu \Delta \mathbf{u} + \mathbf{f} &= 0, \\ \nabla \cdot \mathbf{u} &= 0. \end{aligned} \tag{1}$$

We here limit ourselves to spheroidal particles with semi-axes a and c , and assume $\mu = 1$. The surface of such a particle, oriented along the z coordinate axis, can in Cartesian coordinates be described as

$$\frac{x^2 + y^2}{a^2} + \frac{z^2}{c^2} = 1. \tag{2}$$

If parametrized in the spherical coordinates $\theta \in [0, \pi]$ and $\varphi \in [0, 2\pi)$, the same surface is given by

$$\begin{aligned} x &= a \sin(\theta) \cos(\varphi), \\ y &= a \sin(\theta) \sin(\varphi), \\ z &= c \cos(\theta). \end{aligned} \tag{3}$$

The particles can be classified into three distinct subgroups: prolate ($a < c$), spherical ($a = c$), and oblate ($a > c$).

2.1. Boundary integral formulation

We are considering Stokes flow as described by the Stokes equations (1), which by virtue of being a set of linear partial differential equations with constant coefficients have solutions that can be represented using boundary integrals. For these representations, the essential Green's functions are the stokeslet S , the stresslet T and the rotlet R ¹.

$$S_{ij}(\mathbf{x}, \mathbf{y}) = \frac{\delta_{ij}}{r} + \frac{r_i r_j}{r^3}, \tag{4}$$

$$T_{ijk}(\mathbf{x}, \mathbf{y}) = -6 \frac{r_i r_j r_k}{r^5}, \tag{5}$$

$$R_{ij}(\mathbf{x}, \mathbf{y}) = \epsilon_{ijk} \frac{r_k}{r^3}, \tag{6}$$

¹Throughout this paper we use the Einstein convention that an index appearing twice in an expression implies a summation over the set $\{1, 2, 3\}$, except for when there is a \sum explicitly defining the index.

where $\mathbf{r} = \mathbf{x} - \mathbf{y}$ and $r = |\mathbf{r}|$.

To compute the flow associated with a rigid particle in a viscous fluid, we use the completed boundary integral representation of Power & Miranda [27], where the flow in the domain is expressed as a double layer potential \mathcal{D} plus a completion flow \mathcal{V} and a background flow \mathbf{u}_{bg} ,

$$\mathbf{u}(\mathbf{x}) = \mathcal{D}[\Gamma, \mathbf{q}](\mathbf{x}) + \mathcal{V}(\mathbf{x}) + \mathbf{u}_{\text{bg}}(\mathbf{x}). \quad (7)$$

The double layer potential is the flow associated with a double layer density \mathbf{q} , defined on the particle surface Γ . It is computed by integrating \mathbf{q} over Γ together with the stresslet singularity T and the outward unit normal \mathbf{n} ,

$$\mathcal{D}_i[\Gamma, \mathbf{q}](\mathbf{x}) = \int_{\Gamma} T_{ijk}(\mathbf{x}, \mathbf{y}) q_j(\mathbf{y}) n_k(\mathbf{y}) dS_{\mathbf{y}}. \quad (8)$$

The integral in (8) is singular for $\mathbf{x} \in \Gamma$. It is then taken as the principal value integral, which is weakly singular (assuming Γ Lyapunov smooth). Letting \mathbf{x} approach the surface from either the interior or exterior domain, there is a jump in the double layer potential,

$$\lim_{\epsilon \rightarrow 0} \mathcal{D}[\Gamma, \mathbf{q}](\mathbf{x} \pm \epsilon \mathbf{n}) = \mp 4\pi \mathbf{q}(\mathbf{x}) + \mathcal{D}[\Gamma, \mathbf{q}](\mathbf{x}), \quad \mathbf{x} \in \Gamma. \quad (9)$$

The completion flow \mathcal{V} is added to the formulation due to the double layer potential's inability to represent a net force \mathbf{f} and torque \mathbf{t} on the particle [27]. For a rigid particle, a suitable completion flow is that generated by a stokeslet singularity S of magnitude $\mathbf{f}/8\pi\mu$ and a rotlet singularity R of magnitude $\mathbf{t}/8\pi\mu$, both located at the particle center \mathbf{x}_c ,

$$\mathcal{V}_i(\mathbf{x}) = \frac{1}{8\pi\mu} (S_{ij}(\mathbf{x}, \mathbf{x}_c) f_j + R_{ij}(\mathbf{x}, \mathbf{x}_c) t_j). \quad (10)$$

For a particle undergoing rigid body motion with translational velocity \mathbf{U} and rotational velocity $\mathbf{\Omega}$, we set a no-slip boundary condition,

$$\mathbf{u}(\mathbf{x}) = \mathbf{U} + \mathbf{\Omega} \times (\mathbf{x} - \mathbf{x}_c). \quad (11)$$

Letting \mathbf{x} in our flow field representation (7) go to the surface of the particle from the exterior, we get a diagonal term from (9). We then enforce the no-slip boundary condition (11) on the surface. This results in a BIE of the second kind in the density \mathbf{q} ,

$$-4\pi \mathbf{q}(\mathbf{x}) + \mathcal{D}[\Gamma, \mathbf{q}](\mathbf{x}) + \mathcal{V}(\mathbf{x}) + \mathbf{u}_{\text{bg}}(\mathbf{x}) = \mathbf{U} + \mathbf{\Omega} \times (\mathbf{x} - \mathbf{x}_c). \quad (12)$$

The fact that (12) is second kind with a compact integral operator is the main benefit of using a double layer formulation. After discretization using a quadrature rule and the Nyström method, the corresponding linear system can be said to be well conditioned in two senses; one is that the condition number is bounded (and usually very low), the other is that the condition number stays constant under grid refinement (this is in contrast to single layer formulations, where the condition number increases under refinement).

These properties make the system suitable for iterative solution using the generalized minimum residual method [31] (GMRES), which typically converges rapidly.

The equation system in (12) needs to be closed by specifying either the rigid body motion $(\mathbf{U}, \boldsymbol{\Omega})$ or the completion flow \mathcal{V} . This corresponds to using two different formulations: the resistance problem formulation or the mobility problem formulation.

2.1.1. The resistance problem

When the rigid body motion $(\mathbf{U}, \boldsymbol{\Omega})$ of the particle is known, the problem is to find the resulting force \mathbf{f} and torque \mathbf{t} exerted by the fluid on the particle. These can be related to the double layer density \mathbf{q} through the formulation by Power & Miranda [27],

$$\mathbf{f}(\mathbf{q}) = \int_{\Gamma} \mathbf{q}(\mathbf{y}) dS_{\mathbf{y}}, \quad (13)$$

$$\mathbf{t}(\mathbf{q}) = \int_{\Gamma} \mathbf{q}(\mathbf{y}) \times (\mathbf{y} - \mathbf{x}_c) dS_{\mathbf{y}}, \quad (14)$$

such that the completion flow is a functional of \mathbf{q} ,

$$\mathcal{V}_i(\mathbf{x}) = \mathcal{V}_i[\Gamma, \mathbf{q}](\mathbf{x}) = \frac{1}{8\pi\mu} (S_{ij}(\mathbf{x}, \mathbf{x}_c) f_j(\mathbf{q}) + R_{ij}(\mathbf{x}, \mathbf{x}_c) t_j(\mathbf{q})). \quad (15)$$

2.1.2. The mobility problem

When the external forcing (\mathbf{f}, \mathbf{t}) on the particle is known, the problem is to find the rigid body motion $(\mathbf{U}, \boldsymbol{\Omega})$ of the particle. Using the formulation available in e.g. Pozrikidis [28], the rigid body motion vectors can be computed as functionals of \mathbf{q} ,

$$\mathbf{V}(\mathbf{q}) = -\frac{4\pi}{S_{\Gamma}} \int_{\Gamma} \mathbf{q}(\mathbf{y}) dS_{\mathbf{y}}, \quad (16)$$

$$\boldsymbol{\Omega}(\mathbf{q}) = -4\pi \sum_{n=1}^3 \frac{\boldsymbol{\omega}^{(n)}}{A_n} \left(\boldsymbol{\omega}^{(n)} \cdot \int_{\Gamma} (\mathbf{y} - \mathbf{x}_c) \times \mathbf{q}(\mathbf{y}) dS_{\mathbf{y}} \right), \quad (17)$$

where S_{Γ} is the surface area of Γ and

$$A_n = \int_{\Gamma} \left[\boldsymbol{\omega}^{(n)} \times (\mathbf{y} - \mathbf{x}_c) \right] \cdot \left[\boldsymbol{\omega}^{(n)} \times (\mathbf{y} - \mathbf{x}_c) \right] dS_{\mathbf{y}}. \quad (18)$$

The vectors $\boldsymbol{\omega}^{(n)}$, $n = 1, 2, 3$, are independent unit vectors satisfying

$$\frac{1}{\sqrt{A_n A_m}} \int_{\Gamma} \left[\boldsymbol{\omega}^{(m)} \times (\mathbf{y} - \mathbf{x}_c) \right] \cdot \left[\boldsymbol{\omega}^{(n)} \times (\mathbf{y} - \mathbf{x}_c) \right] dS_{\mathbf{y}} = \delta_{mn}. \quad (19)$$

2.2. Particle systems

For a system of M particles, the linearity of the Stokes equations allows us to express the velocity field as a superposition of the fields from the individual particles,

$$\mathbf{u}(\mathbf{x}) = \sum_{\alpha=1}^M (\mathcal{D}^{\alpha}(\mathbf{x}) + \mathcal{V}^{\alpha}(\mathbf{x})) + \mathbf{u}_{\text{bg}}(vx), \quad (20)$$

where Γ_α is the surface of particle α ,

$$\mathcal{D}^\alpha(\mathbf{x}) = \mathcal{D}[\Gamma_\alpha, \mathbf{q}](\mathbf{x}), \quad (21)$$

and $(\mathcal{V}^\alpha, \mathbf{U}^\alpha, \mathbf{\Omega}^\alpha)$ are defined either as in the resistance formulation or the mobility formulation.

Evaluating (20) and enforcing the no-slip boundary condition (11) on the surface of each particle gives the boundary integral equation for a system of particles,

$$-4\pi\mathbf{q}(\mathbf{x}) + \sum_{\alpha=1}^M (\mathcal{D}^\alpha(\mathbf{x}) + \mathcal{V}^\alpha(\mathbf{x})) + \mathbf{u}_{\text{bg}}(\mathbf{x}) = \mathbf{U}^\beta + \mathbf{\Omega}^\beta \times (\mathbf{x} - \mathbf{x}_c^\beta), \quad (22)$$

$$\mathbf{x} \in \Gamma_\beta, \quad \beta = 1, \dots, M.$$

2.3. Discrete formulation

To solve the (single particle) BIE (12), we first need a way of numerically evaluating integrals over the particle surface, which we denote

$$\mathbb{I}[f] = \int_{\Gamma} f(\mathbf{y}) dS_{\mathbf{y}}. \quad (23)$$

This we do using a quadrature rule \mathbb{Q}_N , which defines a set of N nodes \mathbf{x}_i and weights w_i on Γ , such that

$$\mathbb{I}[f] \approx \mathbb{Q}_N[f] = \sum_{i=1}^N f(\mathbf{y}_i) w_i. \quad (24)$$

The details of the quadrature used will be discussed further in section 3. We denote by the superscript h a quantity computed using \mathbb{Q}_N , e.g.

$$\mathcal{D}_i^h[\Gamma, \mathbf{q}](\mathbf{x}) = \mathbb{Q}_N[T_{ijk}(\mathbf{x}, \cdot) q_j(\cdot) n_k(\cdot)]. \quad (25)$$

Applying the Nyström method, where the integral equation is enforced at the quadrature nodes, the BIE for the single-particle system (12) is approximated by the $3N \times 3N$ linear system in the density values $\mathbf{q}(\mathbf{x}_i)$,

$$-4\pi\mathbf{q}(\mathbf{x}_i) + \mathcal{D}_i^h[\Gamma, \mathbf{q}](\mathbf{x}_i) + \mathcal{V}(\mathbf{x}_i) = \mathbf{U} + \mathbf{\Omega} \times (\mathbf{x}_i - \mathbf{x}_c), \quad (26)$$

$$i = 1, \dots, N,$$

where either $\mathcal{V} = \mathcal{V}^h$ or $\mathbf{U} = \mathbf{U}^h$ and $\mathbf{\Omega} = \mathbf{\Omega}^h$, depending on if the problem is a resistance problem or a mobility problem.

The matrix corresponding to the linear system (26) is dense, so the cost of applying a direct solution method to it would be $\mathcal{O}(N^3)$. However, when solving with GMRES the number of iterations required to reach a certain tolerance is largely independent of the discretization. The cost is then $\mathcal{O}(N^2)$, with a constant that depends on the problem geometry. This is acceptable for a single particle, since the solution \mathbf{q} usually

can be resolved with a limited number of surface nodes. For a problem with multiple particles, however, the corresponding linear system (derived from (22)) is $3MN \times 3MN$ and dense. The cost then grows quadratically with the number of particles, which severely limits the size of the problem that can be considered. The remedy for this is to use a fast method, which exploits the structure of the problem to evaluate the left hand side of the problem (also known as the matrix-vector product, or *matvec*) in less than quadratic time. Common fast methods are the FMM and, for periodic problems, fast Ewald summation. The respective complexities for these two methods are $\mathcal{O}(M)$ and $\mathcal{O}(M \log M)$ as the number of particles is increased. We have here used fast Ewald summation, which will be further discussed in section 5.

3. Quadrature

To derive the linear system (26) corresponding to the BIE (12), we needed a quadrature method \mathbf{Q}_N (24) for evaluating the double layer potential \mathcal{D} (8). We will in this section describe such a method, which is able to accurately resolve \mathcal{D} also when the integral is nearly singular or singular.

3.1. Spheroidal grid

When performing quadrature on a spheroidal particle, it is natural to use the parametrized description (3) in (θ, φ) and to discretize it using an $n_\theta \times n_\varphi$ grid in those coordinates, with the grid aspect ratio given by

$$\frac{n_\varphi}{n_\theta} = \frac{a}{c}. \quad (27)$$

This gives us a total of $N = n_\theta n_\varphi$ quadrature points on the surface. We let $(\theta_i, \lambda_i^\theta)$, $i \in \{1 \dots n_\theta\}$, be the nodes and quadrature weights of an n_θ -point Gauss-Legendre quadrature rule on the interval $[0, \pi]$, and $(\varphi_j, \lambda_j^\varphi)$, $j \in \{1 \dots n_\varphi\}$, be the nodes and quadrature weights of the trapezoidal rule on the interval $[0, 2\pi)$. This gives us the quadrature rule


$$\mathbf{Q}_{n_\theta n_\varphi}[f] = \sum_{i=1}^{n_\theta} \sum_{j=1}^{n_\varphi} f(\mathbf{y}(\theta, \varphi)) W(\theta, \varphi) \lambda_i^\theta \lambda_j^\varphi \quad (28)$$

where $W(\theta, \varphi)$ is the area element. For ease of notation we can simplify this expression by letting one index cover all quadrature points and introducing \mathbf{y}_i and w_i such that

$$\mathbf{Q}_N[f] = \sum_{i=1}^N f(\mathbf{y}_i) w_i = \mathbf{Q}_{n_\theta n_\varphi}[f]. \quad (29)$$

We denote this grid and the accompanying quadrature rule the *spheroidal grid*, and will refer to the quadrature rule as the direct quadrature.

The direct quadrature rule defined in (28) allows us to approximate integrals over the surface with spectral accuracy for *smooth* and *well-resolved* integrands, and works



(a) Stokes flow around a single translating spheroid, exact solution from [12]. (b) Error when evaluating double layer potential from spheroid using direct quadrature. The scale is logarithmic, indicating an exponential growth rate as $\boldsymbol{x} \rightarrow \Gamma$.

Figure 1: Example of error due to nearly singular integration (b), when evaluating flow due to a single particle (a).

well for evaluating the double layer potential (8) when the target point \boldsymbol{x} is far enough away from the surface Γ . However, when \boldsymbol{x} is on Γ the integral is singular and can not be computed using a quadrature rule for smooth functions. There is also a difficulty when \boldsymbol{x} lies in the domain, but is close to Γ . The singularity in the stresslet then causes the integral kernel to be sharply peaked, and trying to evaluate it using a direct quadrature rule gives an error which grows exponentially as \boldsymbol{x} approaches Γ (see Figure 1 and results in [3]). We refer to this as the integral being *nearly singular*.

The singular case ($\boldsymbol{x} \in \Gamma$) can for the double layer potential be evaluated using a singularity subtraction method, which gives third order accuracy at virtually no increased cost [2]. For higher order accuracy in the singular case, and for evaluating the nearly singular case, specialized quadrature methods are required.

In what follows of this section we will first give a brief introduction to an existing quadrature method for singular and nearly singular layer potentials that is based on local expansions. We will then show how this method can be used to evaluate the Stokes double layer potential. Lastly, we will show how this method in our case can be accelerated by geometric considerations.

3.2. Quadrature by expansion

Quadrature by expansion (QBX) is a new quadrature scheme first introduced for the Helmholtz equation in two dimensions by Klöckner et al. [20]. The method is based on the observation that a direct quadrature rule applied to a layer potential is inaccurate when evaluated close to or on the surface, but well resolved when sufficiently far away, as illustrated in Figure 1. Since the layer potential is a smooth function away from the boundary, we can pick any point \boldsymbol{c} in the domain and create a local expansion of the




Figure 2: Local expansion and domain of convergence.

potential about that point. Picking that point in the well-resolved part of the domain, we can compute the coefficients of the expansion using our direct quadrature rule. The domain of convergence for the local expansion is the ball of radius r ,

$$r = \min_{\mathbf{y} \in \Gamma} |\mathbf{c} - \mathbf{y}|, \quad (30)$$

centered on \mathbf{c} , and so the expansion can be used to accurately evaluate the potential close to the surface and at the point where the ball touches the surface (see Figure 2 for an example geometry). This is the essential idea of QBX. For further reading regarding the convergence and evaluation of these local expansions, we refer to Epstein et al. [15] and Barnett [7].

In three dimensions, a good starting point for the construction of local expansions for layer potentials is the expansion of the Green's function for the Laplace equation, called the Laplace expansion,

$$\frac{1}{|\mathbf{x} - \mathbf{y}|} = \sum_{l=0}^{\infty} \frac{4\pi}{2l+1} \sum_{m=-l}^l r_x^l Y_l^{-m}(\theta_x, \varphi_x) \frac{1}{r_y^{l+1}} Y_l^m(\theta_y, \varphi_y), \quad (31)$$

where Y_l^m is the spherical harmonics function of degree l and order m , and $(r_x, \theta_x, \varphi_x)$ and $(r_y, \theta_y, \varphi_y)$ are the coordinates of \mathbf{x} and \mathbf{y} in a spherical coordinate system centered at \mathbf{c} . This expansion gives us the separation of source (\mathbf{y}) and target (\mathbf{x}) that is necessary for the method. As a first application, we will consider a general form of the double layer potential of the Laplace equation, which we henceforth will refer to as the dipole potential. For a smooth vector density $\boldsymbol{\rho}$ defined on the surface Γ , we define it as

$$\mathcal{L}[\Gamma, \boldsymbol{\rho}](\mathbf{x}) := \int_{\Gamma} \boldsymbol{\rho} \cdot \nabla_{\mathbf{y}} \frac{1}{|\mathbf{x} - \mathbf{y}|} dS_{\mathbf{y}}. \quad (32)$$

Inserting the expansion (31) into (32) and moving the terms related to the target point \mathbf{x} out of the integration, we get

$$\mathcal{L}[\Gamma, \boldsymbol{\rho}](\mathbf{x}) = \sum_{l=0}^{\infty} \sum_{m=-l}^l r_x^l Y_l^{-m}(\theta_x, \varphi_x) z_{lm}, \quad (33)$$

where z_{lm} are the local expansion coefficients of the potential,

$$z_{lm} = \frac{4\pi}{2l+1} \int_{\Gamma} \boldsymbol{\rho} \cdot \nabla \frac{1}{r_y^{l+1}} Y_l^m(\theta_y, \varphi_y) dS_{\mathbf{y}}. \quad (34)$$

Once the expansion coefficients z_{lm} are known, they can be used to evaluate the potential at any target point \mathbf{x} in the domain of convergence. At this stage we wish to highlight two things: (i) We have so far made no approximations, evaluating the potential through expressions (33) and (34) is still exact. The sum in l will need to be truncated, and we will discuss the truncation error that it introduces in section 3.5. (ii) The integrand in (34) contains no singularity, and can be evaluated using a direct quadrature rule. The quadrature errors related to this direct quadrature will also be discussed in section 3.5.

3.3. QBX for the stresslet

To evaluate the Stokes double layer potential (8) using QBX, we need to be able to expand its kernel in a way that separates source and target, just as we could for the dipole potential. One way of doing this is to express the double layer potential in terms of the dipole potential, for which we already have the expansion (33). Here we use the result by Tornberg and Greengard [35], which was developed with the purpose of using a harmonic FMM to compute the flow due to a collection of stokeslets and stresslets. Denoting by F the kernel of the dipole potential,

$$F_i(\mathbf{x}, \mathbf{y}) = \nabla_{\mathbf{y}} \frac{1}{r} = \frac{r_i}{r^3}, \quad \mathbf{r} = \mathbf{x} - \mathbf{y}, \quad (35)$$

the kernel of the double layer potential can be expressed as

$$\begin{aligned} T_{ijk}(\mathbf{x}, \mathbf{y}) n_k &= \left((x_j - y_j) \frac{\partial}{\partial x_i} - \delta_{ij} \right) F_k(\mathbf{x}, \mathbf{y}) n_k \\ &\quad + \left((x_k - y_k) \frac{\partial}{\partial x_i} - \delta_{ik} \right) F_j(\mathbf{x}, \mathbf{y}) n_k. \end{aligned} \quad (36)$$

Rearranging terms, the integrand of (8) can then be written

$$\begin{aligned} T_{ijk}(\mathbf{x}, \mathbf{y}) q_j n_k &= \left(x_j \frac{\partial}{\partial x_i} - \delta_{ij} \right) F_k(\mathbf{x}, \mathbf{y}) (q_j n_k + n_j q_k) \\ &\quad - \frac{\partial}{\partial x_i} F_j(\mathbf{x}, \mathbf{y}) (y_k q_k n_j + y_k n_k q_j). \end{aligned} \quad (37)$$

Consequently, the double layer potential can be expressed in terms of dipole potentials as

$$\begin{aligned} \mathcal{D}_i[\Gamma, \mathbf{q}](\mathbf{x}) &= \left(x_j \frac{\partial}{\partial x_i} - \delta_{ij} \right) \mathcal{L}[\Gamma, q_j \mathbf{n} + n_j \mathbf{q}](\mathbf{x}) \\ &\quad - \frac{\partial}{\partial x_i} \mathcal{L}[\Gamma, y_k q_k \mathbf{n} + y_k n_k \mathbf{q}](\mathbf{x}) \end{aligned} \quad (38)$$

With this result we can create an expansion of the double layer potential by creating four expansions of the dipole potential with four different densities (including the implicit summation over j). This means that the three vector components of the double layer potential are computed using four scalar expansions. The expansion coefficients that need to be computed for a given expansion center \mathbf{c} are then

$$\begin{aligned} z_{lm}^j &= \int_{\Gamma} (q_j \mathbf{n} + n_j \mathbf{q}) \cdot \nabla \frac{1}{r_y^{l+1}} Y_l^m(\theta_y, \varphi_y) dS_{\mathbf{y}}, \quad j = 1, 2, 3, \\ z_{lm}^4 &= \int_{\Gamma} (y_k q_k \mathbf{n} + y_k n_k \mathbf{q}) \cdot \nabla \frac{1}{r_y^{l+1}} Y_l^m(\theta_y, \varphi_y) dS_{\mathbf{y}}. \end{aligned} \quad (39)$$

The spherical harmonics are conjugate symmetric in m , $Y_l^{-m}(\theta, \varphi) = Y_l^m(\theta, \varphi)^*$, and since the density \mathbf{q} is real the same conjugate symmetry holds for the expansion coefficients,

$$z_{l,-m}^j = \left(z_{lm}^j \right)^*, \quad j = 1, \dots, 4. \quad (40)$$

Hence, we only need to compute the coefficients for $0 \leq m \leq l$. Truncating the expansion at some $p = l_{\max}$, the total number of coefficients that we need to compute for each component j is

$$N_p = \frac{p^2 + 3p + 2}{2}. \quad (41)$$

3.4. Numerical scheme

To solve the discrete BIE for a single particle (26), we need to evaluate the double layer potential at the quadrature points $\mathbf{x}_1, \dots, \mathbf{x}_N$. Each local expansion of the potential is valid at a single point on the surface, where the ball of convergence touches the surface, so we need N expansion centers $\mathbf{c}_1, \dots, \mathbf{c}_N$ located at a distance r normally from the quadrature points,

$$\mathbf{c}_i = \mathbf{x}_i + r \mathbf{n}_i. \quad (42)$$

The same expansions are used to evaluate the potential on the surface and in the domain. A natural choice of expansion radius is then

$$r = \frac{d_{\epsilon}}{2}, \quad (43)$$

where d_{ϵ} has been determined as the minimum distance from the surface at which the direct quadrature can be used for a given error tolerance ϵ . For target points \mathbf{x} such that $\min_{\mathbf{y} \in \Gamma} |\mathbf{x} - \mathbf{y}| < d_{\epsilon}$, the potential is then evaluated using the closest expansion center.

3.5. Error analysis

Here we give a sketch of an error analysis for the dipole potential, with the argument that it carries over to the double layer potential, as the latter is constructed as a combination of the former. For more in-depth discussions we refer to [15] and [3].

When evaluating the dipole potential (32) using QBX, we truncate the expansion (33) to some order p and use our direct Gauss-Legendre/trapezoidal quadrature rule (28) to approximate the coefficients z_{lm} (34). Denoting the approximate coefficients \tilde{z}_{lm} , the error is

$$E = \left| \mathcal{L}[\Gamma, \boldsymbol{\rho}](\mathbf{x}) - \sum_{l=0}^p \sum_{m=-l}^l r_x^l Y_l^{-m}(\theta_x, \varphi_x) \tilde{z}_{lm} \right|. \quad (44)$$

Adding and subtracting the exact coefficients z_{lm} to the expansion, the error can be split into two components,

$$E \leq E_T + E_Q, \quad (45)$$

where E_T is the truncation error due to using a finite number of terms in the expansion,

$$E_T = \left| \mathcal{L}[\Gamma, \boldsymbol{\rho}](\mathbf{x}) - \sum_{l=0}^p \sum_{m=-l}^l r_x^l Y_l^{-m}(\theta_x, \varphi_x) z_{lm} \right|, \quad (46)$$

and E_Q is the quadrature error due to using a discrete quadrature rule,

$$E_Q = \left| \sum_{l=0}^p \sum_{m=-l}^l r_x^l Y_l^{-m}(\theta_x, \varphi_x) (\tilde{z}_{lm} - z_{lm}) \right|. \quad (47)$$

From the results in [15] we have that the truncation error can be expected to behave as

$$E_T = \mathcal{O}(r^{p+1}), \quad (48)$$

assuming that the lengths have been normalized such that $r < 1$. In [3] an estimate was derived for the QBX quadrature error, when considering the Laplace single layer potential on a spheroid. Simplifying those results, we can expect that the quadrature error should behave approximately as


$$E_Q = \mathcal{O}\left(\left(\beta \frac{r}{h}\right)^p e^{-\alpha r/h}\right), \quad (49)$$

for some $\alpha, \beta > 0$ and h a measure of the grid size; for the spheroidal grid we define


$$h := 2\pi a/n_\varphi. \quad (50)$$

The constants can be expected to take values such that $\beta r/h > 1$. Putting the two expressions together, we see that the total error has one component that decays with p and one component that grows with p ,

$$E = \mathcal{O}(r^{p+1}) + \mathcal{O}\left(\left(\beta \frac{r}{h}\right)^p e^{-\alpha r/h}\right). \quad (51)$$



(a) L^2 error in the density when solving the mobility problem on an oblate spheroid ($a/c = 2/1$) with a 40×20 grid using $r/h = 3/2$ and a range of p and κ .



(b) Expansion centers are used on both sides of Γ , to accurately capture the principal value integral.

Figure 3

This means that if we fix r and h for a given problem, then there exists an optimal value for p that minimizes the error. If we instead fix the relation r/h at some constant value, e.g. $r = 3h$, then the quadrature error stays constant ($E_Q = C_Q$) such that for a given p

$$E = \mathcal{O}(h^{p+1}) + C_Q. \quad (52)$$

The scheme can then be viewed as having order of accuracy $(p+1)$ under grid refinement, up to the point where the truncation and quadrature errors are equal. After that we have to refine the grid (while maintaining a constant r) to get any additional improvements.

In practice we want to be able to get arbitrarily high accuracy out of QBX, without having to introduce more degrees of freedom. To overcome the problem of the quadrature error growing with p , we therefore *upsample* the surface density to a factor κ finer grid before computing the coefficients. The fine grid is of the same type as the original, but with $\kappa n_\theta \times \kappa n_\varphi$ points, and the density on the fine grid is computed using trigonometric interpolation and barycentric Lagrange interpolation [8]. This allows us to write the total error as

$$E = \mathcal{O}(r^{p+1}) + \mathcal{O}\left(\left(\beta \frac{\kappa r}{h}\right)^p e^{-\alpha \kappa r/h}\right). \quad (53)$$

The truncation and quadrature errors can then be controlled separately, since the exponential term will rapidly reduce the quadrature error when κ is increased.

To illustrate the effects of the parameters p and κ , we solve the mobility problem for a single oblate spheroid ($a/c = 2$) with a 20×40 grid using $r/h = 3/2$ and a range of p and κ . Measuring the L^2 norm of the error in the density compared to a reference solution, we can see in Figure 3a that the error decreases exponentially in p until it hits

a floor given by the resolution of the underlying grid, but only for κ large enough. If κ is taken too small, the error increases almost exponentially in p after reaching a minimum. Finding the optimal κ for a given p is not trivial, and picking κ larger than necessary is costly, since it adds a factor κ^2 to the computational complexity.

3.6. Two-sided expansions

When evaluating the singular double layer potential on a surface, we want the principal value integral. However, what we get when evaluating it using QBX is the one-sided limit from the side of the expansion center, which we denote \mathcal{D}^+ or \mathcal{D}^- depending on side,

$$\mathcal{D}^\pm[\Gamma, \mathbf{q}](\mathbf{x}) = \lim_{\epsilon \rightarrow 0} \mathcal{D}[\Gamma, \mathbf{q}](\mathbf{x} \pm \epsilon \mathbf{n}). \quad (54)$$

The relation between the principal value and the one-sided limit is known through the jump relation (9), which can be used together with the QBX result to recover the principal value. However, and this was originally noted in [20], it turns out that the spectrum of the linear system better approximates that of the original operator if we compute the principal value through the relation

$$\mathcal{D} = \frac{\mathcal{D}^+ + \mathcal{D}^-}{2}. \quad (55)$$

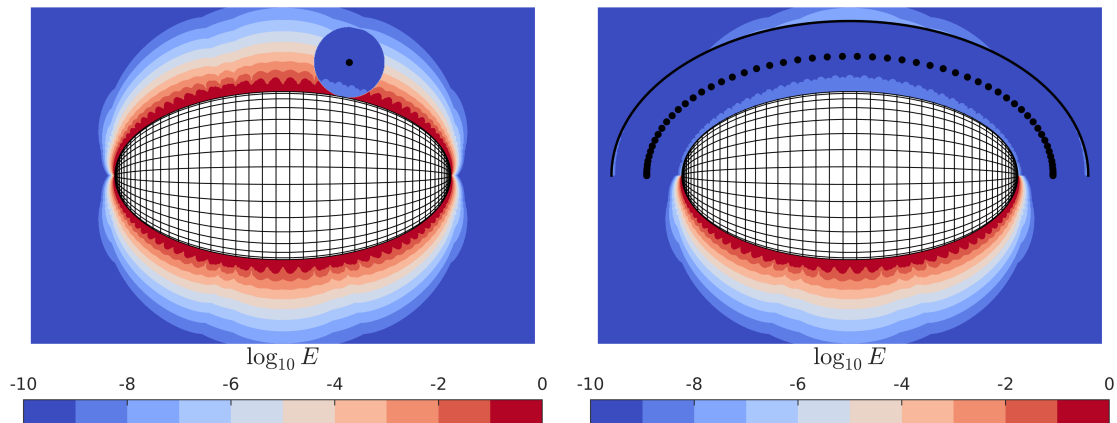
To do this in practice, we need to put one expansion center on each side of the target point on the surface and take the average potential of the two to get the principal value, as illustrated in Figure 3b. This seemingly doubles the amount of work needed, but through the acceleration scheme of section 4 that extra cost is hidden from the computations.

3.7. Parameter selection

For a given geometry, grid resolution h and tolerance ϵ , an empirical process for selecting the QBX parameters p , κ and r is as follows:

1. Set $r = d_\epsilon/2$, such that the error from the direct quadrature will be less than ϵ when the evaluation point is at least a distance d_ϵ from Γ .
2. Set p large enough such that the truncation error E_T from the QBX expansion is smaller than ϵ for points closer than d_ϵ from Γ .
3. Set κ large enough to match p , such that the quadrature error E_Q is kept below ϵ .

Figure 4 shows an example where the QBX parameters have been tuned for high accuracy.



(a) Error when using direct quadrature, punctured in disc around an expansion center. (b) Error when QBX is used in top half plane. Solid black line marks threshold d_ϵ .

Figure 4: Relative error in solution around prolate spheroid, $(a, c) = (1, 2)$, compared to exact solution from [12]. Grid is 60×30 , QBX parameters are $p = 40$, $k = 10$, $r/h = 2$. Note the error pattern inside the QBX area, which is due to an error in the solution, not the quadrature.

4. Acceleration of QBX on spheroidal grid

Evaluating the potential from one particle using the QBX scheme described in section 3 involves a lot of work. For a spheroidal grid with N points, a given upsampling factor κ and expansion order p , we need to interpolate the density to a factor κ^2 larger grid and then compute $\mathcal{O}(p^2)$ expansion coefficients at every expansion center. This puts the cost for computing the potential from a particle on itself, the "self-interaction", at $\mathcal{O}(\kappa^2 p^2 N^2)$. One natural way of speeding up the process is to incorporate the QBX scheme into an FMM, which hierarchically computes potentials using expansions of the same kind as those used in QBX (the method was in fact conceived with this in mind). This makes it possible to compute everything on the fly at a computational cost that is linear in N (but still includes $\kappa^2 p^2$). The required modifications to the FMM are however nontrivial, and the only unified FMM/QBX method published to date (by Rachh et al. [30]) is for two-dimensional problems.

Another way of achieving speedup is to precompute as much as possible of the work associated with QBX, and then use precomputed values whenever possible. This comes at a potentially large storage cost, and still has a computational complexity that is quadratic in N , albeit with a significantly smaller leading constant. In this work we have used precomputations in combination with a fast Ewald summation method, which computes the long-range interactions between M particles in $\mathcal{O}(M \log M)$ time (more on this in section 5). This is feasible partly because the number N of grid points required on a spheroid is limited, such that the overall computational cost scales like $\mathcal{O}(M \log M)$, and partly because the spheroidal geometry of the particles in our problem significantly reduces the storage required for the precomputed values. We will in this section outline the principles and main results of our precomputation strategy; the details

of the implementation are available for reference in appendix A.

4.1. Self-interaction

When solving the boundary integral equation (12) on a single particle using the spheroidal grid of section 3.1, we typically use the Nyström method to create a linear system which we solve iteratively. We then have N grid points denoted $\mathbf{x}_1, \dots, \mathbf{x}_N$ on the surface, and for every iteration we need to compute the self-interaction at those points. With $\mathbf{Q} \in \mathbb{R}^{3N}$ containing $\{\mathbf{q}(\mathbf{x}_i)\}_{i=1}^N$, we can let $R_i \in \mathbb{R}^{3 \times 3N}$ be the map $\mathbf{Q} \rightarrow \mathbf{u}(\mathbf{x}_i)$ using QBX, such that

$$\mathbf{u}(\mathbf{x}_i) = R_i \mathbf{Q}. \quad (56)$$

The matrix R_i then represents the resulting action after upsampling the density to a fine grid, using the fine grid to compute the expansion coefficients at the expansion center \mathbf{c}_i , and finally evaluating that expansion at \mathbf{x}_i . Precomputing all the matrices R_1, \dots, R_N requires $3N \times 3N$ memory storage, but once computed they allow the computation of $\mathbf{Q} \rightarrow \mathbf{u}(\mathbf{x}_i)$ at all points \mathbf{x}_i at a the cost of a matrix-vector product, i.e. $\mathcal{O}(N^2)$. This means that the $\mathcal{O}(p^2 \kappa^2)$ cost related to the QBX parameters — and ultimately to the accuracy — only appears in the precomputation step. Once the R_i matrices are computed the self-interaction evaluation has a fixed cost, independent of which QBX parameters were used.

The memory cost of precomputing all the self interaction matrices can be reduced by taking advantage of the fact that the spheroidal grid is rotationally symmetric about the polar axis. This symmetry means that the matrices R_a and R_b related to two points \mathbf{x}_a and \mathbf{x}_b on the same latitude ($\theta_a = \theta_b$) can be related as


$$R_a = T_z R_b (T_z^{-1} \otimes P), \quad (57)$$

where $T_z \in \mathbb{R}^{3 \times 3}$ represents a rotation about the polar axis of the particle, and $P \in \mathbb{R}^{N \times N}$ is a permutation matrix, i.e. a row permutation of the identity matrix that represents a permutation of the point ordering. The symbol \otimes denotes the Kronecker product. As a consequence, it is enough to precompute matrices related to the n_θ points on the first longitude, and then use rotations and permutations to compute the potential at all grid points with those matrices.

The savings in memory and precomputation time can be further increased by taking advantage of the fact the spheroidal grid has a mirror symmetry about the equator. For two points \mathbf{x}_b and \mathbf{x}_c on the same longitude ($\varphi_b = \varphi_c$) but different sides of the equator, such that $\theta_b = \pi - \theta_c$, we can relate the matrices R_b and R_c as

$$R_b = T_x R_c (T_x^{-1} \otimes F), \quad (58)$$

where $T_x \in \mathbb{R}^{3 \times 3}$ represents a rotation that mirrors the points \mathbf{x}_b and \mathbf{x}_c into each other, and $F \in \mathbb{R}^{N \times N}$ is a permutation matrix. Taking advantage of this mirror symmetry allows a further reduction of the required number of precomputed matrices to $n_\theta/2$, corresponding to the points marked in Figure 5a.



(a) Spheroidal grid with the $n_\theta/2$ first points circled.

(b) Example of configuration where near-singular quadrature is necessary.

Figure 5

4.2. Nearly singular evaluation

When particles are close to each other, such as in Figure 5b, we need to use QBX to evaluate their interaction accurately, since at least some of the quadrature points on one particle will be too close to the other particle for the direct quadrature to be accurate without substantial refinement. To be able to evaluate the potential at arbitrary points close to a particle using QBX, we need the expansion coefficients for the potential at all the expansion centers in the domain outside the surface.

We let $\mathbf{z}^j(\mathbf{c}_i) \in \mathbb{C}^{N_p}$, $j \in \{1 \dots 4\}$, denote the four vectors of N_p (41) coefficients each at an expansion center \mathbf{c}_i , as defined in (39). Analogously to (56), the process of upsampling the density and computing the expansion coefficients at all the expansion coefficients can then be represented by a set of matrices $M_i^j \in \mathbb{C}^{N_p \times 3N}$, such that

$$\mathbf{z}^j(\mathbf{c}_i) = M_i^j \mathbf{Q}, \quad \begin{aligned} i &= 1, \dots, N, \\ j &= 1, \dots, 4. \end{aligned} \quad (59)$$

Using the same arguments of symmetry as above, we can relate these matrices by rotational symmetry as

$$M_a^j = \sum_{i=1}^4 A_{ij} E M_b^i (T_z^{-1} \otimes P), \quad (60)$$

and by mirror symmetry as

$$M_b^j = \sum_{i=1}^4 B_{ij} \Theta^j M_c^i (T_x^{-1} \otimes F), \quad (61)$$

where $\{\Theta, E\} \in \mathbb{R}^{N_p \times N_p}$ are diagonal, and $\{A, B\} \in \mathbb{R}^{4 \times 4}$ orthogonal. Just as for the matrices R_i , it is therefore sufficient to precompute the matrices M_i^j for the first $n_\theta/2$ expansion points, since the remaining matrices can be reconstructed by symmetry operations.

The precomputation of the matrices M_i^j hides the cost κ^2 related to the upsampling from the nearly singular evaluation, though the $\mathcal{O}(p^2)$ expansion order cost still affects the evaluation and storage cost. This is because it reflects the number of expansion coefficients used, and we need all the coefficients to be able to evaluate the expansion at arbitrary target points. Only when precomputing for a specific target point, as in the self-interaction, can we get rid of the p^2 factor.

4.3. Storage complexity

Precomputing the density-to-coefficient matrices M_i^j for a single expansion center \mathbf{c}_i requires storage of $12N_p N$ complex numbers, while the density-to-potential matrix R_i related to a target point \mathbf{x}_i only requires storage of $9N$ real numbers. Considering a spheroidal body, where $N = n_\varphi n_\theta$, and taking both symmetries into account, the storage required for computing all necessary M_i^j is then $6N_p n_\theta^2 n_\varphi$ complex numbers, while the storage required for all necessary R_i is $\frac{9}{2} n_\theta^2 n_\varphi$ real numbers.

As an example, using a set of 50×50 quadrature points and $p = 15$ requires approximately 1.6 GB for the M_i^j and 4.5 MB for the R_i , using double precision. The number of expansion coefficients scales as $N_p = \mathcal{O}(p^2)$, so assuming $n_\varphi \propto n_\theta = n$ allows us to write the scaling of these expressions as $\mathcal{O}(p^2 n^3)$ and $\mathcal{O}(n^3)$. Alternatively, introducing a grid size $h \propto 1/n$ gives $\mathcal{O}(p^2 h^{-3})$ and $\mathcal{O}(h^{-3})$.

4.4. Summary of precomputation scheme

The precomputation scheme outlined above minimizes the cost for QBX when computing it directly, as opposed to accelerating it by an FMM. For a given setup of particle shape, quadrature grid and QBX parameters, we can precompute the matrices R_i and M_i^j , associated with self-interaction and nearly singular evaluation, and store them to disk. The symmetries of the spheroidal grid allow us to do this at a manageable cost in terms of memory. Once we have these precomputed values, we can load and use them for simulations with a large number of particles that all have identical shape and grid. The single set of precomputed values can then be used for all particles in the simulation, since they will differ from each other only by a rigid body rotation.

For the self-interaction, the costs of upsampling and expansion order are hidden from the evaluation step by the precomputations. We can therefore pick κ and p large enough to make sure that the accuracy in the evaluation of the double layer potential is only limited by how well the density \mathbf{q} is resolved on the underlying quadrature grid. For the nearly singular evaluation the expansion order p still affects the computational cost, so there is still a cost/accuracy tradeoff there. However, the upsampling factor κ is hidden by the precomputations, so we can always make sure that we pick it large enough to resolve integrands associated with large p .

5. Periodic systems and Ewald summation

To model an infinitely large particle system while minimizing the effects of a finite-size computational domain, we use the common method of applying periodic boundary

conditions. We then consider a system of M particles contained in a box of size $L_1 \times L_2 \times L_3$, which is periodically replicated in all three spatial directions. The velocity field is then periodic,

$$\mathbf{u}^P(\mathbf{x}) = \mathbf{u}^P(\mathbf{x} + \boldsymbol{\tau}), \quad (62)$$

where the superscript P denotes the periodicity and $\boldsymbol{\tau}$ is a periodic shift,

$$\boldsymbol{\tau} = (p_1 L_1, p_2 L_2, p_3 L_3), \quad \mathbf{p} \in \mathbb{Z}^3. \quad (63)$$

The periodic velocity field is computed as a superposition of the field from all periodic replications of all particles,

$$\mathbf{u}^P(\mathbf{x}) = \sum_{\alpha=1}^M [\mathcal{D}^P[\Gamma_\alpha, \mathbf{q}](\mathbf{x}) + \mathcal{V}^{P,\alpha}(\mathbf{x})], \quad (64)$$

where \mathcal{D}^P is the periodic double layer potential,

$$\mathcal{D}_i^P[\Gamma, \mathbf{q}](\mathbf{x}) = \sum_{\boldsymbol{\tau}} \int_{\Gamma} T_{ijk}(\mathbf{x}, \mathbf{y} + \boldsymbol{\tau}) q_j(\mathbf{y}) n_k(\mathbf{y}) dS_{\mathbf{y}}, \quad (65)$$

and $\mathcal{V}^{P,\alpha}$ is the periodic completion flow,

$$\mathcal{V}_i^{P,\alpha}(\mathbf{x}) = \sum_{\boldsymbol{\tau}} \mathcal{V}_i^\alpha(\mathbf{x} + \boldsymbol{\tau}). \quad (66)$$

Computing \mathbf{u}^P by direct summation over all periodic images is not feasible since the sums over $\boldsymbol{\tau}$ in (65) and (66) decay slowly and depend on the order of summation. Additionally, the convergence of the sum over stokeslets depends on the assumption that there is a pressure gradient maintaining the force balance. Here we compute the sums using a fast Ewald summation method [2], which corresponds to using a spherical summation order. The stresslet T in (65) is then decomposed into two parts,

$$T(\mathbf{x}, \mathbf{y}) = T^{(R)}(\mathbf{x}, \mathbf{y}, \xi) + T^{(F)}(\mathbf{x}, \mathbf{y}, \xi), \quad (67)$$

where $T^{(R)}$ is short-ranged and singular, while $T^{(F)}$ is long-ranged and smooth. The Ewald parameter ξ determines how short-ranged and smooth the two respective parts are. Away from the source \mathbf{y} the short-ranged part $T^{(R)}$ decays exponentially,

$$T^{(R)}(\mathbf{x}, \mathbf{y}, \xi) \sim e^{-\xi^2 |\mathbf{x} - \mathbf{y}|^2}. \quad (68)$$

This makes it possible to introduce a truncation radius r_c and only consider near-neighbor interactions within that radius ($|\mathbf{x} - \mathbf{y}| \leq r_c$). The long-ranged part $T^{(F)}$ converges rapidly when computed in Fourier space, where its Fourier coefficients $\widehat{T}^{(F)}(\mathbf{k}, \xi)$ decay exponentially due to its smoothness,

$$\widehat{T}^{(F)}(\mathbf{k}, \xi) \sim e^{-k^2/4\xi^2}. \quad (69)$$

The Fourier expansion is truncated at a maximum wave number $K = \max_{\mathbf{k}} \|\mathbf{k}\|_\infty$ and computed using the spectral Ewald (SE) method [23, 24], which uses a fast Fourier transform (FFT) to compute the long-range interactions to spectral accuracy. In [2] we describe in detail how this is done for the stresslet, and provide accurate truncation error estimates which can be used for selecting the Ewald parameter ξ and the cutoffs r_c and K . The equivalent development for the Stokeslet and rotlet singularities are available in [23] and [1]. Computing $T^{(F)}$ using the SE method and only computing $T^{(R)}$ for near neighbors gives a method with $\mathcal{O}(N \log N)$ complexity, where N is the total number of degrees of freedom when the number of particles in the system is increased, under the assumption that the average number of near neighbors of each discrete point is kept constant [2].

5.1. QBX and Ewald summation

In section 3 we developed a framework for computing the double layer potential (8) to high accuracy for both the singular and nearly singular cases, based on an expansion for the stresslet $T = r_i r_j r_k / r^5$. When computing the periodic double layer potential (65) using Ewald summation, we have split the stresslet into a short-ranged part $T^{(R)}$ and a long-ranged part $T^{(F)}$ (67). The explicit form of the short-ranged part is [2]

$$\begin{aligned} T_{jlm}^{(R)}(\mathbf{x}, \mathbf{y}, \xi) = & -\frac{2}{r^4} \left(\frac{3}{r} \operatorname{erfc}(\xi r) + \frac{2\xi}{\sqrt{\pi}} (3 + 2\xi^2 r^2 - 4\xi^4 r^4) e^{-\xi^2 r^2} \right) r_j r_l r_m \\ & + \frac{8\xi^3}{\sqrt{\pi}} (2 - \xi^2 r^2) e^{-\xi^2 r^2} (\delta_{jl} r_m + \delta_{lm} r_j + \delta_{mj} r_l), \end{aligned} \quad (70)$$

where $\mathbf{r} = \mathbf{x} - \mathbf{y}$ and $r = |\mathbf{r}|$. All of the singular behavior of the stresslet is contained in $T^{(R)}$, since a series expansion of $T^{(R)}$ in the limit of $r \rightarrow 0$ shows that

$$\lim_{\mathbf{x} \rightarrow \mathbf{y}} \left[T(\mathbf{x}, \mathbf{y}) - T^{(R)}(\mathbf{x}, \mathbf{y}, \xi) \right] = 0. \quad (71)$$

This means that we need a way of using QBX for evaluating the short-ranged potential, which we define as

$$\mathcal{D}_i^{(R)}[\Gamma, \mathbf{q}](\mathbf{x}) = \int_{\Gamma} T_{ijk}^{(R)}(\mathbf{x}, \mathbf{y}, \xi) q_j(\mathbf{y}) n_k(\mathbf{y}) dS_{\mathbf{y}}. \quad (72)$$

Deriving a local expansion for $T^{(R)}$, as defined in (70), is a daunting proposition. Instead, we rewrite the short-ranged potential as

$$\mathcal{D}_i^{(R)}[\Gamma, \mathbf{q}](\mathbf{x}) = \int_{\Gamma} T_{ijk}(\mathbf{x}, \mathbf{y}) q_j(\mathbf{y}) n_k(\mathbf{y}) dS_{\mathbf{y}} \quad (73)$$

$$- \int_{\Gamma} \underbrace{\left[T_{ijk} - T_{ijk}^{(R)} \right]}_{T_{ijk}^{(F)}}(\mathbf{x}, \mathbf{y}, \xi) q_j(\mathbf{y}) n_k(\mathbf{y}) dS_{\mathbf{y}}. \quad (74)$$

The first term is now the ordinary double layer potential (8), for which we can use our existing QBX framework. The second term is the smooth and long-ranged part of the potential, which we can evaluate using our direct quadrature rule. In our implementation we use (74) together with QBX if \mathbf{x} is close to or on Γ , and (72) truncated at r_c otherwise.

5.2. Computational complexity

We now consider a system of M particles with N degrees of freedom each, which we solve using our accelerated QBX scheme combined with SE. The computational complexity of solving this system then has three distinct components: (i) the interactions of particles with themselves (the self-interactions), (ii) the nearly singular interactions between particles, and (iii) the well-separated interactions. The self-interactions are completely precomputed, and can therefore be evaluated at a $\mathcal{O}(MN^2)$ cost. The nearly singular interactions use the precomputed M_i^j -matrices, and contribute with a $\mathcal{O}(p^2MN^2)$ cost, p being the QBX expansion order and the leading constant depending strongly on the interparticle spacing. The well-separated interactions are handled by SE, at a cost which scales as $\mathcal{O}(MN \log(MN))$ as we increase the size of the particle system. This puts the total computational complexity at

$$\mathcal{O}(MN^2) + \mathcal{O}(p^2MN^2) + \mathcal{O}(MN \log(MN)). \quad (75)$$

For a given particle system with fixed N , p and average interparticle spacing, the cost then grows as $\mathcal{O}(M \log M)$ as the system size is increased by increasing the number of particles M .

6. Numerical results

We can now summarize our numerical scheme as follows:

- **Spheroidal grid:** For a spheroidal body with axes a and c , we introduce a spheroidal grid (section 3.1) of dimensions $N = n_\theta \times n_\varphi$. For a system with M particles, an identical grid is used for each particle.
- **QBX:** We place expansion centers \mathbf{c}_i (42) at a normal distance r from the grid points \mathbf{x}_i , and use an order p local expansion (section 3.2) to evaluate the double layer potential (8) in the vicinity of each expansion center. The expansion is computed using a factor κ upsampled grid. The QBX computations are accelerated by the precomputation scheme of section 4.
- **Spectral Ewald:** For periodic problems, we compute the velocity field by combining QBX with a fast and spectrally accurate Ewald summation method (section 5).
- **Solution:** We solve the discrete system (26) with $3MN$ unknowns iteratively using GMRES, with the left hand side computed using QBX and Ewald.
- **Preconditioning:** For systems with many particles, we construct a simple block-diagonal preconditioner by computing the explicit inverse of the corresponding single-particle system and using it to create the diagonal blocks (through rotation). This requires $(3N)^2$ memory storage, but reduces the number of iterations by as much as a factor three for some problems.

This scheme allows us to efficiently compute the Stokes flow around systems of particles where the particles are very close to each other, and also to accurately evaluate the flow everywhere in the domain. An example is shown in Figure 6, where one can clearly see the benefit using QBX when evaluating the potential close to the particles.

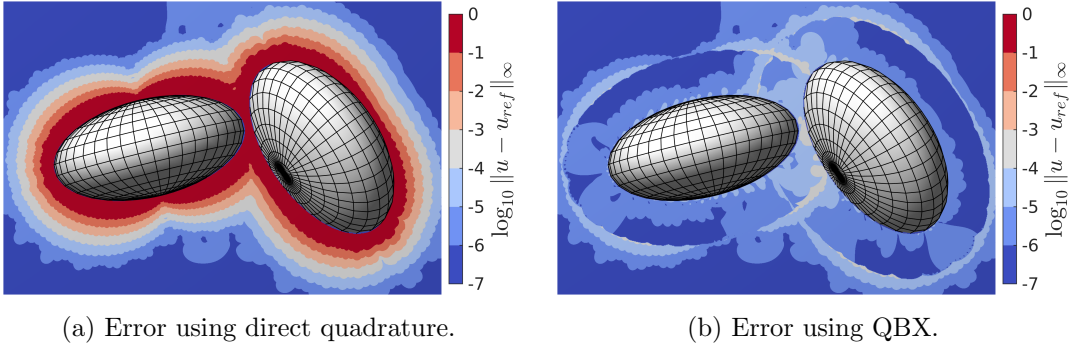


Figure 6: Example of quadrature error (vs. a numerically obtained reference solution) when evaluating the flow around two close particles, with and without QBX. The particles are oblates, $a/c = 2$, with a 20×40 grid and QBX parameters $r/h = 1.5$, $\kappa = 10$ and $p = 30$.

The method has currently been implemented in MATLAB, with time-critical components written in C and called through the MEX interface. Currently only the SE implementation is available as open source [22], though the authors plan to release the full implementation at a future date. All tests are run on a regular workstation with a 3.4 GHz quad-core Intel Haswell CPU. For the large system in section 6.2 (131 particles, 1800 points/particle, $p = 30$) singular interactions are evaluated at about 1000 targets/second and nearly singular interactions at about 250 targets/second. A full matvec including the local QBX treatment and the spectral Ewald method on the full set of points takes 2–3 minutes.

6.1. Validation

To validate that our method is correct, we have compared it against several analytical, semi-analytical and numerical results. For a single particle in free space, analytic results are available in Chwang & Wu [12, 11] and Jeffery [19]. Figure 4 shows an example of a comparison against [12]. For systems of three spheres in free space, semi-analytic results for several configurations are available in Wilson [36]. For a periodic array of spheres, numeric results are available in Zick & Homsy [39] and Sangani & Acrivos [32]. The described method can solve all these problems to arbitrary accuracy, provided that the grid is fine enough to represent the solution, and that the QBX parameter selection of section 3.7 is followed. We will below discuss results for a three-sphere system and for a periodic array of spheres.

Our implementation of the method has also been used in a separate work by Bagge [5], where it was thoroughly validated against the solutions by Chwang [11] and Jeffery [19].

6.1.1. Triangle of spheres

To test the ability of our method to resolve particle interactions in the lubrication limit, we set up the following test case from Wilson [36]: Three spheres of radius a are arranged in an equilateral triangle with sides sa , $s > 2$. A force of magnitude $6\mu\pi a$ is applied to sphere 1 towards the centroid of triangle, pushing the sphere into the other two spheres. An example of this geometry for $s = 2.1$ is shown in Figure 7. The following quantities are measured:

- Velocity U_1 of sphere 1, in direction of force.
- Velocity U_2 of spheres 2 and 3, in direction of force.
- Velocity U_3 of spheres 2 and 3, in direction perpendicular to force.
- Angular velocity $a\Omega$ of spheres 2 and 3.

s	Wilson				Current method			
	U_1	U_2	U_3	Ω	U_1	U_2	U_3	Ω
2.01	0.65528	0.63461	0.00498	0.037336	0.655 <u>01</u>	0.634 <u>24</u>	0.00 <u>201</u>	0.03 <u>6978</u>
2.10	0.73857	0.59718	0.03517	0.052035	0.73857	0.59718	0.03517	0.0520 <u>36</u>
2.50	0.87765	0.49545	0.07393	0.045466	0.87765	0.49545	0.07393	0.045466
3.00	0.93905	0.41694	0.07824	0.035022	0.93905	0.41694	0.07824	0.035022

Table 1: Translational and angular velocities for triangle of spheres, computed using 64×64 grid, $r/h = 3$, $\kappa = 5$ and $p = 20$. Deviations from reference results by Wilson [36] are underlined.

As can be seen in Table 1, we are able to reproduce the results by Wilson to high accuracy for $s \geq 2.10$ using a 64×64 grid. For $s = 2.01$ the results are inaccurate even though the quadrature is accurately resolved everywhere, and a finer grid is required if the results are to be improved. This is because the double layer density \mathbf{q} gets sharply peaked as the particles get close to each other, with a peak that has spatial scale $\mathcal{O}(d^{1/2})$ when the separation distance is d [6, 33], and we can not hope to get good results if our grid is too coarse to resolve the density. The problem is therefore one of resolution, not quadrature. To solve the problem efficiently as $s \rightarrow 2$, an adaptive refinement strategy for the surface grid would be required.

6.1.2. Periodic array of spheres

As a validation test for our periodic computations, we consider a periodic array of spheres of radius a arranged in a simple cubic lattice. This is modeled using a single sphere in a cubic box with sides L . We subject the sphere to an external force, and solve the mobility problem to get its velocity for various values of the concentration $\rho = 4\pi a^3/3L^3$. We then compute the dimensionless drag coefficient $K = F/6\pi aU$ and compare it to the results available from Sangani & Acrivos [32]. Using a 30×30 grid and





Figure 7: Flow velocity around triangle of spheres at configuration $s = 2.1$. Sphere 1 is the top sphere. A feature of boundary integral methods is that once the solution \mathbf{q} is known, it is possible to zoom in and study the details of the flow field everywhere, as is done here in the right picture.

tuning all parameters (QBX+Ewald) for a tolerance of $\epsilon = 10^{-7}$, we get the results shown in Figure 8. Our results are identical to the reference results in all the digits presented in [32], all the way up to the limit $\rho = \rho_{\max}$, where the spheres touch ($L = 2a$). Our results for K also agree in all six decimals with the series expansion in $\chi = (\rho/\rho_{\max})^{1/3}$ published in [32, eq. 60], up to and including $\chi = 0.5$, after which the series expansion loses precision.

6.2. Application: Porous media flow

As a demonstration of a possible application for our method, we consider a periodic cell of dimensions $40 \times 15 \times 15$ containing 131 randomly positioned, oblate particles with size $(a, c) = (2, 1)$, for a volume particle concentration of about 24%. This system can be used as a model of a three-dimensional porous medium with non-spherical grains. To this end, we set a background flow $\mathbf{u}_{\text{bg}} = (1, 0, 0)$ and solve the resistance problem for fixed particle positions. We use a 30×60 grid and tune our parameters for a tolerance $\epsilon = 10^{-4}$. The resulting system has 707,400 unknowns and is solved in about 2 hours using our current implementation and hardware.

Once we have the solution \mathbf{q} on the particles, it is easy to evaluate the velocity field anywhere in the domain. We can then place a Lagrangian point randomly in the domain and let it be advected by the flow, using a high-order ODE solver (MATLAB's `ode45`). Doing this for a few particles as they pass through the domain allows us to draw the streamlines of Figure 9a. Doing this for many (~ 100) particles over a long time interval ($t \in [0, 10^3]$), it is possible to extract statistics about the flow. We shall here perform an




(a) Simple cubic array.

χ	K	K_{ref}
0.3	1.699884	1.7000
0.4	2.151801	2.1518
0.5	2.842022	2.8420
0.6	3.973781	3.9738
0.7	6.004034	6.004
0.8	10.054098	10.05
0.9	19.161078	19.16
0.95	27.918287	27.9
1.00	42.102343	42.1


(b) Dimensionless drag.

Figure 8: Dimensionless drag coefficient K for a simple cubic array of spheres, compared to reference results from [32]. The dimensionless parameter $\chi = (\rho/\rho_{\text{max}})^{1/3}$ measures how close the configuration is to the touching configuration ($\chi = 1$).

analysis that closely follows that by de Anna et al. [13] for a two-dimensional medium with circular grains. Our intention here is not to draw any new conclusions about this particular case, only to show that our method makes the analysis possible and that the results correspond to those in [13].



(a)



(b)

Figure 9: (a) Streamlines around a periodic system with 131 oblate particles. (b) Mean-squared displacement in longitudinal (σ_x) and transversal (σ_{yz}) directions. The transition from linear to square root (Fickian) scaling appears to begin around $t = t_A$.

The hydrodynamic particle dispersion can be measured through the particle displacement $\Delta x_i(t) = x_i(t) - x_i(0)$. Computing the variance of the displacement², $\sigma_{x_i}^2(t) = \langle [\Delta x_i(t) - \langle \Delta x_i(t) \rangle]^2 \rangle$, we can study the longitudinal displacement in σ_x^2 and the transver-

²Brackets denote average over all Lagrangian particles.

sal displacement in $\sigma_{yz}^2 := \sigma_y^2 + \sigma_z^2$. These are shown in Figure 9b, with time normalized by $t_A = \bar{d}/\bar{v}$, the mean particle diameter over the mean velocity. Both σ_x and σ_{yz} can be seen to have an initial linear growth, until entering a transition regime beginning at $t = t_A$, after which the growth approaches a rate proportional to $t^{1/2}$ (this is known as the Fickian regime). This is consistent with [13].




Figure 10: Probability density function of the normalized longitudinal Lagrangian velocity increments $\Delta_\tau u/\sigma_{\Delta u}(\tau)$ when $\tau < t_A$ and $\tau > t_A$. Clearly, a transition occurs around $\tau = t_A$.

For a given time lag τ , the longitudinal Lagrangian velocity increment is defined as $\Delta_\tau u = u(t + \tau) - u(t)$. The probability density function of the velocity increment, normalized by its standard deviation, is shown in Figure 10 for two different time lags. For large time lags ($\tau \gg t_A$) the distribution approaches the Gaussian distribution, while for small time lags ($\tau < t_A$) the distribution is similar to that of the correlated continuous time random walk (CCTRW) model in de Anna [13].

7. Conclusions

We have in this paper presented a complete boundary integral framework for simulating periodic Stokes flow around spheroidal bodies. It is based on representing the flow using the Stokes double layer potential, which results in a well-conditioned system that converges rapidly when solved using GMRES. Singular and nearly singular integrals are computed using QBX, which we have adapted for the double layer potential. This allows us to evaluate the velocity to high accuracy everywhere in domain, such that we can have accurate quadrature also for nearly touching geometries. By exploiting the symmetries of the spheroidal bodies, we have been able to develop a precomputation scheme that minimizes storage requirements, while allowing us to rapidly compute the potential using QBX. By integrating the precomputed QBX with a fast and spectrally accurate Ewald summation method, we get a fast method which scales favorably as we go to systems with many particles.

The results presented in section 6 show that our method is able to accurately solve reference problems available in the literature. The example involving a porous media model also indicates that our method could be useful when studying such models, with both spherical and non-spherical grains. A natural application extension of our method

would be to use it also for moving rigid particles, for example by studying sedimentation in particle suspensions, as was done in [2]. Then particles can come arbitrary close to each other, so close that lubrication forces start to dominate. This is not an issue for our quadrature, but may cause the double layer density to be sharply peaked, as we saw in the example in section 6.1.1. To deal with such cases, one could either (i) introduce an adaptive surface quadrature, that is able to resolve the density peak using a limited number of surface nodes, or (ii) couple the method to a lubrication model that removes the need for the double layer potential to resolve the strong interactions [21, 25, 26, 33].

Acknowledgments

This work has been supported by the Göran Gustafsson Foundation for Research in Natural Sciences and Medicine and by the Swedish Research Council under grant no. 2011-3178. The authors gratefully acknowledge this support.

Appendix A Precomputing QBX on a spheroidal grid

In this appendix we develop the details of how to accelerate QBX for the double layer potential on a spheroidal grid by exploiting symmetries of the grid. To derive the expressions that we need, we first consider acceleration on a spheroidal grid for the Laplace potentials.

A.1 Symmetries in precomputed layer potentials

We begin by considering a general, scalar layer potential from a radially symmetric kernel K , integrated over a surface Γ together with a scalar density q ,

$$u(\mathbf{x}) = \int_{\Gamma} K(\mathbf{x} - \mathbf{y})q(\mathbf{y})dS_{\mathbf{y}}. \quad (\text{A.1})$$

The surface is discretized using a spheroidal grid, as defined in section 3.1, which gives us $N = n_{\theta}n_{\varphi}$ points $\mathbf{x}_1, \dots, \mathbf{x}_N$ on the surface. At each point we have a discrete density $q_i = q(\mathbf{x}_i)$. All density values are stored in the vector $\mathbf{Q} \in \mathbb{R}^N$,

$$\mathbf{Q} = \begin{pmatrix} q(\mathbf{x}_1) \\ \vdots \\ q(\mathbf{x}_N) \end{pmatrix} = \begin{pmatrix} q(\theta_1, \varphi_1) \\ q(\theta_2, \varphi_1) \\ \vdots \\ q(\theta_{n_{\theta}}, \varphi_1) \\ q(\theta_1, \varphi_2) \\ \vdots \\ q(\theta_{n_{\theta}}, \varphi_{n_{\varphi}}) \end{pmatrix}. \quad (\text{A.2})$$

The discrete approximation of the layer potential is then

$$u^h(\mathbf{x}) = \sum_{i=1}^N K(\mathbf{x} - \mathbf{x}_i)q_iw_i. \quad (\text{A.3})$$

When solving a boundary integral equation, we typically need to compute the potential at all the points \mathbf{x}_i on the surface,

$$u^h(\mathbf{x}_i) = \sum_{j=1}^N K(\mathbf{x}_i - \mathbf{x}_j) q_j w_j, \quad i = 1, \dots, N, \quad (\text{A.4})$$

which we can write more compactly as

$$\mathbf{U} = \mathbf{A}\mathbf{Q}, \quad (\text{A.5})$$

where $\mathbf{A} \in \mathbb{R}^{N \times N}$ is a (typically) full matrix and $\mathbf{U} \in \mathbb{R}^N$ is the layer potential at all the points on the surface, $U_i = u(\mathbf{x}_i)$.

A.1.1 Rotational symmetry

Now, since the spheroid is rotationally symmetric and we use the trapezoidal rule with uniform spacing in the azimuthal direction, the grid will look identical from the point of view of all points which lie on the same latitude (θ). If we denote the rows of \mathbf{A} by R_i , $i = 1, \dots, N$, then this must mean that rows i and $i + n_\theta$ contains the same information, and differ only by a permutation. This can be viewed as R_i and R_{i+n_θ} being stencils that cover the spheroidal grid in the θ - φ -plane. Having their target point at the same θ -coordinate they then have identical coefficients, but they differ by a permutation due to the periodic wrap-around at $\phi = 2\pi$.

We introduce the permutation $\tau_n(\mathbf{Q})$ ³,

$$(\tau_n(\mathbf{Q}))_i = Q_{\sigma_n(i)}, \quad \sigma_n(i) = (i + n - 1 \pmod{N}) + 1, \quad (\text{A.6})$$

such that

$$\tau_n(\mathbf{Q}) = \begin{pmatrix} Q_{1+n} \\ \vdots \\ Q_N \\ Q_1 \\ \vdots \\ Q_n \end{pmatrix}. \quad (\text{A.7})$$

We represent this permutation using a (sparse) $N \times N$ matrix P_n ,

$$\tau_n(\mathbf{Q}) = P_n \mathbf{Q}. \quad (\text{A.8})$$

For $1 \leq i \leq n_\theta$ we then have that

$$u(\mathbf{x}_i) = R_i \mathbf{q}, \quad (\text{A.9})$$

$$u(\mathbf{x}_{i+n_\theta}) = R_i P_{n_\theta} \mathbf{q}, \quad (\text{A.10})$$

$$u(\mathbf{x}_{i+2n_\theta}) = R_i P_{n_\theta}^2 \mathbf{q}, \quad (\text{A.11})$$

$$\dots, \quad (\text{A.12})$$

³ $\text{mod}(a, b)$ refers to the operation *a modulo b*, which returns the remainder of the division of *a* by *b*, s.t. $0 \leq \text{mod}(a, b) < b$.

which we can summarize as

$$R_{i+\alpha n_\theta} = R_i P_{n_\theta}^\alpha, \quad i \in [1, n_\theta]. \quad (\text{A.13})$$

So we only need to store the first n_θ rows of A , as the remaining ones are given by successive permutations of that set. This reduces our memory use from $N \times N$ to $n_\theta \times N$.

A.1.2 Mirror symmetry

For the rotational symmetry we could show that we only need to store coefficients related to the first n_θ points in order to make computations for all $n_\theta \times n_\varphi$ points. An additional reduction can be made by observing that the spheroidal grid also has a mirror symmetry around its equator. This means that we can use coefficients related to target points on the northern hemisphere to make computations for target points on the southern hemisphere. Hence, we only need to store the first $n_\theta/2$ rows of the matrix A in memory to be able to compute \mathbf{U} from \mathbf{Q} .

Mathematically this is done for the current ordering of points by introducing a mirroring permutation matrix F_{n_θ} such that

$$F_{n_\theta} \mathbf{Q} = \begin{pmatrix} Q_{n_\theta} \\ \vdots \\ Q_1 \\ Q_N \\ \vdots \\ Q_{n_\theta+1} \end{pmatrix}. \quad (\text{A.14})$$

We can then recover rows $n_\theta/2 + 1, \dots, n_\theta$ by using rows $1, \dots, n_\theta/2$,

$$R_i = R_{n_\theta-i+1} F_{n_\theta}, \quad i \in [n_\theta/2 + 1, n_\theta], \quad (\text{A.15})$$

and from there recover all the rows of the matrix using the rotational symmetry. For notational simplicity we here assume n_θ to be even, extending to odd n_θ is trivial.

A.2 Precomputing QBX

Now we have established the symmetries available when evaluating a layer potential on a spheroidal grid. We will now see how those can be used in conjunction with QBX, to create a fast method.

A.2.1 Laplace single layer potential

We begin by considering the Laplace single layer potential, having kernel $K(\mathbf{r}) = |\mathbf{r}|^{-1}$, as the analysis is simplified by all the involved quantities being scalar. Using the Laplace expansion (31) about an expansion center \mathbf{c} , the potential

$$u(\mathbf{x}) = \int_{\Gamma} \frac{q(\mathbf{y})}{|\mathbf{x} - \mathbf{y}|} dS_{\mathbf{y}}, \quad (\text{A.16})$$

can be written as

$$u(\mathbf{x}) = \sum_{l=0}^{\infty} \sum_{m=-l}^l r_x^l Y_l^{-m}(\theta_x, \varphi_x) z_{lm}(\mathbf{c}), \quad (\text{A.17})$$

where the expansion coefficients are computed as

$$z_{lm}(\mathbf{c}) = \frac{4\pi}{2l+1} \int_{\Gamma} \frac{q(\mathbf{y})}{r_y^{l+1}} Y_l^m(\theta_y, \varphi_y) dS_{\mathbf{y}}. \quad (\text{A.18})$$

When this is done discretely on a spheroidal grid, there is a series of steps involved. The first step is to create a fine grid by increasing the number of grid points by a factor κ in each direction, so that we get a grid with points $\tilde{\mathbf{x}}_i$, $i = 1, \dots, \kappa^2 N$. The density is upsampled to the fine grid by a suitable interpolation scheme, which is represented as the matrix operation

$$\tilde{\mathbf{Q}} = U\mathbf{Q}, \quad U \in \mathbb{R}^{\kappa^2 N \times N}. \quad (\text{A.19})$$

In the second step, we use the upsampled density to compute the local expansion coefficients up to order p at expansion points \mathbf{c}_i (located at a distance r in the normal direction from \mathbf{x}_i),

$$z_{lm}(\mathbf{c}_i) = \sum_{n=1}^{\kappa^2 N} \tilde{q}_n \tilde{w}_n \frac{1}{r_n^{l+1}} Y_l^m(\theta_n, \varphi_n), \quad (\text{A.20})$$

where \tilde{w}_n is the quadrature weight at $\tilde{\mathbf{x}}_n$ and $(r_n, \theta_n, \varphi_n) = \tilde{\mathbf{x}}_n - \mathbf{c}$. This gives us for each point \mathbf{c}_i a coefficient vector $\mathbf{z}(\mathbf{c}_i) \in \mathbb{C}^{N_p}$,

$$\mathbf{z}(\mathbf{c}_i) = \begin{pmatrix} z_{0,0}(\mathbf{c}_i) \\ z_{1,0}(\mathbf{c}_i) \\ z_{1,1}(\mathbf{c}_i) \\ \vdots \\ z_{p,p}(\mathbf{c}_i) \end{pmatrix}, \quad (\text{A.21})$$

where $N_p = (p^2 + 3p + 2)/2$ is the total number of coefficients at each expansion center that we need to store⁴. In matrix notation we write this as

$$\mathbf{z}(\mathbf{c}_i) = \mathcal{S}_i \tilde{\mathbf{Q}} = \mathcal{S}_i U \mathbf{Q} = M_i \mathbf{Q}, \quad (\text{A.22})$$

where

$$\mathcal{S}_i \in \mathbb{C}^{N_p \times \kappa^2 N}, \quad (\text{A.23})$$

$$M_i \in \mathbb{C}^{N_p \times N}. \quad (\text{A.24})$$

⁴As mentioned in section 3.3, the coefficients for negative m need not be stored, since $z_{l,-m} = z_{lm}^*$ for a real density.

As the third step, we let the vector $\mathbf{s}(\mathbf{x}_i) \in \mathbb{C}^{N_p}$ contain the coefficients for evaluating the expansion at \mathbf{c}_i ,

$$\mathbf{s}(\mathbf{x}_i) = \begin{pmatrix} s_{0,0}(\mathbf{x}_i) \\ s_{1,0}(\mathbf{x}_i) \\ s_{1,1}(\mathbf{x}_i) \\ \vdots \\ s_{p,p}(\mathbf{x}_i) \end{pmatrix}, \quad s_{lm}(\mathbf{x}_i) = \begin{cases} \frac{1}{2} r_x^l Y_l^m(\theta_x, \varphi_x), & \text{if } m = 0, \\ r_x^l Y_l^m(\theta_x, \varphi_x) & \text{if } m > 0 \end{cases}, \quad (\text{A.25})$$

such that, using $(Y_l^m)^* = Y_l^{-m}$, we can write

$$u(\mathbf{x}_i) = \sum_{l=0}^p \sum_{m=0}^l (s_{lm}(\mathbf{x}_i)^* z_{lm}(\mathbf{c}) + s_{lm}(\mathbf{x}_i) z_{lm}(\mathbf{c})^*). \quad (\text{A.26})$$

Defining the conjugate product

$$\langle \mathbf{x}, \mathbf{y} \rangle = (\mathbf{x}^*)^T \mathbf{y} + \mathbf{x}^T \mathbf{y}^*, \quad (\text{A.27})$$

we can write this as

$$u(\mathbf{x}_i) = \langle \mathbf{s}(\mathbf{x}_i), \mathbf{z}(\mathbf{c}_i) \rangle. \quad (\text{A.28})$$

Putting (A.22) and (A.28) together, we can conclude that the matrix R_i , such that $u(\mathbf{x}_i) = R_i \mathbf{q}$, is given by

$$R_i = \mathbf{s}(\mathbf{x}_i)^{*T} M_i + \mathbf{s}(\mathbf{x}_i)^T M_i^*, \quad i = 1, \dots, N. \quad (\text{A.29})$$

This R_i has the symmetries discussed in section A.1. From rotation and mirroring when then have

$$R_i = R_{n_\theta - i + 1} F_{n_\theta}, \quad i \in [n_\theta/2 + 1, n_\theta], \quad (\text{A.30})$$

$$R_{i + \alpha n_\theta} = R_i P_{n_\theta}^\alpha, \quad i \in [1, n_\theta], \quad (\text{A.31})$$

such that it is enough to directly compute $R_i, \dots, R_{n_\theta/2}$, as the remaining R_i can be recovered through these operations.

The symmetries that apply to the rows R_i also apply to the matrices M_i , such that we can compute the coefficients at all expansion centers using M_i for the first n_θ expansion centers. Here, however, it is not enough to permute the right hand side vector; we also need to account for the phase of the factor $e^{im\varphi}$ in the spherical harmonic. An expansion center \mathbf{c}_{i+n_θ} will "see" all the grid points in the same angles φ as an expansion center \mathbf{c}_i does, plus a rotation of $\Delta\varphi$. To account for this rotation, we introduce a phase matrix $E \in \mathbb{C}^{N_p \times N_p}$ with diagonal elements

$$E_{jj} = e^{im\Delta\varphi}. \quad (\text{A.32})$$

For $1 \leq i \leq n_\theta$ this lets us write

$$\mathbf{z}(\mathbf{c}_i) = M_i \mathbf{Q}, \quad (\text{A.33})$$

$$\mathbf{z}(\mathbf{c}_{i+n_\theta}) = E M_i P_{n_\theta} \mathbf{Q}, \quad (\text{A.34})$$

$$\mathbf{z}(\mathbf{c}_{i+2n_\theta}) = E^2 M_i P_{n_\theta}^2 \mathbf{Q}, \quad (\text{A.35})$$

$$\dots, \quad (\text{A.36})$$

and

$$\mathbf{s}(\mathbf{x}_{i+n_\theta}) = \mathbf{s}(\mathbf{x}_i) E \quad (\text{A.37})$$

$$\mathbf{s}(\mathbf{x}_{i+2n_\theta}) = \mathbf{s}(\mathbf{x}_i) E^2 \quad (\text{A.38})$$

$$\dots, \quad (\text{A.39})$$

which can be summarized as

$$M_{i+\alpha n_\theta} = E^\alpha M_i P_{n_\theta}^\alpha, \quad (\text{A.40})$$

$$\mathbf{s}(\mathbf{x}_{i+\alpha n_\theta}) = \mathbf{s}(\mathbf{x}_i) E^\alpha. \quad (\text{A.41})$$

From this we can recover the original rotational symmetry of R_i ,

$$R_{i+\alpha n_\theta} = \mathbf{s}(\mathbf{x}_{i+\alpha n_\theta})^* M_{i+\alpha n_\theta} + \mathbf{s}(\mathbf{x}_{i+\alpha n_\theta}) M_{i+\alpha n_\theta}^* \quad (\text{A.42})$$

$$= (\mathbf{s}(\mathbf{x}_i) E^\alpha)^* E^\alpha M_i P_{n_\theta}^\alpha + \mathbf{s}(\mathbf{x}_i) E^\alpha (E^\alpha M_i P_{n_\theta}^\alpha)^* \quad (\text{A.43})$$

$$= \mathbf{s}(\mathbf{x}_i)^* M_i P_{n_\theta}^\alpha + \mathbf{s}(\mathbf{x}_i) M_i^* P_{n_\theta}^\alpha \quad (\text{A.44})$$

$$= R_i P_{n_\theta}^\alpha, \quad (\text{A.45})$$

since $E^* E = E E^* = I$.

To also make use of the mirror symmetry for M_i , we have to keep in mind that the mirroring operation makes an expansion center $\mathbf{c}_{n_\theta/2+i}$ "see" the other grid points in a coordinate system which has been rotated 180° about the x -axis, compared to what \mathbf{c}_i sees. This amounts to θ and φ in (A.20) having counterparts $\theta' = -\theta$ and $\varphi' = -\varphi$ in the mirrored system. Inserting this into the relevant terms in (A.20), we see that

$$Y_l^m(\theta', \varphi') = \tilde{P}_l^m(\cos \theta') e^{im\varphi'} = \tilde{P}_l^m(-\cos \theta) e^{-im\varphi}. \quad (\text{A.46})$$

Changing sign in the exponential is equal to taking the complex conjugate of the entire expression. To handle the change of sign in the argument of the associated Legendre polynomial, we apply the parity relation available in Arfken & Weber [4, p. 776],

$$P_l^m(-x) = (-1)^{l+m} P_l^m(x). \quad (\text{A.47})$$

Defining the diagonal matrix $\Theta \in \mathbb{R}^{N_p \times N_p}$ such that

$$\Theta_{jj} = (-1)^{l+m}, \quad (\text{A.48})$$

where $l = l(j)$ and $m = m(j)$ as $j = 1, \dots, N_p$, we finally arrive at the relation

$$M_i = \Theta M_{n_\theta-i+1}^* P_{n_\theta}, \quad i \in [n_\theta/2 + 1, n_\theta], \quad (\text{A.49})$$

which allows us to compute all expansion coefficients for all expansion points using precomputed results for the first $n_\theta/2$ points.

A.2.2 Dipole potential

We will now extend the above results for the Laplace single layer potential to the dipole potential, defined in (32) for a vector density \mathbf{q} as

$$\mathcal{L}[\Gamma, \mathbf{q}](\mathbf{x}) := \int_{\Gamma} \mathbf{q}(\mathbf{y}) \cdot \nabla_{\mathbf{y}} \frac{1}{|\mathbf{x} - \mathbf{y}|} dS_{\mathbf{y}}. \quad (\text{A.50})$$

A discrete quadrature by expansion for this potential is formed from the expressions (33) and (34). We store the pointwise vector densities in linear form in $\mathbf{Q} \in \mathbb{R}^{3N}$,

$$\mathbf{Q} = \begin{pmatrix} \mathbf{Q}_1 \\ \mathbf{Q}_2 \\ \mathbf{Q}_3 \end{pmatrix}, \quad \text{where} \quad \mathbf{Q}_j = \begin{pmatrix} q_j(\mathbf{x}_1) \\ \vdots \\ q_j(\mathbf{x}_N) \end{pmatrix}. \quad (\text{A.51})$$

Just as for the single layer potential, the expansion coefficients are computed using an upsampled density $\tilde{\mathbf{Q}}$, which here is computed as

$$\tilde{\mathbf{Q}} = (I_3 \otimes U)\mathbf{Q}, \quad (\text{A.52})$$

where \otimes is the Kronecker product. The expansion coefficients are here computed as

$$z_{lm}(\mathbf{c}) = \sum_{n=1}^{\kappa^2 N} \tilde{w}_n \tilde{\mathbf{q}}(\tilde{\mathbf{x}}_n) \cdot \nabla_{\tilde{\mathbf{x}}} \frac{1}{r^{l+1}} \tilde{Y}_l^m(\theta_n, \varphi_n). \quad (\text{A.53})$$

This can be represented in matrix form as

$$\mathbf{z}(\mathbf{c}_i) = \mathcal{S}_i (I_3 \otimes U)\mathbf{Q} = M_i \mathbf{Q}, \quad (\text{A.54})$$

where

$$\mathbf{z}(\mathbf{c}_i) \in \mathbb{C}^{N_p}, \quad (\text{A.55})$$

$$M_i \in \mathbb{C}^{N_p \times 3N}. \quad (\text{A.56})$$

The difference from the single layer case is that the matrix $\mathcal{S}_i \in \mathbb{C}^{N_p \times 3\kappa^2 N}$ here includes the gradient of the spherical harmonic on the surface. Evaluation of the expansion is identical to the single layer case (A.28),

$$u(\mathbf{x}_i) = \langle \mathbf{s}(\mathbf{x}_i), \mathbf{z}(\mathbf{c}_i) \rangle, \quad (\text{A.57})$$

with $\mathbf{s}(\mathbf{x}_i)$ defined as in (A.25), and R_i given by

$$R_i = \mathbf{s}(\mathbf{x}_i)^{*T} M_i + \mathbf{s}(\mathbf{x}_i)^T M_i^*, \quad i = 1, \dots, N. \quad (\text{A.58})$$

Due to the density now being vector-valued, there is difference between the single layer potential and the dipole potential in how we can make use of the grid symmetries. To reuse M_i of the first n_{θ} points by rotational symmetry we must now not only take

the permutation P_{n_θ} into account, but also the rotation of the frame of reference, since \mathbf{Q} now represents a quantity which is pointwise vector-valued. Introducing the rotation matrix T_z ,

$$T_z(\alpha) = \begin{pmatrix} \cos(\alpha) & -\sin(\alpha) & 0 \\ \sin(\alpha) & \cos(\alpha) & 0 \\ 0 & 0 & 1 \end{pmatrix}, \quad (\text{A.59})$$

the rotation of M_i can be formulated as

$$M_{i+\alpha n_\theta} = E^\alpha M_i (T_z(-\Delta\varphi) \otimes P_{n_\theta})^\alpha. \quad (\text{A.60})$$

Similarly, for the mirror symmetry we must take the 180° rotation about the x -axis into account. Representing this by the rotation matrix

$$T_x = \begin{pmatrix} 1 & 0 & 0 \\ 0 & -1 & 0 \\ 0 & 0 & -1 \end{pmatrix}, \quad (\text{A.61})$$

we arrive at the expression

$$M_i = \Theta M_{n_\theta-i+1}^* (T_x \otimes F_{n_\theta}), \quad i \in [n_\theta/2 + 1, n_\theta], \quad (\text{A.62})$$

for the mirror symmetry. These results also carry over to the symmetry relations for R_i ,

$$R_i = R_{n_\theta-i+1} (T_x \otimes F_{n_\theta}), \quad i \in [n_\theta/2 + 1, n_\theta], \quad (\text{A.63})$$

$$R_{i+\alpha n_\theta} = R_i (T_z(-\Delta\varphi) \otimes P_{n_\theta})^\alpha, \quad i \in [1, n_\theta]. \quad (\text{A.64})$$

A.2.3 Gradient of expansion

To compute the gradient of a (single or double layer) potential that is locally represented by a local expansion at \mathbf{c} , we differentiate (A.28) or (A.57) with respect to \mathbf{x} ,

$$\nabla u(\mathbf{x}) = \langle \mathbf{g}(\mathbf{x}), \mathbf{z}(\mathbf{c}) \rangle, \quad (\text{A.65})$$

where $\mathbf{g}(\mathbf{x}) \in \mathbb{C}^{3 \times N_p}$ has elements

$$g_{ij}(\mathbf{x}) = \frac{\partial}{\partial x_i} s_j(\mathbf{x}). \quad (\text{A.66})$$

We can also write

$$\frac{\partial}{\partial x_i} u(\mathbf{x}) = \langle \mathbf{g}_i(\mathbf{x}), \mathbf{z}(\mathbf{c}) \rangle, \quad (\text{A.67})$$

where $\mathbf{g}_i(\mathbf{x}) \in \mathbb{C}^{N_p}$,

$$\mathbf{g}_i(\mathbf{x}) = \frac{\partial}{\partial x_i} \mathbf{s}_i(\mathbf{x}). \quad (\text{A.68})$$

If we want to reuse $\mathbf{g}(\mathbf{x}_i)$ from the first n_θ points for the remainder of the points, we have to take into account that the output is vector-valued and has to be rotated into the correct frame of reference. So, if we have the entire process condensed into the form

$$\nabla u(\mathbf{x}_i) = R_i \mathbf{q}, \quad 1 \leq i \leq n_\theta, \quad (\text{A.69})$$

then

$$\nabla u(\mathbf{x}_{i+\alpha n_\theta}) = T_z^\alpha(\Delta\varphi) R_i P_{n_\theta}^\alpha \mathbf{q}, \quad 1 \leq \alpha < n_\varphi. \quad (\text{A.70})$$

For the mirror symmetry we get

$$\nabla u(\mathbf{x}_{n_\theta/2+i}) = T_x R_{n_\theta/2-i+1} F_{n_\theta}, \quad 1 \leq i \leq n_\theta/2. \quad (\text{A.71})$$

A.2.4 Stokes double layer potential

We now turn to our potential of interest, the Stokes double layer potential (8),

$$\mathcal{D}_i[\Gamma, \mathbf{q}](\mathbf{x}) = \int_\Gamma T_{ijk}(\mathbf{x}, \mathbf{y}) q_j(\mathbf{y}) n_k(\mathbf{y}) dS_{\mathbf{y}}. \quad (\text{A.72})$$

Our goal is now to derive the matrix representations necessary for precomputing QBX for this potential on a spheroidal grid. We have from (38) that we can represent the double layer potential as a combination of dipole potentials,

$$\begin{aligned} \mathcal{D}_i[\Gamma, \mathbf{q}](\mathbf{x}) &= \left(x_j \frac{\partial}{\partial x_i} - \delta_{ij} \right) \mathcal{L}[\Gamma, q_j \mathbf{n} + n_j \mathbf{q}](\mathbf{x}) \\ &\quad - \frac{\partial}{\partial x_i} \mathcal{L}[\Gamma, y_k q_k \mathbf{n} + y_k n_k \mathbf{q}](\mathbf{x}), \end{aligned} \quad (\text{A.73})$$

which means that we can derive the expressions needed for the double layer potential using the results derived so far in sections A.2.1– A.2.3. Taking (A.73) apart, we get

$$\mathcal{D}_i[\Gamma, \mathbf{q}](\mathbf{x}) = \left(x_j \frac{\partial}{\partial x_i} - \delta_{ij} \right) u^j(\mathbf{x}) - \frac{\partial}{\partial x_i} u^4(\mathbf{x}), \quad (\text{A.74})$$

where

$$u^j(\mathbf{x}) = \mathcal{L}[\Gamma, q_j \mathbf{n} + n_j \mathbf{q}](\mathbf{x}), \quad j = 1, \dots, 3, \quad (\text{A.75})$$

$$u^4(\mathbf{x}) = \mathcal{L}[\Gamma, y_k q_k \mathbf{n} + y_k n_k \mathbf{q}](\mathbf{x}). \quad (\text{A.76})$$

When computing these components on a spheroidal grid, we can improve the accuracy by taking into account that the geometrical quantities \mathbf{n} and \mathbf{y} are explicitly known, such that we only need to upsample the density \mathbf{q} . Using (A.54), the expansion coefficients at a center \mathbf{c} for the components u^1 – u^4 are then given by

$$\mathbf{z}^j(\mathbf{c}) = \mathcal{S}(\mathbf{c}) \left[\widetilde{\mathbf{N}}^T U \mathbf{Q}_j + \left(I_3 \otimes (\text{diag}(\widetilde{\mathbf{N}}_j) U) \right) \mathbf{Q} \right], \quad j = 1, 2, 3, \quad (\text{A.77})$$

$$\mathbf{z}^4(\mathbf{c}) = \mathcal{S}_i(\mathbf{c}) \left[\widetilde{\mathbf{N}}^T \widetilde{\mathbf{X}} (I_3 \otimes U) + (I_3 \otimes \widetilde{\mathbf{P}} U) \right] \mathbf{Q}, \quad (\text{A.78})$$

where the geometrical quantities on the fine grid are given by

$$\widetilde{\mathbf{X}}_j = \begin{pmatrix} (\tilde{\mathbf{x}}_1)_j \\ \vdots \\ (\tilde{\mathbf{x}}_{\kappa^2 N})_j \end{pmatrix}, \quad \widetilde{\mathbf{N}}_j = \begin{pmatrix} (\tilde{\mathbf{n}}_1)_j \\ \vdots \\ (\tilde{\mathbf{n}}_{\kappa^2 N})_j \end{pmatrix} \quad (\text{A.79})$$

and

$$\widetilde{\mathbf{N}} = \left(\text{diag}(\widetilde{\mathbf{N}}_1) \quad \text{diag}(\widetilde{\mathbf{N}}_2) \quad \text{diag}(\widetilde{\mathbf{N}}_3) \right), \quad (\text{A.80})$$

$$\widetilde{\mathbf{X}} = \left(\text{diag}(\widetilde{\mathbf{X}}_1) \quad \text{diag}(\widetilde{\mathbf{X}}_2) \quad \text{diag}(\widetilde{\mathbf{X}}_3) \right), \quad (\text{A.81})$$

$$\widetilde{\mathbf{P}} = \text{diag}(\tilde{\mathbf{x}}_1 \cdot \tilde{\mathbf{n}}_1, \dots, \tilde{\mathbf{x}}_{\kappa^2 N} \cdot \tilde{\mathbf{n}}_{\kappa^2 N}). \quad (\text{A.82})$$

From the coefficient vectors $\mathbf{z}^1 - \mathbf{z}^4$ at \mathbf{c} , the double layer potential at a point \mathbf{x} can be computed as

$$\mathcal{D}_i[\Gamma, \mathbf{q}](\mathbf{x}) = \langle x_j \mathbf{g}_i(\mathbf{x}) - \delta_{ij} \mathbf{s}(\mathbf{x}), \mathbf{z}_j(\mathbf{c}) \rangle - \langle \mathbf{g}_i(\mathbf{x}), \mathbf{z}_4(\mathbf{c}) \rangle, \quad (\text{A.83})$$

where the evaluation vectors $\mathbf{s}(\mathbf{x})$ and $\mathbf{g}_i(\mathbf{x})$ are as defined in (A.25) and (A.68).

For the points \mathbf{x}_i and centers \mathbf{c}_i related to the spheroidal grid, the operations (A.77), (A.78) and (A.83) can be represented using a set of transfer matrices $M_i^{1-4} \in \mathbb{C}^{N_p \times 3N}$ and $R_i \in \mathbb{R}^{3 \times 3N}$, such that

$$\mathbf{z}^j(\mathbf{c}_i) = M_i^j Q, \quad j = 1, 2, 3, 4, \quad (\text{A.84})$$

$$\mathcal{D}[\Gamma, \mathbf{q}](\mathbf{x}_i) = R_i Q, \quad (\text{A.85})$$

where $i = 1, \dots, N$. We can, just as for the single layer and dipole potentials, compute the matrices R_i and M_i^j for all points by applying rotation and mirror operations to the matrices from the first $n_\theta/2$ points. For R_i , we can combine (A.63), (A.64), (A.70) and (A.71) directly to get

$$R_i = T_x R_{n_\theta - i + 1} (T_x \otimes F_{n_\theta}), \quad i \in [n_\theta/2 + 1, n_\theta], \quad (\text{A.86})$$

$$R_{i + \alpha n_\theta} = T_z^\alpha (\Delta\varphi) R_i (T_z(-\Delta\varphi) \otimes P_{n_\theta})^\alpha, \quad i \in [1, n_\theta]. \quad (\text{A.87})$$

For the matrices M_i^j , we need to take into account that $M_i^1 - M_i^3$ are directly related to the coordinate directions x, y and z . This means that when applying the symmetry operations to these matrices, we have to take this into account. For the mirror symmetry we then get

$$\begin{aligned} M_i^1 &= \Theta M_{n_\theta - i + 1}^{1,*} (T_x \otimes F_{n_\theta}), \\ M_i^2 &= -\Theta M_{n_\theta - i + 1}^{2,*} (T_x \otimes F_{n_\theta}), \\ M_i^3 &= -\Theta M_{n_\theta - i + 1}^{3,*} (T_x \otimes F_{n_\theta}), \\ M_i^4 &= \Theta M_{n_\theta - i + 1}^{4,*} (T_x \otimes F_{n_\theta}), \end{aligned} \quad i \in [n_\theta/2 + 1, n_\theta], \quad (\text{A.88})$$

while for the rotational symmetry

$$\begin{aligned}
M_{i+\alpha n_\theta}^1 &= E^\alpha (\cos(\alpha\Delta\varphi)M_i^1 - \sin(\alpha\Delta\varphi)M_i^2) (T_z(-\Delta\varphi) \otimes P_{n_\theta})^\alpha, \\
M_{i+\alpha n_\theta}^2 &= E^\alpha (\sin(\alpha\Delta\varphi)M_i^1 + \cos(\alpha\Delta\varphi)M_i^2) (T_z(-\Delta\varphi) \otimes P_{n_\theta})^\alpha, \\
M_{i+\alpha n_\theta}^3 &= E^\alpha M_i^3 (T_z(-\Delta\varphi) \otimes P_{n_\theta})^\alpha, \\
M_{i+\alpha n_\theta}^4 &= E^\alpha M_i^4 (T_z(-\Delta\varphi) \otimes P_{n_\theta})^\alpha.
\end{aligned} \tag{A.89}$$

With this we have all we need for using precomputed QBX on a spheroidal grid, while only storing precomputed data for the first $n_\theta/2$ points.

- [1] L. af Klinteberg. Ewald summation for the rotlet singularity of Stokes flow. arXiv:1603.07467 [physics.flu-dyn], 2016.
- [2] L. af Klinteberg and A.-K. Tornberg. Fast Ewald summation for Stokesian particle suspensions. *Int. J. Numer. Methods Fluids*, 76(10):669–698, 2014, doi:10.1002/fld.3953.
- [3] L. af Klinteberg and A.-K. Tornberg. Estimation of quadrature errors in layer potential evaluation using quadrature by expansion. arXiv:1603.08366 [math.NA], 2016.
- [4] G. B. Arfken and H. J. Weber. *Mathematical Methods for Physicists*. Elsevier Academic Press, London, 6th edition, 2005, ISBN 0120598760.
- [5] J. Bagge. *Numerical simulation of an inertial spheroidal particle in Stokes flow*. Master’s thesis, KTH, 2015, <http://urn.kb.se/resolve?urn=urn:nbn:se:kth:diva-180290>, .
- [6] A. Barnett, B. Wu, and S. Veerapaneni. Spectrally Accurate Quadratures for Evaluation of Layer Potentials Close to the Boundary for the 2D Stokes and Laplace Equations. *SIAM J. Sci. Comput.*, 37(4):B519–B542, 2015, doi:10.1137/140990826.
- [7] A. H. Barnett. Evaluation of Layer Potentials Close to the Boundary for Laplace and Helmholtz Problems on Analytic Planar Domains. *SIAM J. Sci. Comput.*, 36(2):A427–A451, 2014, doi:10.1137/120900253.
- [8] J.-P. Berrut and L. N. Trefethen. Barycentric Lagrange Interpolation. *SIAM Rev.*, 46(3):501–517, 2004, doi:10.1137/S0036144502417715.
- [9] J. Bremer and Z. Gimbutas. A Nyström method for weakly singular integral operators on surfaces. *J. Comput. Phys.*, 231(14):4885–4903, 2012, doi:10.1016/j.jcp.2012.04.003.
- [10] J. Bremer and Z. Gimbutas. On the numerical evaluation of the singular integrals of scattering theory. *J. Comput. Phys.*, 251:327–343, 2013, doi:10.1016/j.jcp.2013.05.048.
- [11] A. T. Chwang. Hydromechanics of low-Reynolds-number flow. Part 3. Motion of a spheroidal particle in quadratic flows. *J. Fluid Mech.*, 72(01):17, 1975, doi:10.1017/S0022112075002911.
- [12] A. T. Chwang and T. Y.-T. Wu. Hydromechanics of low-Reynolds-number flow. Part 2. Singularity method for Stokes flows. *J. Fluid Mech.*, 67(04):787, 1975, doi:10.1017/S0022112075000614.
- [13] P. de Anna, T. Le Borgne, M. Dentz, A. M. Tartakovsky, D. Bolster, and P. Davy. Flow Intermittency, Dispersion, and Correlated Continuous Time Random Walks in Porous Media. *Phys. Rev. Lett.*, 110(18):184502, 2013, doi:10.1103/PhysRevLett.110.184502.

- [14] M. Deserno and C. Holm. How to mesh up Ewald sums. I. A theoretical and numerical comparison of various particle mesh routines. *J. Chem. Phys.*, 109(18):7678, 1998, doi:10.1063/1.477414.
- [15] C. L. Epstein, L. Greengard, and A. Klöckner. On the Convergence of Local Expansions of Layer Potentials. *SIAM J. Numer. Anal.*, 51(5):2660–2679, 2013, doi:10.1137/120902859.
- [16] L. Greengard and V. Rokhlin. A new version of the Fast Multipole Method for the Laplace equation in three dimensions. *Acta Numer.*, 6:229, 1997, doi:10.1017/S0962492900002725.
- [17] É. Guazzelli and J. Hinch. Fluctuations and Instability in Sedimentation. *Annu. Rev. Fluid Mech.*, 43(1):97–116, 2011, doi:10.1146/annurev-fluid-122109-160736.
- [18] J. Helsing and R. Ojala. On the evaluation of layer potentials close to their sources. *J. Comput. Phys.*, 227(5):2899–2921, 2008, doi:10.1016/j.jcp.2007.11.024.
- [19] G. B. Jeffery. The Motion of Ellipsoidal Particles Immersed in a Viscous Fluid. *Proc. R. Soc. A Math. Phys. Eng. Sci.*, 102(715):161–179, 1922, doi:10.1098/rspa.1922.0078.
- [20] A. Klöckner, A. Barnett, L. Greengard, and M. O’Neil. Quadrature by expansion: A new method for the evaluation of layer potentials. *J. Comput. Phys.*, 252:332–349, 2013, doi:10.1016/j.jcp.2013.06.027.
- [21] A. Lefebvre-Lepot, B. Merlet, and T. N. Nguyen. An accurate method to include lubrication forces in numerical simulations of dense Stokesian suspensions. *J. Fluid Mech.*, 769:369–386, 2015, doi:10.1017/jfm.2015.101.
- [22] D. Lindbo, L. af Klinteberg, and D. Saffar Shamshirgar. The Spectral Ewald Unified package, http://github.com/ludvigak/SE_unified, 2016.
- [23] D. Lindbo and A.-K. Tornberg. Spectrally accurate fast summation for periodic Stokes potentials. *J. Comput. Phys.*, 229(23):8994–9010, 2010, doi:10.1016/j.jcp.2010.08.026.
- [24] D. Lindbo and A.-K. Tornberg. Spectral accuracy in fast Ewald-based methods for particle simulations. *J. Comput. Phys.*, 230(24):8744–8761, 2011, doi:10.1016/j.jcp.2011.08.022.
- [25] A. A. Mammoli. The treatment of lubrication forces in boundary integral equations. *Proc. R. Soc. A Math. Phys. Eng. Sci.*, 462(2067):855–881, 2006, doi:10.1098/rspa.2005.1600.
- [26] S. Nasser, N. Phan-Thien, and X.-J. Fan. Lubrication approximation in completed double layer boundary element method. *Comput. Mech.*, 26(4):388–397, 2000, doi:10.1007/s004660000188.

- [27] H. Power and G. Miranda. Second Kind Integral Equation Formulation of Stokes' Flows Past a Particle of Arbitrary Shape. *SIAM J. Appl. Math.*, 47(4):689–698, 1987, doi:10.1137/0147047.
- [28] C. Pozrikidis. *Boundary Integral and Singularity Methods for Linearized Viscous Flow*. Cambridge University Press, Cambridge, 1992, ISBN 9780511624124, doi:10.1017/CBO9780511624124.
- [29] B. Quaife and G. Biros. High-volume fraction simulations of two-dimensional vesicle suspensions. *J. Comput. Phys.*, 274:245–267, 2014, doi:10.1016/j.jcp.2014.06.013.
- [30] M. Rachh, A. Klöckner, and M. O'Neil. Fast algorithms for Quadrature by Expansion I: Globally valid expansions. arXiv:1602.05301v1 [math.NA], 2016.
- [31] Y. Saad and M. H. Schultz. GMRES: A Generalized Minimal Residual Algorithm for Solving Nonsymmetric Linear Systems. *SIAM J. Sci. Stat. Comput.*, 7(3):856–869, 1986, doi:10.1137/0907058.
- [32] A. Sangani and A. Acrivos. Slow flow through a periodic array of spheres. *Int. J. Multiph. Flow*, 8(4):343–360, 1982, doi:10.1016/0301-9322(82)90047-7.
- [33] A. S. Sangani and G. Mo. Inclusion of lubrication forces in dynamic simulations. *Phys. Fluids*, 6(5):1653–1662, 1994, doi:10.1063/1.868228.
- [34] S. Tlupova and J. T. Beale. Nearly Singular Integrals in 3D Stokes Flow. *Commun. Comput. Phys.*, 14(5):1207–1227, 2013, doi:10.4208/cicp.020812.080213a.
- [35] A.-K. Tornberg and L. Greengard. A fast multipole method for the three-dimensional Stokes equations. *J. Comput. Phys.*, 227(3):1613–1619, 2008, doi:10.1016/j.jcp.2007.06.029.
- [36] H. J. Wilson. Stokes flow past three spheres. *J. Comput. Phys.*, 245:302–316, 2013, doi:10.1016/j.jcp.2013.03.020.
- [37] L. Ying, G. Biros, and D. Zorin. A high-order 3D boundary integral equation solver for elliptic PDEs in smooth domains. *J. Comput. Phys.*, 219(1):247–275, 2006, doi:10.1016/j.jcp.2006.03.021.
- [38] H. Zhao, A. H. Isfahani, L. N. Olson, and J. B. Freund. A spectral boundary integral method for flowing blood cells. *J. Comput. Phys.*, 229(10):3726–3744, 2010, doi:10.1016/j.jcp.2010.01.024.
- [39] A. A. Zick and G. M. Homsy. Stokes flow through periodic arrays of spheres. *J. Fluid Mech.*, 115:13–26, 1982, doi:10.1017/S0022112082000627.