An accurate integral equation method for Stokes flow with piecewise smooth boundaries

Lukas Bystricky¹ · Sara Pålsson¹ · Anna-Karin Tornberg¹

Received: 2 December 2019 / Accepted: 27 May 2020 / Published online: 7 August 2020
© The Author(s) 2020

Abstract

Two-dimensional Stokes flow through a periodic channel is considered. The channel walls need only be Lipschitz continuous, in other words they are allowed to have corners. Boundary integral methods are an attractive tool for numerically solving the Stokes equations, as the partial differential equation can be reformulated into an integral equation that must be solved only over the boundary of the domain. When the boundary is at least C^1 smooth, the boundary integral kernel is a compact operator, and traditional Nyström methods can be used to obtain highly accurate solutions. In the case of Lipschitz continuous boundaries, however, obtaining accurate solutions using the standard Nyström method can require high resolution. We adapt a technique known as recursively compressed inverse preconditioning to accurately solve the Stokes equations without requiring any more resolution than is needed to resolve the boundary. Combined with a periodic fast summation method we construct a method that is $\mathcal{O}(N \log N)$ where N is the number of quadrature points on the boundary. We demonstrate the robustness of this method by extending an existing boundary integral method for viscous drops to handle the movement of drops near corners.

Keywords Numerical analysis · Boundary integral equations · Creeping flow · Multiphase flow

Mathematics Subject Classification 65R20 · 76D07 · 76T20

1 Introduction

The Stokes equations are used to model slowly moving, highly viscous, fluid flow. They can be thought of as the zero Reynolds number limit of the Navier–Stokes equations.

Communicated by Gunilla Kreiss.

✉ Lukas Bystricky
lukasby@kth.se

¹ Department of Mathematics, KTH Royal Institute of Technology, Stockholm, Sweden

Of the many applications of the Stokes equations, they are often used to model particle suspensions including solid particles [4,19], drops [25,31,32], or vesicles [33]. Even for non-zero Reynolds numbers, the Stokes equations often describe very well the flow near a solid boundary, and can therefore be used to derive effective slip boundary conditions for problems at higher Reynolds numbers [1,5].

An advantage of using the Stokes equations over the Navier–Stokes equations is that the Stokes equations are linear and elliptic, allowing us to recast them as a boundary integral equation (BIE) [20,25,34]. BIEs have several nice properties. All the information needed to solve a BIE is confined to the boundary of the domain; this leads to an immediate dimension reduction. If the boundary of the domain is at least C^1 smooth, the Stokes equations can be represented as a second-kind Fredholm equation [28]. After discretization the condition number of the resulting linear system is independent of the number of discretization points used, meaning that very highly accurate solutions are obtainable. Traditionally one major drawback of BIEs was the need to solve dense linear systems. However, by using efficient iterative solvers such as GMRES [35], combined with fast matrix vector products [8,23], the cost to solve the $N \times N$ dense linear system can be reduced to $\mathcal{O}(N)$ or $\mathcal{O}(N \log N)$, where N is the number of discretization points on the boundary of the domain.

In this paper, we will be considering wall-bounded, periodic Stokes flow. For particulate flows, such models are useful because they allow for the computation of various time averaged quantities over a relatively small reference cell, without the need to simulate an unfeasibly large domain. BIEs have been successfully applied to such problems [24,32,41].

Lipschitz continuous boundaries admit corners, specifically they allow for countable number of corners with the angle θ of each corner satisfying $0 < \theta < 2\pi$ (thus prohibiting cusps). Solving PDEs on Lipschitz domains to high accuracy everywhere in the domain is in general quite challenging, independent of the numerical method used; see for example the discussion in [7]. When solving boundary integral equations on domains with corners, the standard Nyström method fails to achieve optimal accuracy [3]. While the underlying equation has a unique solution, the layer density defined on the boundary can become weakly singular at corner points, thus reducing the accuracy of regular quadrature rules, such as composite Gauss–Legendre quadrature. One approach to solving this problem is to cluster additional quadrature points near the corners. This of course dramatically increases the number of unknowns, while at the same time the accuracy of such an approach is limited [3,12]. A recent paper [40] demonstrates an approach that automates both the spatial adaptivity and the order of the quadrature (similar to hp adaptivity) to achieve high accuracy for complicated domains containing several corners. Other approaches have been developed, some of which use elegant kernel-dependent custom quadrature rules [34,36]. In this paper, we use a kernel-independent method known as *recursively compressed inverse preconditioning* [10,12,13]. This method has the advantage of being relatively simple to implement on top of existing code, and does not introduce any additional unknowns to the linear system that must be solved. We demonstrate the robustness of this method when applied to periodic Stokes flow by solving problems involving moving viscous drops.

Fig. 1 Sketch of a periodic channel with a Lipschitz boundary. The period of the channel is L in the x_1 direction, and the minimum height of the channel is h . The normal vector on the channel walls points into the bulk fluid

2 Governing equations

The governing equations are the steady incompressible Stokes equations,

$$-\mu \Delta \mathbf{u}(\mathbf{x}) + \nabla p(\mathbf{x}) = \mathbf{0}, \quad \mathbf{x} \in \Omega, \quad (1a)$$

$$\nabla \cdot \mathbf{u}(\mathbf{x}) = 0, \quad \mathbf{x} \in \Omega, \quad (1b)$$

where $\mathbf{u}$ is the velocity of the fluid, p is its pressure and μ is its viscosity.

In this paper we will restrict our attention to periodic channel flow as depicted in Fig. 1. For boundary conditions, we will prescribe Dirichlet conditions on the velocity, and enforce periodicity in the x_1 direction on the velocity. In addition, we impose a pressure drop across the reference cell, so that the pressure itself is not periodic, but the gradient of the pressure is,

$$\mathbf{u}(\mathbf{x}) = \mathbf{g}(\mathbf{x}), \quad \mathbf{x} \in \Gamma, \quad (2a)$$

$$\mathbf{u}(\mathbf{x}) = \mathbf{u}(\mathbf{x} + L), \quad \mathbf{x} \in \Omega, \quad (2b)$$

$$p(\mathbf{x}) - p(\mathbf{x} + L) = p_0 - p_1, \quad \mathbf{x} \in \Omega. \quad (2c)$$

To ensure the incompressibility condition is satisfied, the boundary data $\mathbf{g}$ must satisfy the compatibility condition,

$$\int_{\Gamma} \mathbf{u} \cdot \mathbf{n} \, dS = 0.$$

Fig. 2 Sketch of non-periodic, interior Lipschitz domain. The normal vector $\mathbf{n}$ always points into the fluid

3 Boundary integral formulation

For clarity of exposition, in this section and the next we will present the boundary integral formulation, and the treatment for the corners on a non-periodic domain. The periodicity will be addressed in Sect. 5.

Consider the Stokes equations (1) defined inside a Lipschitz domain, along with Dirichlet boundary conditions on the velocity. For the non-periodic case, we will not be concerned with the pressure. The normal vector along the boundary Γ always points into the fluid. See Fig. 2 for a sketch of such a domain.

As given in [29, Chapter 2], a solution to the forced Stokes equations

$$-\mu \Delta \mathbf{u} + \nabla p = \mathbf{g} \delta(\mathbf{x} - \mathbf{x}_0),$$

is given by the velocity and pressure pair

$$u_j(\mathbf{x}) = \frac{1}{4\pi\mu} G_{j\ell}(\mathbf{x} - \mathbf{x}_0) g_\ell, \quad p(\mathbf{x}) = \frac{1}{4\pi} p_j(\mathbf{x} - \mathbf{x}_0) g_j.$$

The stress tensor $\sigma_{j\ell}$ at a point $\mathbf{x}$ is given by the combination of the velocity and pressure,

$$\begin{aligned} \sigma_{jm} &= \frac{1}{4\pi} \left(-\delta_{jm} p_\ell(\mathbf{x} - \mathbf{x}_0) + \frac{\partial G_{j\ell}}{\partial x_m}(\mathbf{x} - \mathbf{x}_0) + \frac{\partial G_{m\ell}}{\partial x_j}(\mathbf{x} - \mathbf{x}_0) \right) g_\ell \\ &= \frac{1}{4\pi} T_{j\ell m}(\mathbf{x} - \mathbf{x}_0) g_\ell. \end{aligned}$$

Here and in the remainder of this paper we have used the Einstein summation convention, where summation over repeated indices is implied. The letters i and k are reserved for later use and are therefore not used as indices.

In $\mathbb{R}^2$, we have that

$$G_{j\ell}(\mathbf{r}) = -\delta_{j\ell} \log(r) + \frac{r_j r_\ell}{r^2}, \quad p_j(\mathbf{r}) = 2 \frac{r_j}{r^2}, \quad T_{j\ell m}(\mathbf{r}) = -4 \frac{r_j r_\ell r_m}{r^4}.$$

Here we have used r to denote the Euclidean norm of $\mathbf{r} := \mathbf{x} - \mathbf{x}_0$.

The tensors $G_{j\ell}$ and $T_{j\ell m}$ are called the Stokeslet and stresslet respectively. From the Stokeslet and stresslet we can define the single- and double-layer potentials, $\mathbb{S}[\mathbf{q}](\mathbf{x})$

and $\mathbb{D}[\mathbf{q}](\mathbf{x})$, as a convolution of the Stokeslet and stresslet with a density function $\mathbf{q}(\mathbf{x})$ defined over Γ . Explicitly,

$$\begin{aligned}\mathbb{S}[\mathbf{q}](\mathbf{x}) &= \frac{1}{4\pi\mu} \int_{\Gamma} q_j(\mathbf{y}) G_{j\ell}(\mathbf{x} - \mathbf{y}) \, ds(\mathbf{y}), & \mathbf{x} \in \Omega, \\ \mathbb{D}[\mathbf{q}](\mathbf{x}) &= -\frac{1}{4\pi} \int_{\Gamma} q_j(\mathbf{y}) T_{j\ell m}(\mathbf{x} - \mathbf{y}) n_m(\mathbf{y}) \, ds(\mathbf{y}), & \mathbf{x} \in \Omega,\end{aligned}$$

where $\mathbf{n}$ is the unit normal vector pointing into the fluid.

Following [10], we will express the solution to (1) as a combination of the single- and double-layer potentials

$$\mathbf{u}(\mathbf{x}) = 2\mathbb{D}[\mathbf{q}](\mathbf{x}) + \eta\mathbb{S}[\mathbf{q}](\mathbf{x}) := \mathbb{K}[\mathbf{q}](\mathbf{x}), \quad \mathbf{x} \in \Omega, \quad (3)$$

where $\eta > 0$ is an arbitrary constant which governs the relation between the single and double layer potentials. To obtain a well-conditioned system η cannot be too large; we will always take $\eta = 1$.

This equation is valid for $\mathbf{x} \in \Omega$. To obtain a boundary integral equation (BIE) we will take the limit of (3) as $\mathbf{x}$ approaches a point $\mathbf{x}_0 \in \Gamma$. To do this, we will need the limiting values of the single- and double-layer potentials,

$$\begin{aligned}\lim_{\mathbf{x} \rightarrow \mathbf{x}_0 \in \Gamma} \mathbb{S}[\mathbf{q}](\mathbf{x}) &= \mathbb{S}[\mathbf{q}](\mathbf{x}_0), \\ \lim_{\mathbf{x} \rightarrow \mathbf{x}_0 \in \Gamma} \mathbb{D}[\mathbf{q}](\mathbf{x}) &= -\frac{1}{2}\mathbf{q}(\mathbf{x}_0) + \mathbb{D}[\mathbf{q}](\mathbf{x}_0).\end{aligned}$$

Applying these limits to (3) and matching it to the boundary condition $\mathbf{g}$ yields the BIE

$$-\mathbf{q}(\mathbf{x}_0) + \mathbb{K}[\mathbf{q}](\mathbf{x}_0) = \mathbf{g}(\mathbf{x}_0), \quad (4)$$

where $\mathbf{x}_0 \in \Gamma$, and $\mathbb{K}$ is the operator defined in (3) restricted to the boundary.

If Γ is a Lyapunov curve (see [4, Chapter 1] for a precise definition), then $\mathbb{K}$ is a compact operator, with eigenvalues accumulating at zero [?, Chapter 15]. In this case (4) can be analyzed using Fredholm theory. In particular the Fredholm alternative applies, and in [10] this is used to demonstrate the existence and uniqueness of solutions. An alternative to the formulation (3) is the so-called completed double-layer formulation [?], [28]. This formulation has the advantage of not requiring the single-layer potential, which is weakly singular on the boundary. We have chosen (3) because numerical experiments [42] have shown it to be more stable for squeezing drops, and in the periodic case it allows us to naturally impose a global pressure gradient, as will be discussed in Sect. 5.

If, as in our case, Γ is only Lipschitz smooth, then $\mathbb{K}$ is not compact, and Fredholm theory cannot be applied. Furthermore, the double-layer potential involves the normal vector, meaning that the double-layer kernel cannot be evaluated at any corner points on Γ . Nonetheless, it can be shown that (4) has a unique solution, even when Γ is only Lipschitz continuous (see for example [?]). In this case the density function $\mathbf{q}$ is

singular at the corner points, however, (3) can still be evaluated for any $\mathbf{x} \in \Omega$, and (4) can be evaluated anywhere on Γ except at the corner points.

4 Numerical methods

To evaluate (4), we use the Nyström method, as described in [2, Chapter 4]. Let $\gamma(s)$, $s \in [0, 2\pi]$ parameterize Γ . The BIE (4) can be written in the abstract form

$$-q_\ell(t) + \int_0^{2\pi} K_{j\ell}(t, s) q_j(s) |\gamma'(s)| ds = g_\ell(t), \quad t \in [0, 2\pi], \quad (5)$$

where K in this case is the kernel of $\mathbb{K}$, i.e.,

$$K_{j\ell}(s, t) = \frac{1}{2\pi} T_{j\ell m}(\gamma(t) - \gamma(s)) n_m(s) + \frac{\eta}{4\pi\mu} G_{j\ell}(\gamma(t) - \gamma(s)). \quad (6)$$

We will approximate the integral in (5) using a composite Gauss–Legendre quadrature scheme of n_{pan} panels and n_q quadrature points per panel,

$$-q_\ell(t) + \sum_{n=1}^N K_{j\ell}(t, s_n) q_j(s_n) |\gamma'(s_n)| w_n = g_\ell(t), \quad t \in [0, 2\pi], \quad (7)$$

where $N = n_{\text{pan}} n_q$ is the total number of quadrature points, and w_n is the quadrature weight corresponding to the quadrature point s_n . For the remainder of this paper we will assume $n_q = 16$.

We will then enforce (7) at the quadrature points s_m , $m = 1, \dots, N$ to get the linear system

$$-q_\ell(s_m) + \sum_{n=1}^N K_{j\ell}(s_m, s_n) q_j(s_n) |\gamma'(s_n)| w_n = g_\ell(s_m), \quad m = 1, \dots, N.$$

Here the point values of the density function q are unknown. When setting up the composite quadrature rule, care must be taken to ensure that each corner on Γ is at the intersection of two panels. In this way we avoid difficulties arising from trying to evaluate the normal vector on a corner, as the Gauss–Legendre quadrature points cluster near but are never located at panel endpoints.

4.1 Singular quadrature

When $t = s$, the kernel of the double-layer potential $T_{j\ell m}(0) n_m(s)$ has a removable singularity, provided that t is not a corner point. The Stokeslet, however, must be handled using a specialized quadrature technique. We will use the approach given in

[26]. We begin by writing the single-layer potential in complex variables,

$$\begin{aligned} \mathbb{S}[q](z) = & \frac{1}{8\pi\mu} \int_{\Gamma} q(\tau) |d\tau| + \frac{1}{4\pi\mu} \int_{\Gamma} q(\tau) \log(|\tau - z|) |d\tau| \\ & - \frac{1}{8\pi\mu} \int_{\Gamma} \overline{q(\tau)} \frac{\tau - z}{\overline{\tau} - \overline{z}} |d\tau|, \end{aligned} \quad (8)$$

where we have introduced a slight abuse of notation to denote $q(z)$ as density function $\mathbf{q}$ written as a complex number, i. e., $z = x_1 + ix_2$ and $q(z) = q_1(\mathbf{x}) + iq_2(\mathbf{x})$.

The only term in the complex formulation of the Stokeslet that does not have a finite limit as $\tau \rightarrow z$ is the term with the logarithm. We will consider this integral over a single panel, Γ_k . The panel extends from α_1 to α_2 , with $\alpha_1, \alpha_2 \in \mathbb{C}$, and is parameterized by $\alpha \in [\alpha_1, \alpha_2]$. Let α be parameterized as,

$$\alpha(t) = \frac{\alpha_1 + \alpha_2}{2} + \psi t,$$

where $t \in [-1, 1]$ and $\psi = (\alpha_2 - \alpha_1)/2$. The log term in the integral above can be rewritten in the form,

$$\begin{aligned} \log(|\tau(\alpha) - z|) &= \log \left(\left| \frac{\psi(\tau(\alpha) - z)}{\alpha - \alpha_z} \frac{\alpha - \alpha_z}{\psi} \right| \right) \\ &= \log \left(\left| \psi \frac{\tau(\alpha) - z}{\alpha - \alpha_z} \right| \right) + \log \left(\left| \frac{\alpha - \alpha_z}{\psi} \right| \right), \end{aligned}$$

where $\tau(\alpha_z) = z$. Define t_z to be such that $\alpha(t_z) = \alpha_z$. Then we have that

$$\frac{\alpha - \alpha_z}{\psi} = t - t_z,$$

which allows us to rewrite the log integral in (8) as

$$\begin{aligned} \int_{\Gamma_k} q(\tau) \log(|\tau - z|) |d\tau| &= \int_{\Gamma_k} q(\tau(\alpha)) \log(|\tau(\alpha) - z|) |\tau'(\alpha)| d\alpha \\ &= \int_{\Gamma_k} q(\alpha) \log \left(\left| \psi \frac{\tau(\alpha) - z}{\alpha - \alpha_z} \right| \right) d\alpha \\ &\quad + \int_{-1}^1 q(\alpha(t)) \log(|t - t_z|) |\tau'(\alpha(t))| dt. \end{aligned}$$

The first integral can be evaluated using the fact that

$$\lim_{\alpha \rightarrow \alpha_z} \frac{\tau(\alpha) - z}{\alpha - \alpha_z} = \tau'(\alpha_z).$$

To evaluate the second integral, we expand the density $q(t)$ as a polynomial series to obtain,

$$I(t_z) = \int_{-1}^1 q(\alpha(t)) \log(|t - t_z|) |\tau'(\alpha(t))| dt \approx \sum_{j=0}^{n_q-1} c_j \int_{-1}^1 t^j \log(|t - t_z|) |\tau'(\alpha(t))| dt.$$

Denoting $\mathbf{h} = h_j$, $j = 0, \dots, n_q - 1$ as

$$h_j(t_z) = \int_{-1}^1 t^j \log(|t - t_z|) |\tau'(\alpha(t))| dt,$$

allows us to write I in the compact form

$$I(t_z) \approx \sum_{j=0}^{n_q-1} c_j h_j = \mathbf{c}^T \mathbf{h}(t_z).$$

The integrals in $\mathbf{h}$ can all be computed analytically using a recursion relation. The vector $\mathbf{c} = \{c_j\}_{j=0}^{n_q-1}$ can be computed by solving the Vandermonde system

$$\mathbf{q} = V\mathbf{c},$$

where $\mathbf{q} = \{q(t_0), \dots, q(t_{n_q-1})\}$, and V is the Vandermonde matrix. Thus I can be approximated by

$$I(t_z) \approx (V^{-1}\mathbf{q})^T \mathbf{h}(t_z) = \mathbf{q}^T (V^{-1})^T \mathbf{h}(t_z) = \mathbf{q}^T \boldsymbol{\omega}(t_z),$$

where $\boldsymbol{\omega}(t_z)$ is the solution to the linear system $V^T \boldsymbol{\omega}(t_z) = \mathbf{h}(t_z)$. The stability of these computations is discussed in [14, Appendix A], here we simply note that the computations are stable up to at least $n_q = 40$. Note that this linear system is independent of $\mathbf{q}$. For now we will need the values of I only at the quadrature points $t_0, \dots, t_{n_q-1}$. Since we have rescaled and rotated the panel Γ_k to run from -1 to 1 , the corrected weights $\boldsymbol{\omega}(t_z)$ can be precomputed for each of the quadrature points $t_0, \dots, t_{n_q-1}$.

Quadrature points on adjacent panels will also need to be corrected. How many points to correct depends on accuracy considerations, but for $n_q = 16$, numerical results have shown that if all the panels are the same size (in parameter space), then correcting the four closest quadrature points in each adjacent panel is sufficient. Again, these corrections can be precomputed.

4.2 Recursively compressed inverse preconditioning

As previously mentioned, in the case of Lipschitz domains, (4) has a unique solution. However at the corners, the density function $\mathbf{q}$ becomes singular. Therefore Gauss–Legendre quadrature fails to accurately integrate the integrands in (4) near the corners. Local panel refinement around the corners is one way to mitigate this issue; however,

Fig. 3 Local panel refinement around a corner. Each addition subdivision adds two extra panels and $4n_q$ additional unknowns

this approach adds a potentially large number of new unknowns, see Fig. 3. In addition, the condition number of the resulting locally refined linear system grows with the number of quadrature points.

An alternative approach, *recursively compressed inverse preconditioning* (RCIP) [12,13], can achieve the same accuracy as local refinement by performing a simple precomputation. As the desired accuracy requirements increase, the precomputation grows linearly with n_{sub} , however, both the size and the condition number of the final linear system remain fixed.

The main idea behind RCIP is to transform the density function $\mathbf{q}$ in (4) into a transformed density $\tilde{\mathbf{q}}$ that is piecewise smooth everywhere on Γ . Then Gauss–Legendre quadrature can be effectively used. To create this transformation, the operator $\mathbb{K}$ is written as

$$\mathbb{K} = \mathbb{K}^\circ + \mathbb{K}^*, \quad (9)$$

where $\mathbb{K}^\circ$ is compact away from the corners, and $\mathbb{K}^*$ describes the corner interactions (note that $\mathbb{K}^*$ is *not* the adjoint of $\mathbb{K}$). These operators will be defined precisely shortly. We then define the transformed density as

$$\tilde{\mathbf{q}} = (-\mathcal{I} + \mathbb{K}^*)\mathbf{q}, \quad (10)$$

where $\mathcal{I}$ is the identity operator.

Using the split (9), and the transformed density (10), we can convert (4) to the transformed BIE for $\tilde{\mathbf{q}}$,

$$(\mathcal{I} + \mathbb{K}^\circ \mathcal{R})\tilde{\mathbf{q}}(\mathbf{x}_0) = \mathbf{g}(\mathbf{x}_0), \quad \mathbf{x}_0 \in \Gamma, \quad (11)$$

where $\mathcal{R} = (-\mathcal{I} + \mathbb{K}^*)^{-1}$. If we assume $\mathbf{g}$ is piecewise smooth, it can be immediately seen that $\tilde{\mathbf{q}}$ must also be piecewise smooth. Since $\mathbb{K}^\circ$ is a smoothing operator, then $\mathbb{K}^\circ \mathcal{R}\tilde{\mathbf{q}}$ will be smooth everywhere. Then, by contradiction, in order for (11) to hold, $\mathcal{I}\tilde{\mathbf{q}} = \tilde{\mathbf{q}}$ must be piecewise smooth.

It remains to discretize (11). One possibility for discretization is to use two meshes: a *coarse* mesh Γ_{coa} on which $\mathbb{K}$ is discretized, and a *fine* mesh Γ_{fin} , on which $\mathcal{R}$ is discretized. To get Γ_{fin} from Γ_{coa} we first designate the two panels on either side of corner j to be the subset $\Gamma_j^* \subset \Gamma_{\text{coa}}$, and define Γ° to be $\Gamma_{\text{coa}} \setminus \bigcup_{j=1}^{n_c} \Gamma_j^*$, where n_c is the number of corners. The panels on either side of corner j in Γ_j^* are then dyadically

Fig. 4 Panel discretization using two meshes. The boundary Γ is first discretized with a coarse mesh, Γ_{coa} . The two panels on either side of each corner are denoted Γ^* , and the remaining panels are denoted Γ° . The panels in Γ^* closest to the corners are dyadically refined n_{sub} times to obtain Γ_{fin} . Each panel, regardless of its size, contains n_q Gauss–Legendre quadrature points

Fig. 5 Nested meshes in parameter space. Let $\Gamma_{j,\ell}^*$ be parameterized by $s \in [-1, 1]$, with corner j at $s = 0$. Then the panel junctions on $\Gamma_{j,\ell c}^*$ will be at $s = \{-1, 0.5, 0, 0.5, 1\}$, and the panel junctions of $\Gamma_{j,\ell b}^*$ will be at $s = \{-1, -0.5, -0.25, 0, 0.25, 0.5, 1\}$. The panels are nested so that the junctions for $\Gamma_{j,(\ell-1)c}^*$ in the parameterization of $\Gamma_{j,\ell}^*$ will be at $s = \{-0.5, -0.25, 0, 0.25, 0.5\}$

refined n_{sub} times to get $\Gamma_{j,\text{fin}}^*$. The fine mesh is defined as $\Gamma_{\text{fin}} = \Gamma^\circ \cup \left(\bigcup_{j=1}^{n_c} \Gamma_{j,\text{fin}}^* \right)$. An example of this is shown in Fig. 4.

We will define $\mathbb{K}^\circ$ and $\mathbb{K}^*$ to be the operator $\mathbb{K}$ restricted to the domains Γ° and Γ^* respectively. Note that $\mathbb{K}^\circ$ is a compact operator. The operator $\mathbb{K}^\circ$ can be discretized on Γ_{coa} to get the matrix $\mathbf{K}^\circ$. This matrix is equivalent to the discretization of $\mathbb{K}$ on Γ_{coa} , but with the entries corresponding to both source and target being outside Γ° set to zero.

The operator $\mathcal{R}$ could be discretized on Γ_{fin} . This would, however, not be much use, since it would introduce a large number of unknowns. Instead we will exploit a forward recursion relation to construct $\mathbf{R}$ on a sequence of meshes covering larger and larger portions of Γ_j^* . To define the recursion relation, we will need to provide some definitions of different types of meshes. For each corner j , define a sequence of meshes $\Gamma_{j,\ell}^*$, $\ell = 1, \dots, n_{\text{sub}}$, with $\Gamma_{j,\ell-1}^* \subset \Gamma_{j,\ell}^*$ and $\Gamma_{j,n_{\text{sub}}}^* = \Gamma_j^*$. On each $\Gamma_{j,\ell}^*$ we will have a six panel type-b mesh, $\Gamma_{j,\ell b}^*$, and a four panel type-c mesh $\Gamma_{j,\ell c}^*$. The type-b and type-c panels will be related as shown in Fig. 5.

To interpolate between type-b and type-c meshes defined over the same interval, we introduce the prolongation matrix $\mathbf{P}_{bc}$. In addition, defining $\mathbf{W}_b$ and $\mathbf{W}_c$ to be a diagonal matrix containing quadrature weights on the type-b and type-c meshes

respectively, we can define the weighted prolongation matrix $\mathbf{P}_{Wbc}$ as

$$\mathbf{P}_{Wbc} = \mathbf{W}_b \mathbf{P} \mathbf{W}_c^{-1}.$$

Let $\mathbf{K}_{j,\ell b}$ be the discretization of $\mathbb{K}$ on $\Gamma_{j,\ell b}^*$. This matrix will be of size $6n_q$ for scalar problems, or $12n_q$ for vector problems in $\mathbb{R}^2$. Define $\mathbf{R}_{j,1}$ as

$$\mathbf{R}_{j,1} = \mathbf{P}_{Wbc}^T (-\mathbf{I}_b + \mathbf{K}_{j,1b}) \mathbf{P}_{bc}.$$

We then compute the sequence $\mathbf{R}_{j,2}, \dots, \mathbf{R}_{j,n_{\text{sub}}}$ using the recursion relation

$$\mathbf{R}_{j,\ell} = \mathbf{P}_{Wbc}^T \left(\mathbb{F}\{\mathbf{R}_{j,\ell-1}^{-1}\} - \mathbf{I}_b^\circ + \mathbf{K}_{j,\ell b}^\circ \right)^{-1} \mathbf{P}_{bc}, \quad (12)$$

where $\mathbf{I}_b^\circ$ and $\mathbf{K}_{j,\ell b}^\circ$ are the identity matrix and $\mathbf{K}_{j,\ell b}$ with the entries in the two panels around the corner zeroed out respectively. The operator $\mathbb{F}\{\cdot\}$ zero pads the argument to turn a matrix defined on $\Gamma_{j,\ell c}^*$ into one defined on $\Gamma_{j,(\ell+1)b}^*$.

Once we have the matrices $\mathbf{R}_{j,n_{\text{sub}}}$, we can construct $\hat{\mathbf{R}}$, the discretization of $\mathcal{R}$ on Γ_{coa} . Outside of Γ^* , $\mathbb{K}_*$ is zero, so from the definition of $\mathcal{R}$, we obtain that $\hat{\mathbf{R}}$ must be the identity matrix over Γ° . Over Γ_j^* , the discretization of $\mathcal{R}$ is just $\mathbf{R}_{j,n_{\text{sub}}}$.

To handle the log singularity when assembling the matrix $\mathbf{K}_{j,1b}$ and $\mathbf{K}_{j,\ell b}^\circ$, the same techniques as described in Sect. 4.1 can be used. Some bookkeeping is needed to account for the fact that the panels are not equally sized in parameter space. In some cases adjacent panels will be double or half the size; however, the modifications for the quadrature weights can still be precomputed for each case.

For wedge shaped corners, the domain segments $\Gamma_{j,\ell b}$ are self similar for $\ell = 1, \dots, n_{\text{sub}}$ and $j = 1, \dots, n_c$. If the operator $\mathbb{K}$ is scale invariant, then $\mathbf{K}_{j,\ell b}^\circ$ will be independent of ℓ , and the recursion relation (12) becomes a fixed point iteration which can be iterated to find $\mathbf{R}_j$ to a desired tolerance without specifying n_{sub} [12, Section 12]. Unfortunately, in our case, because the single-layer potential contains a logarithmic term, $\mathbb{K}$ is not scale invariant, so we cannot use this idea, and n_{sub} must be specified a priori. A formulation involving just the double-layer potential [?], [28] would not suffer from this drawback and could be formulated as fixed point iteration.

4.3 Fast multiplication

The fully discrete transformed BIE defined on Γ_{coa} is

$$(\mathbf{I} + \mathbf{K}^\circ \hat{\mathbf{R}}) \tilde{\mathbf{q}} = \mathbf{g}, \quad (13)$$

which can be solved using GMRES. To accelerate the required matrix-vector products from $\mathcal{O}(N^2)$ to something computationally feasible, fast summation methods are a

necessity. To do this, we will rewrite $\mathbf{K}$ and $\hat{\mathbf{R}}$ as the block matrices

$$\mathbf{K} = \begin{pmatrix} \mathbf{K}^{**} & \mathbf{K}^{*\circ} \\ \mathbf{K}^{\circ*} & \mathbf{K}^{\circ\circ} \end{pmatrix}, \quad \hat{\mathbf{R}} = \begin{pmatrix} \mathbf{R} & \mathbf{0} \\ \mathbf{0} & \mathbf{I}^\circ \end{pmatrix}.$$

It follows that $\mathbf{K}^\circ$ can be expressed as

$$\mathbf{K}^\circ = \mathbf{K} - \begin{pmatrix} \mathbf{K}^{**} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{pmatrix} = \mathbf{K} - \mathbf{K}^*.$$

Then (13) can be rewritten as

$$(\mathbf{I} + \mathbf{K}\hat{\mathbf{R}} - \mathbf{K}^*\hat{\mathbf{R}}) \tilde{\mathbf{q}} = \mathbf{g}.$$

The matrix vector product $\hat{\mathbf{R}}\tilde{\mathbf{q}}$ can be done directly, since $\hat{\mathbf{R}}$ is just the identity matrix with one $8n_q \times 8n_q$ block for each corner. The remaining matrix vector products $\mathbf{K}(\hat{\mathbf{R}}\tilde{\mathbf{q}})$ and $\mathbf{K}^*(\hat{\mathbf{R}}\tilde{\mathbf{q}})$ can be computed with a fast summation method, reducing the cost from $\mathcal{O}(N^2)$ to $\mathcal{O}(N \log N)$ or $\mathcal{O}(N)$, depending on the choice of fast summation method. We will use an $\mathcal{O}(N \log N)$ method that can naturally handle periodic problems, as will be described in Sect. 5.

4.4 Post-processing

After we have computed $\tilde{\mathbf{q}}$, we evaluate the velocity at a point $\mathbf{u} \in \Omega$ by using the regular quadrature rule

$$u_\ell(\mathbf{x}) \approx \sum_{n=1}^N K_{j\ell}(\mathbf{x}, s_n) \hat{q}_j(s_n) |\gamma'(s_n)| w_n,$$

where $\hat{\mathbf{q}} = \hat{\mathbf{R}}\tilde{\mathbf{q}}$, and $K_{j\ell}$ is defined in (6). This quadrature is as accurate as if we had access to the fine density defined on Γ_{fin} . From the definitions of $\tilde{\mathbf{q}}$ and $\hat{\mathbf{R}}$, it is clear that outside of Γ^* , $\hat{\mathbf{q}}$ and $\mathbf{q}$ are equivalent. On Γ^* , the transformed density $\hat{\mathbf{q}}$ can be thought of as a weight corrected density function.

Both the double- and single-layer potentials contain integrals that become nearly singular when $|t - s|$ is small. This causes the numerical errors from standard Gauss–Legendre quadrature to grow quite large when evaluating the integrals at any point close to a boundary. To handle this, the integrals are treated in the manner described in [32]. In short, the main idea is similar to that explained in Sect. 4.1 where the density is expanded as a polynomial series, and the integrals are computed analytically using recursion relations.

5 Periodicity

We now address the periodicity in (2). The periodic single- and double-layer potentials, $\mathbb{S}^P[\mathbf{q}](\mathbf{x})$ and $\mathbb{D}^P[\mathbf{q}](\mathbf{x})$ are defined as an infinite sum of the single- and double-layer potentials defined in Sect. 3,

$$\mathbb{S}^P[\mathbf{q}](\mathbf{x}) = \frac{1}{4\pi\mu} \sum_{\mathbf{p} \in \mathbb{Z}^2} \int_{\Gamma} q_j(\mathbf{y}) G_{j\ell}(\mathbf{x} - \mathbf{y} - L\mathbf{p}) \, ds(\mathbf{y}), \quad \mathbf{x} \in \Omega, \quad (14a)$$

$$\mathbb{D}^P[\mathbf{q}](\mathbf{x}) = \frac{1}{4\pi} \sum_{\mathbf{p} \in \mathbb{Z}^2} \int_{\Gamma} q_j(\mathbf{y}) T_{j\ell m}(\mathbf{x} - \mathbf{y} - L\mathbf{p}) n_m(\mathbf{y}) \, ds(\mathbf{y}), \quad \mathbf{x} \in \Omega. \quad (14b)$$

The interpretation of these periodic sums will be given in Sect. 5.1. With these operators we can define the periodic operator $\mathbb{K}^P[\mathbf{q}]$,

$$\mathbb{K}^P[\mathbf{q}](\mathbf{x}) = 2\mathbb{D}^P[\mathbf{q}](\mathbf{x}) + \eta\mathbb{S}^P[\mathbf{q}](\mathbf{x}). \quad (15)$$

Note that periodic sum is over $\mathbf{p} \in \mathbb{Z}^2$; in other words we are implying periodicity in both spatial dimensions, even though in our actual problems the periodicity will be only in one direction. The reason for this is that it allows us to exploit a more efficient fast-summation method. To enforce the periodicity in one dimension we will embed the channel, which is only periodic over a length L_1 in the x_1 direction, in a doubly periodic box of size $L_1 \times L_2$. A similar approach is used in [41] where a periodic flow in one direction is embedded in a two-dimensional periodic box. For simplicity we will assume here that $L_1 = L_2 = L$, but this need not be the case, and only minor modifications are needed to use a rectangular periodic box.

As mentioned in Sect. 2, a constant pressure drop $p_0 - p_1$ in the x_1 direction is applied. In order to impose this we will add on an unknown mean velocity $\langle \mathbf{u} \rangle$ to the layer formulation (4),

$$\mathbf{u}(\mathbf{x}) = \mathbb{K}^P[\mathbf{q}](\mathbf{x}) + \langle \mathbf{u} \rangle.$$

Note that $\langle \cdot \rangle$ denotes a volume average over a full periodic cell, including the parts outside Ω . This quantity has no physical meaning beyond imposing a pressure gradient. For a channel with flat parallel walls, a simple alternative would be to add a predetermined background flow, as done in [32].

The mean pressure gradient $\langle \nabla p \rangle$ must balance the net effect of the wall friction [41]. From (15) the wall friction force is equal to $\eta \int_{\Gamma} \mathbf{q} \, d\Gamma$ [10]. The system we have to solve is thus given by

$$\begin{aligned} -\mathbf{q}(\mathbf{x}) + \mathbb{K}^P[\mathbf{q}](\mathbf{x}) + \langle \mathbf{u} \rangle &= \mathbf{g}(\mathbf{x}), \quad \mathbf{x} \in \Gamma, \\ -\frac{1}{L^2} \eta \int_{\Gamma} q_1 \, d\Gamma &= \langle \nabla p \rangle, \\ \int_{\Gamma} q_2 \, d\Gamma &= 0. \end{aligned}$$

The splits for the two dimensional Stokeslet and stresslet, as well as truncation estimates are derived in [31]. Here we list the results.

To compute the infinite sums in (14) we will use the spectral Ewald method [22,23]. In the spectral Ewald method the infinite sums in $\mathbb{S}^P$ and $\mathbb{D}^P$ are split into two parts based on a so-called Ewald decomposition: one that decays exponentially fast in real space, and one that decays exponentially fast in Fourier space. The real space sum can be truncated in real space, while the Fourier sum can be truncated in Fourier space. The splits for the two-dimensional Stokeslet and stresslet, as well as truncation estimates are derived in [32]. Here we list the results.

5.1 Ewald splits

The discretized periodic single- and double layer potentials are given by

$$\begin{aligned}\mathbb{S}^P[\mathbf{q}](\mathbf{x}) &\approx u_\ell^G(\mathbf{x}) = \frac{1}{4\pi\mu} \sum_{\mathbf{p} \in \mathbb{Z}^2}^* \sum_{m=1}^M G_{j\ell}(\mathbf{x} - \mathbf{x}_m - L\mathbf{p}) q_j(\mathbf{x}_m) w_m, \\ \mathbb{D}^P[\mathbf{q}](\mathbf{x}) &\approx u_\ell^T(\mathbf{x}) = \frac{1}{4\pi} \sum_{\mathbf{p} \in \mathbb{Z}^2}^* \sum_{m=1}^M T_{j\ell s}(\mathbf{x} - \mathbf{x}_m - L\mathbf{p}) q_j(\mathbf{x}_m) n_s(\mathbf{x}_m) w_m,\end{aligned}$$

where the $*$ in the summations over $\mathbf{p}$ indicates that we are excluding any $\mathbf{p}$ that makes $\mathbf{x} - \mathbf{x}_m - L\mathbf{p} = \mathbf{0}$. As written, these sums are not well-defined, as they are not convergent. By assuming that a pressure gradient is balancing the forces acting on the fluid, well-defined but slowly converging sums can be found in Fourier space. See [32, Section 4] for a discussion. Below we will introduce an Ewald decomposition for these sums that yields the same well-defined results, but achieves them by a split into two rapidly converging sums: one in real space, and one in Fourier space.

The discretized single-layer potential can be rewritten as the split

$$\begin{aligned}u_\ell^G(\mathbf{x}) &= \sum_{\mathbf{p} \in \mathbb{Z}^2}^* \sum_{m=1}^M G_{j\ell}^R(\mathbf{x} - \mathbf{x}_m + L\mathbf{p}, \xi) q_j(\mathbf{x}_m) w_m \\ &\quad + \frac{1}{L^2} \sum_{\mathbf{k} \neq \mathbf{0}} \hat{G}_{j\ell}^F(\mathbf{k}, \xi) \sum_{m=1}^M e^{-i\mathbf{k} \cdot (\mathbf{x} - \mathbf{x}_m)} q_j(\mathbf{x}_m) w_m,\end{aligned}$$

where

$$\begin{aligned}G_{j\ell}^R(\mathbf{r}, \xi) &= \frac{1}{4\pi} \left(e^{-\xi^2 r^2} \left(\frac{r_j r_\ell}{r^2} - \delta_{j\ell} \right) + \frac{\delta_{j\ell}}{2} E_1(\xi^2 r^2) \right) \\ \hat{G}_{j\ell}^F(\mathbf{k}, \xi) &= \frac{1}{k^2} e^{-k^2/(4\xi^2)} \left(\delta_{j\ell} - \frac{k_j k_\ell}{k^2} \right) \left(1 + \frac{k^2}{4\xi^2} \right), \quad \mathbf{k} \neq \mathbf{0}, \quad k = |\mathbf{k}|.\end{aligned}$$

The $\mathbf{k} = \mathbf{0}$ mode in the Fourier expansion has been shown to be 0 [30]. The decomposition parameter $\xi > 0$ is called the Ewald parameter and determines the relative sizes of the real and Fourier parts of $\mathbf{u}^G(\mathbf{x})$. Note that $\mathbf{u}^G(\mathbf{x})$ itself is independent of ξ .

When applying the Nyström method to solve for the unknown pointwise density values, we do not wish to evaluate singular cases where $\mathbf{x} - \mathbf{x}_n + L\mathbf{p} = \mathbf{0}$. This contribution to $\mathbf{u}^G(\mathbf{x})$ can be skipped in the real space sum, but for the Fourier sum we will have to subtract off the limiting value given by

$$\lim_{r \rightarrow 0} (G_{j\ell}(\mathbf{r}) - G_{j\ell}^R(\mathbf{r}))q_j(\mathbf{x}_n)w_n = \left(\frac{1}{2}\gamma + \log(\xi) + 1\right)q_\ell(\mathbf{x}_n)w_n,$$

where γ is the Euler–Mascheroni constant.

For the discretized double-layer potential, we use the split,

$$\begin{aligned} u_\ell^T(\mathbf{x}) &= \sum_{\mathbf{p} \in \mathbb{Z}^2}^* \sum_{m=1}^M T_{j\ell s}^R(\mathbf{x} - \mathbf{x}_m + L\mathbf{p}, \xi)q_j(\mathbf{x}_m)n_s(\mathbf{x}_m)w_m \\ &\quad + \frac{1}{L^2} \sum_{\mathbf{k} \neq \mathbf{0}} \hat{T}_{j\ell s}^F(\mathbf{k}, \xi) \sum_{m=1}^M e^{-i\mathbf{k} \cdot (\mathbf{x} - \mathbf{x}_m)} q_j(\mathbf{x}_m)n_s(\mathbf{x}_m)w_m \\ &\quad + \frac{1}{L^2} \sum_{m=1}^M \hat{T}_{j\ell s}^{F,0}(\mathbf{x}_m)q_j(\mathbf{x}_m)n_s(\mathbf{x}_m)w_m. \end{aligned}$$

where

$$\begin{aligned} T_{j\ell s}^R(\mathbf{r}, \xi) &= \frac{1}{4\pi} e^{-\xi^2 r^2} \left(2\xi^2 (\delta_{j\ell} r_s + \delta_{js} r_\ell + \delta_{\ell s} r_j) - \frac{4r_j r_\ell r_s}{r^4} (1 + \xi^2 r^2) \right), \\ \hat{T}_{j\ell s}^F(\mathbf{k}, \xi) &= \frac{i}{k^2} e^{-k^2/(4\xi^2)} \left(\delta_{j\ell} k_s + \delta_{js} k_\ell + \delta_{\ell s} k_j - \frac{2k_j k_\ell k_s}{k^2} \right) \left(1 + \frac{k^2}{4\xi^2} \right), \quad \mathbf{k} \neq \mathbf{0}. \end{aligned}$$

For the $\mathbf{k} = \mathbf{0}$ mode, the choice

$$\hat{T}_{j\ell s}^{F,0}(\mathbf{x}) = \delta_{\ell s} x_j, \quad (16)$$

guarantees zero mean-flow though the reference cell [18].

For the stresslet, the limit

$$\lim_{r \rightarrow 0} (T_{j\ell s}(\mathbf{r}) - T_{j\ell s}^R(\mathbf{r}))q_\ell n_s = \mathbf{0},$$

so no limiting value needs to be subtracted when $\mathbf{r} = \mathbf{0}$.

To compute these sums, it is necessary to truncate them. Fortunately, both these sums decay exponentially fast, and they can be truncated to a desired tolerance following the estimates in [32]. With appropriate scaling of the decomposition parameter, the

Table 1 Maximum relative errors in a pipe flow simulation along lines parallel to the top wall. As the target points get closer to the wall, the integrands become more challenging to evaluate accurately, but special quadrature is activated giving high accuracy everywhere using a small amount of panels

n_{pan} per wall	Distance from wall			
	0.5	0.1	0.01	0.001
1	1.64×10^{-1}	2.07×10^0	2.10×10^1	8.44×10^1
2	1.14×10^{-4}	8.15×10^{-1}	4.65×10^0	1.28×10^1
4	5.95×10^{-10}	1.12×10^{-9}	5.03×10^{-9}	7.07×10^{-9}
8	8.88×10^{-15}	1.25×10^{-13}	2.02×10^{-13}	1.08×10^{-11}
16	7.77×10^{-15}	1.06×10^{-13}	1.60×10^{-13}	7.90×10^{-12}

real space part is computed in $\mathcal{O}(N)$ time and the Fourier space part is accelerated to $\mathcal{O}(N \log N)$ using the fast Fourier transform. In order to use FFTs, the source points are spread to a uniform grid where the computations for the Fourier space sum are carried out. The result is then gathered from the uniform grid to the target points. The spreading and gathering is done using truncated Gaussians whose shape parameter is selected to minimize the approximation error for a given support. For efficiency, fast Gaussian gridding [9] can be used in both the spreading and the gathering steps. For more details see [32].

A numerical difficulty in the periodic formulation lies in the fact that the Ewald representation of the stresslet is not translation invariant due to the zero mode (16). This means that the submatrices $\mathbf{R}_j$ are different for each corner, even if the corners have the exact same shape. Therefore to assemble $\mathbf{R}_j$ we first assemble the translation invariant part, and then add on (16) only at the end. Additionally, to avoid roundoff error, for large n_{sub} , (where the smallest $\Gamma_{j,\ell}^*$ is $\mathcal{O}(10^{-16})$ or less) we introduce a local coordinate system for each corner centered at $\mathbf{x} = \mathbf{0}$.

6 Numerical examples

To test the periodization scheme, we can compare to a known exact solution. Pressure driven pipe flow creates the well-known Poiseuille parabolic flow profile in the x_2 direction. The exact solution for flow through a flat channel with top wall at $x_2 = 1$ and bottom wall at $x_2 = 0$ is given by

$$\mathbf{u}^{\text{exact}}(\mathbf{x}) = \left\langle -\frac{(p_1 - p_0)x_2(x_2 - 1)}{2L\mu}, 0 \right\rangle$$

Note that this flow is constant in the x_1 direction, and therefore also periodic in the x_1 direction. As can be seen in Table 1, using only 8 panels per wall allows us to achieve a relative error of 10^{-11} everywhere up to a distance of 10^{-3} from the boundary.

Another test is to prescribe Dirichlet boundary conditions on the solid walls. We will perform two types of tests: convergence towards a known, smooth solution, and self-convergence test towards an unknown solution. For the known smooth solution,

Fig. 6 A shear flow $\mathbf{u} = \langle x_2, 0 \rangle$ is prescribed as a Dirichlet boundary condition on the top and bottom walls. If we take the pressure drop δp to be 0, the BIE should recover the shear flow everywhere in the domain. Contour plots of the base 10 log of the relative error are shown for different values of n_{sub} . Top: $n_{\text{sub}} = 0, 10$, bottom: $n_{\text{sub}} = 20, 30$. In every case each wall has been discretized using 20 panels with $n_q = 16$, meaning $N = 1280$ unknowns. Without any special treatment, the recovered solution differs quite a bit from the exact solution everywhere. With $n_{\text{sub}} = 20$ the BIE can achieve 11 digits accuracy everywhere except very close to the corners. If additional accuracy is required near the corners, the fine density function can be recovered and used in conjunction with the special quadrature described earlier

we will prescribe a constant shear flow $\mathbf{u} = \langle x_2, 0 \rangle$ as boundary conditions on the top and bottom wall. The bottom wall is now only piecewise smooth. To recover a shear flow in the interior from the BIE solution, we must prescribe no pressure drop from inlet to outlet, i.e., $\langle \nabla p \rangle = 0$. We will use this simple example to test the RCIP algorithm for various values of n_{sub} . Figure 6 shows the results. If we use the standard Nyström discretization without RCIP, we get quite large errors everywhere in the domain; the errors are not localized around the corner. For $n_{\text{sub}} = 30$ the RCIP algorithm allows us to achieve 11 digits accuracy everywhere in the domain, except for close to the corners. To gain additional accuracy near the corners, it is possible to recover the actual fine density and use it along with the special quadrature described in Sect. 4.4 [12, Section 10].

For the self-convergence study, we will use the same domain, but prescribe zero boundary conditions on the walls, and set $\langle \nabla p \rangle = 1$. We compare the computed velocity with $n_{\text{sub}} = 0, 10, 20$, and 30, to one with $n_{\text{sub}} = 50$. As can be seen in Fig. 7, the velocity field using $n_{\text{sub}} = 30$ matches the solution with $n_{\text{sub}} = 50$ to 11 digits.

The number of refinements needed to achieve a desired accuracy depends heavily on the geometry. Reentrant corners in particular require more refinement. For example Fig. 8 shows an example where $n_{\text{sub}} = 20$ achieves only a maximum of 6 digits of accuracy. Using $n_{\text{sub}} = 60$ we are able to achieve 11 digits accuracy in the interior.

Fig. 7 The same geometry as in Fig. 6 is used to perform a self-convergence study. We prescribe an average pressure gradient $\langle p \rangle = 1$ and set $\mathbf{u} = \mathbf{0}$ on the walls. We then compare the velocity field for different n_{sub} to one with $n_{\text{sub}} = 50$. Top: $n_{\text{sub}} = 0, 10$, bottom $n_{\text{sub}} = 20, 30$. Each wall has been discretized using 20 panels with $n_q = 16$, meaning $N = 1280$ unknowns for each simulation. As with Fig. 6, the contour plot shows the base 10 log of the relative error

At the finest n_{sub} we are placing quadrature points on panels of size 10^{-19} . Using a local coordinate system centered at zero is essential in this case.

6.1 Remarks on timings

Each subdivision in the RCIP algorithm involves inverting an $\mathbf{R}$ matrix in (12). This matrix will be of the size $8n_q \times 8n_q$; all our simulations use $n_q = 16$, meaning $\mathbf{R}$ will be 128×128 . This is not a very cheap procedure, and for small problems this cost can indeed dominate the cost of solving the linear system to find $\tilde{\mathbf{q}}$. Using Matlab 2017a on a desktop with 32 GB of RAM and an Intel i7 processor each subdivision step takes approximately 1.6 s (this is independent of the problem under consideration). For the problem in Figs. 6 and 7 with $n_{\text{sub}} = 10$, assembling the RCIP matrices (3 of them) takes around 48 s. By comparison solving the linear system of size 1280×1280 using GMRES takes around 0.7 s, and postprocessing the solution at 25,000 points (and applying special quadrature where needed) takes around 4 s.

The cost of a single matrix-vector product is slightly higher as well using RCIP. Without RCIP, for each GMRES iteration it is necessary to compute the product $\mathbf{K}\mathbf{q}$, usually using some kind of fast summation method. As discussed in Sect. 4.3, when using RCIP, each matrix-vector product requires the additional computation of $\hat{\mathbf{R}}\tilde{\mathbf{q}}$, as well as the product $\mathbf{K}^*(\hat{\mathbf{R}}\tilde{\mathbf{q}})$. As mentioned, the first of those products can be done directly, since $\hat{\mathbf{R}}$ is a block diagonal matrix, with one $8n_q \times 8n_q$ block per corner. The second of these products is actually just another $8n_q \times 8n_q$ product per corner. This

Fig. 8 A shear flow $\mathbf{u} = \langle x_2, 0 \rangle$ is prescribed as a Dirichlet boundary condition on the top and bottom walls. If we take the pressure drop δp to be 0, the BIE should recover the shear flow everywhere in the domain. Note that the bottom wall contains several reentrant corners. Contour plots of the base 10 log of the relative error are shown for different values of n_{sub} . Top: $n_{\text{sub}} = 0, 20$, bottom: $n_{\text{sub}} = 40, 60$. Both the top and bottom walls are discretized with 65 panels with $n_q = 16$, meaning 4160 unknowns. For comparison, a locally refined solution with $n_{\text{sub}} = 60$ for this problem would require over 50,000 unknowns

product is effectively zeroing out the entries in the matrix $\mathbf{K}$ corresponding to entries where both the source and target points are close to corners. Therefore it should be done using the same fast method as the full matrix-vector product $\mathbf{K}(\hat{\mathbf{R}}\hat{\mathbf{q}})$. In the example corresponding to Figs. 6 and 7, one full matrix-vector product takes 0.035 s without RCIP and 0.038 s with RCIP, an increase of about 8%. This percentage increase will depend on the n_q , the number of corners, as well as the relative size of Γ^* to Γ , but *not* on n_{sub} .

It should be noted, that without RCIP the cost grows as $\mathcal{O}(N \log N)$ with local refinement. Here N is the number of unknowns, and grows by $4n_q$ per corner for each subdivision. In the examples shown in Figs. 6 and 7 this would correspond to solving a linear system of size $1280 + 4(16)(10)(3) = 3200$; for $n_{\text{sub}} = 50$, the linear system would be of the size 10,880. In addition, the condition number of the refined matrix would grow with the number of quadrature points, and require more GMRES iterations to solve to a desired accuracy. For time dependent problems as in Sect. 7, the larger system would have to be solved at every time step. In this case RCIP has a clear advantage, since as long as the wall geometry remains fixed, the $\mathbf{R}$ matrices need only be assembled once and a smaller linear system can be solved at each time step.

7 Adding viscous drops

Earlier work [32] has looked at modelling the movement of drops inside confined periodic geometries. We now demonstrate the robustness and usefulness of RCIP by extending the method in [32] to model the movement of drops near sharp interfaces. A drop is a packet of fluid that does not mix with the fluid surrounding it. Surface

Fig. 9 Sketch of a problem involving deformable drops. The region Ω_0 is the bulk fluid and Γ_0 are solid walls. The regions Ω_1 , Ω_2 , and Ω_3 are drops, with interfaces Γ_1 , Γ_2 and Γ_3 respectively. On each boundary the normal vector $\mathbf{n}$ points into the bulk fluid. The problem is periodic in the horizontal direction

tension forces prevent the mixing of the drop with the bulk fluid. Both the fluid inside each drop, and the bulk fluid, satisfy the incompressible Stokes equations,

$$\begin{aligned} -\mu_0 \Delta \mathbf{u}(\mathbf{x}) + \nabla p(\mathbf{x}) &= \mathbf{0}, & \mathbf{x} \in \Omega_0, \\ -\mu_\ell \Delta \mathbf{u}(\mathbf{x}) + \nabla p(\mathbf{x}) &= \mathbf{0}, & \mathbf{x} \in \Omega_\ell, \ell = 1, \dots, N_d, \\ \nabla \cdot \mathbf{u} &= 0, & \mathbf{x} \in \bigcup_{\ell=0}^{N_d} \Omega_\ell, \end{aligned}$$

where μ_ℓ denotes the viscosity inside the region bounded by Γ_ℓ . See Fig. 9 for a representative problem. As λ increases the drop behaves more like a rigid particle.

At the interfaces the velocity is continuous; however, in general the surface forces acting on the interface from the inside and the outside of the drop are not equal. We will denote the jump in the normal force across interface ℓ as $\delta \mathbf{f}_\ell(\mathbf{x})$. This jump is related to the curvature $\kappa_\ell(\mathbf{x})$, and the surface tension $\sigma_\ell(\mathbf{x})$ by Pozrikidis [29, Chapter 5],

$$\delta \mathbf{f}_\ell(\mathbf{x}) =: (-p_0 - p_\ell) \mathbf{n}_\ell + 2(\mu_0 \mathbf{e}_0 - \mu_\ell \mathbf{e}_\ell) \cdot \mathbf{n}_\ell = \sigma_\ell(\mathbf{x}) \kappa_\ell(\mathbf{x}) \mathbf{n}(\mathbf{x}) - \nabla_s \sigma_\ell(\mathbf{x}),$$

where $\mathbf{e}_\ell$ denotes the rate of strain tensor in Ω_ℓ and ∇_s is the gradient along the interface.

To nondimensionalize this problem, we will define h to be the minimum height of the channel. We would like the maximum velocity of an empty channel to be 1. To do this, we will introduce a pressure scale $h \langle \nabla p \rangle / 8$, a length scale h , a velocity scale $h^2 \langle \nabla p \rangle / (8\mu_0)$, and a surface tension scale σ_0 to obtain the full nondimensional form of the problem,

$$\begin{aligned}
-\Delta \mathbf{u}(\mathbf{x}) + \nabla p(\mathbf{x}) &= \mathbf{0}, & \mathbf{x} \in \Omega_0, \\
-\lambda_\ell \Delta \mathbf{u}(\mathbf{x}) + \nabla p(\mathbf{x}) &= \mathbf{0}, & \mathbf{x} \in \Omega_\ell, \ell = 1, \dots, N_d, \\
\nabla \cdot \mathbf{u} &= 0, & \mathbf{x} \in \bigcup_{\ell=0}^{N_d} \Omega_\ell, \\
\delta \mathbf{f}_\ell &= \frac{1}{L^p} (\sigma_\ell \kappa_\ell \mathbf{n}_\ell - \nabla_s \sigma_\ell), & \mathbf{x} \in \Gamma_\ell, \\
\mathbf{u}(\mathbf{x}) &= \mathbf{u}(\mathbf{x} + L), \\
\langle \nabla p \rangle &= (8, 0),
\end{aligned}$$

where the Laplace number L_p is the dimensionless quantity given by $h^2 \langle \nabla p \rangle / (8\sigma_0)$ [20]. The time scale is now $h/U = 8\mu_0 / (h \langle \nabla p \rangle)$. The Laplace number serves the same purpose as the Capillary number in other drop simulations [32], but for pressure driven flows it is more natural to nondimensionalize according to the pressure gradient as opposed to a maximum velocity. In [31, 32] chemical surface active agents change the surface tension of the drops; however, for the remainder of this paper we will assume a constant surface tension, i.e. $\nabla_s \sigma = 0$. Allowing for non constant surface tension would not impact the RCIP algorithm in any way.

A BIE formulation [42] for the problem is given by

$$\mathbf{u}(\mathbf{x}) = \sum_{\ell=1}^{N_d} \left(\mathbb{S}_{\Gamma_\ell}^P [\delta \mathbf{f}](\mathbf{x}) + (\lambda_\ell - 1) \mathbb{D}_{\Gamma_\ell}^P [\mathbf{u}](\mathbf{x}) \right) + \mathbb{K}^P [\mathbf{q}](\mathbf{x}) + \langle \mathbf{u} \rangle, \quad (17)$$

where $\mathbb{S}_{\Gamma_\ell}^P$ and $\mathbb{D}_{\Gamma_\ell}^P$ are the periodic single- and double-layer operators defined on Γ_ℓ .

Taking the limit as $\mathbf{x} \rightarrow \mathbf{x}_0 \in \Gamma_0$,

$$\begin{aligned}
-\mathbf{q}(\mathbf{x}_0) + \mathbb{K}^P [\mathbf{q}](\mathbf{x}_0) + \sum_{\ell=1}^{N_d} (\lambda_\ell - 1) \mathbb{D}_{\Gamma_\ell}^P [\mathbf{u}](\mathbf{x}_0) + \langle \mathbf{u} \rangle &= \mathbf{g}(\mathbf{x}_0) \\
-\sum_{\ell=1}^{N_d} \mathbb{S}_{\Gamma_\ell}^P [\delta \mathbf{f}](\mathbf{x}_0), & \quad \mathbf{x}_0 \in \Gamma_0,
\end{aligned}$$

and taking the limit $\mathbf{x} \rightarrow \mathbf{x}_0 \in \Gamma_m$,

$$\begin{aligned}
\mathbf{u}(\mathbf{x}_0) - 2 \sum_{\ell=1}^{N_d} \frac{\lambda_\ell - 1}{\lambda_m + 1} \mathbb{D}_{\Gamma_\ell}^P [\mathbf{u}](\mathbf{x}_0) - \frac{2}{\lambda_m + 1} \left(\mathbb{K}^P [\mathbf{q}](\mathbf{x}_0) + \langle \mathbf{u} \rangle \right) \\
= \sum_{\ell=1}^{N_d} \frac{2}{\lambda_m + 1} \mathbb{S}_{\Gamma_\ell}^P [\delta \mathbf{f}](\mathbf{x}_0) \quad \mathbf{x}_0 \in \Gamma_m.
\end{aligned}$$

As in Sect. 5 we close the system by relating the density function $\mathbf{q}$, still defined only on the solid walls Γ_0 , to the (nondimensionalized) average pressure gradient,

$$\eta \int_{\Gamma_0} \mathbf{q} \, d\Gamma_0 = \frac{1}{L^2} \begin{pmatrix} 8 \\ 0 \end{pmatrix}.$$

Fig. 10 Snapshots of a drop moving inside a channel containing a corner. A pressure drop is applied from left to right. By varying the number of points on the drop and the wall we can perform a numerical convergence study, the results of which are shown in Fig. 11. We have run two simulations, one with $\lambda = 1$ (blue drop), and another with $\lambda = 5$ (red drop). As expected, the drop with $\lambda = 5$ deforms less than the $\lambda = 1$ drop. $Lp = 1$ for all simulations. Note that for presentation purposes the red and blue drops are overlayed on top of each other and appear to be purple where they overlap (colour figure online)

By inspecting (17), we see that if $\lambda_\ell = 1$, $\ell = 1, \dots, N_d$, then the velocity on the drop interfaces only appears on one side of the equation. This means that after solving for $\mathbf{q}$ on the solid walls, $\mathbf{u}$ can be computed as a post-processing step. In general; however, both $\mathbf{u}$ and $\mathbf{q}$ need to be solved for simultaneously. The drop boundaries are discretized in exactly the same way as the solid walls, and the same close evaluation scheme is used to compute both $\mathbb{D}_{\Gamma_\ell}^P[\mathbf{u}]$ and $\mathbb{S}_{\Gamma_\ell}^P[\delta\mathbf{f}]$ for drops that are near each other, or near a solid wall. After computing $\mathbf{u}(\mathbf{x})$ on the drop interfaces, the drops move and deform according to the velocity,

$$\frac{d\mathbf{x}}{dt} = \mathbf{u}(\mathbf{x}), \quad \mathbf{x} \in \bigcup_{\ell=1}^{N_d} \Gamma_\ell.$$

To evaluate this system of ODEs we will use the fourth order adaptive time stepper described in [16]. Further details on time stepping, including a way to maintain consistent arclength spacing as the drop perimeter changes, can be found in [31,32].

It is important to note that since the solid walls are not moving, the RCIP matrices $\mathbf{R}_j$, $j = 1, \dots, n_c$ need only be computed at the start of the simulation. This means that after this precomputation the actual time needed each time step to solve the linear system is independent of n_{sub} .

7.1 Numerical results involving viscous drops

To illustrate the performance of our method, we will consider two different examples. First, Fig. 10 shows snapshots of a simulation of a single drop passing through a narrow piecewise continuous constriction. We begin by creating a high resolution reference solution with RCIP for $\lambda = 1$ and $\lambda = 5$. The reference solution will be discretized using $n_{\text{pan}} = 140$ on both the solid walls and the drop, and $n_{\text{sub}} = 20$. The time stepping tolerance for the high resolution simulation is 10^{-10} .

To compute an error, we will look at the ℓ_∞ difference of a numerical solution with the computed reference solution. As described in [32], when updating the positions of the drops it is advantageous to work on a uniform grid with parameter spacing Δs ,

Fig. 11 Spatial error study with and without RCIP for the simulation shown in Fig. 10. Using RCIP gives much lower errors, and these are near the time stepping tolerance of 10^{-8} . $Lp = 1$ for all simulations

instead of the panel Gauss–Legendre points. Using a uniform grid also allows us to easily upsample the non-reference solutions using an FFT to obtain a solution at the same discretization points as the reference solution.

As can be seen in Fig. 11, without RCIP, as we refine the number of points on both the drop and the wall, the numerical solutions for both $\lambda = 1$ and $\lambda = 5$ converge very slowly towards the high resolution reference solution. When we use RCIP with $n_{\text{sub}} = 20$, the numerical solutions are much more accurate; around the level of the time stepping tolerance for the low resolution simulations of 10^{-8} .

As a second example, we will investigate the movement of multiple drops of different sizes confined in a periodic channel containing several corners. A snapshot of such a simulation at $t = 6$ (after all the drops have passed at least one periodic length) is shown in Fig. 12. As the snapshot makes clear, there is a clear, visible difference in the modelled drops depending on whether or not we use RCIP. In particular, the difference between no RCIP and $n_{\text{sub}} = 10$ is quite large, whereas the difference between the $n_{\text{sub}} = 20$ and $n_{\text{sub}} = 30$ is hardly noticeable.

Further proof of the necessity of the proper handling of corners is shown in Fig. 13. Since the fluid inside the drops is incompressible, the area of each drop should be conserved. Note that this is not enforced explicitly, nor is area conservation used in the criteria for the temporal adaptivity. Without RCIP, the area after each drop has passed one periodic channel length is conserved only up to around 10^{-4} . When applying RCIP with $n_{\text{sub}} = 10$, it is below 10^{-7} , and when applying $n_{\text{sub}} = 20$, the area error is at or below the time stepping tolerance.

Fig. 12 Snapshot of three drops of different sizes in a periodic channel. Four different simulations are run starting from identical initial conditions, each with a different n_{sub} . For the red drops, no RCIP is used; $n_{\text{sub}} = 10$ for blue drops; $n_{\text{sub}} = 20$ for orange drops; $n_{\text{sub}} = 30$ for cyan drops. The area error as a function of time is shown in Fig. 13. There is a visible difference between the $n_{\text{sub}} = 0$ and all the other simulations, thus demonstrating the necessity of the special corner treatment. As n_{sub} increases the magnified regions demonstrate that the drops converge to the same solution (the cyan drops). For all cases $\lambda = 1$ and $Lp = 0.2$ for all the drops. The top wall is discretized with 20 panels, the bottom with 30 panels, and each drop with 60 panels. The time stepping tolerance is set to 10^{-10} (colour figure online)

8 Conclusion

Domains with sharp corners pose challenges for boundary integral methods. The layer density defined on the boundary becomes weakly singular at the corners, and therefore standard quadrature rules cannot be accurately applied. This destroys the accuracy of the post-processed solution everywhere in the domain. We have demonstrated how a technique known as recursively compressed inverse preconditioning (RCIP) can be used to accurately solve the Stokes equations on two-dimensional domains with corners. This method requires a small amount of precomputation, however, it does not add any additional unknowns to the problem, nor does it increase the condition number of the linear system. Numerical experiments have demonstrated its robustness and usefulness for postprocessing the velocity anywhere in the domain. Additionally we have shown that it can be added to an existing drop model [32] to accurately model the movement of drops near corners.

Fig. 13 Area conservation errors for the simulation shown in Fig. 12 as a function of time. The error eventually plateaus around 10^{-4} without RCIP, but remains below 10^{-10} when RCIP with $n_{\text{sub}} = 20$ is used. The time stepping tolerance is set to 10^{-10}

Future work could include extending this model to three dimensions. A boundary integral method for three-dimensional drops in free space has been developed [37,38]. In three dimensions Lipschitz domains admit both corners and edges, so the RCIP algorithm described in Sect. 4.2 must be further developed. In [15] such an approach is used to compute the capacitance of a cube.

Acknowledgements Open access funding provided by Royal Institute of Technology. The authors would like to acknowledge support the Knut and Alice Wallenberg Foundation (Award Number KAW 2017.0410) as well as by the Göran Gustafsson Foundation for Research in Nature and Medicine

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

1. Achdou, Y., Pironneau, O., Valentin, F.: Effective boundary conditions for laminar flows over periodic rough boundaries. *J. Comput. Phys.* **147**(1), 187–218 (1998)
2. Atkinson, K.E.: *The Numerical Solution of Integral Equations of the Second Kind*. Cambridge University Press, Cambridge (1997)
3. Bremer, J.: On the Nyström discretization of integral equations on planar curves with corners. *Appl. Comput. Harmon. Anal.* **32**(1), 45–64 (2012)

4. Bystricky, L., Shanbhag, S., Quaife, B.: Stable and contact-free time stepping for dense rigid particle suspensions. *Int. J. Numer. Methods in Fluids* **92**(2), 94–113 (2019)
5. Dalibard, A.-L., Gérard-Varet, D.: Effective boundary condition at a rough surface starting from a slip condition. *J. Differ. Equ.* **251**(12), 3450–3487 (2011)
6. Gohberg, I., Krupnik, N.: One Dimensional Linear Singular Integral Equations. I. Introduction. Birkhäuser Verlag, Basel (1992)
7. Gopal, A., Trefethen, L.N.: Solving Laplace problems with corner singularities via rational functions. *SIAM J. Numer. Anal.* **57**(5), 2074–2094 (2019)
8. Greengard, L., Rokhlin, V.: A fast algorithm for particle simulations. *J. Comput. Phys.* **73**(2), 325–348 (1987)
9. Leslie Greengard and June Yub Lee: Accelerating the nonuniform fast Fourier transform. *SIAM Rev.* **46**(3), 443–454 (2004)
10. Hebekker, F.-K.: Efficient boundary element methods for three-dimensional exterior viscous flow. *Numer. Methods for PDEs* **2**, 273–297 (1986)
11. Helsing, J., Jiang, S.: On integral equation methods for the first Dirichlet problem of the biharmonic and modified biharmonic equations in nonsmooth domains. *SIAM J. Sci. Comput.* **40**(4), A2609–A2630 (2018)
12. Helsing, J.: Solving integral equations on piecewise smooth boundaries using the RCIP method: a tutorial. *Abstr. Appl. Anal.* **2013** (2013). Article ID 938167
13. Helsing, J., Ojala, R.: Corner singularities for elliptic problems: integral equations, graded meshes, quadrature, and compressed inverse preconditioning. *J. Comput. Phys.* **227**(20), 8820–8840 (2008)
14. Helsing, J., Ojala, R.: On the evaluation of layer potentials close to their sources. *J. Comput. Phys.* **227**(5), 2899–2921 (2008)
15. Helsing, J., Perfekt, K.-M.: On the polarizability and capacitance of the cube. *Appl. Comput. Harmon. Anal.* **34**(3), 445–468 (2013)
16. Kennedy, C.A., Carpenter, M.H.: Additive Runge–Kutta schemes for convection–diffusion–reaction equations. *Appl. Numer. Math.* **44**(1), 139–181 (2003)
17. Kim, S., Karrila, S.J.: *Microhydrodynamics: Principles and Selected Applications*. Dover Publications Inc, Mineola, NY (2005)
18. af Klinteberg, L., Shamshirgar, D.S., Tornberg, A.-K.: Fast Ewald summation for free-space Stokes potentials. *Res. Math. Sci.* **4**(1) (2017)
19. af Klinteberg, L., Tornberg, A.-K.: A fast integral equation method for solid particles in viscous flow using quadrature by expansion. *J. Comput. Phys.* **326**, 420–445 (2016)
20. Kuneš, J.: *Dimensionless Physical Constants in Science and Engineering*. Elsevier, London (2012)
21. Ladyzhenskaya, O.A.: *The Mathematical Theory of Viscous Incompressible Flow*, Second English edn. Gordon and Breach Science Publishers, London (1987)
22. Lindbo, D., Tornberg, A.-K.: Spectrally accurate fast summation for periodic Stokes potentials. *J. Comput. Phys.* **229**(23), 8994–9010 (2010)
23. Lindbo, D., Tornberg, A.-K.: Spectral accuracy in fast Ewald-based methods for particle simulations. *J. Comput. Phys.* **230**(24), 8744–8761 (2011)
24. Marple, G.R., Barnett, A., Gillman, A., Veerapaneni, S.: A fast algorithm for simulating multiphase flows through periodic geometries of arbitrary shape. *SIAM J. Sci. Comput.* **38**(5), 740–772 (2015)
25. Ojala, R., Tornberg, A.K.: An accurate integral equation method for simulating multi-phase Stokes flow. *J. Comput. Phys.* **298**, 145–160 (2015)
26. Ojala, R.: A robust and accurate solver of Laplace’s equation with general boundary conditions on general domains in the plane. *J. Comput. Math.* **30**(4), 433–448 (2012)
27. Power, H.: The completed double layer boundary integral formulation for two-dimensional Stokes flow. *IMA J. Appl. Math.* **51**(2), 123–145 (1993)
28. Power, H., Miranda, G.: Second kind integral equation formulation of Stokes’ flows past a particle of arbitrary shape. *SIAM J. Appl. Math.* **47**(4), 689–698 (1987)
29. Pozrikidis, C.: *Boundary Integral and Singularity Methods for Linearized Viscous Flow*. Cambridge University Press, Cambridge (1992)
30. Pozrikidis, C.: Computation of periodic Green’s functions of Stokes flow. *J. Eng. Math.* **30**(1), 79–96 (1996)
31. Pålsson, S., Siegel, M., Tornberg, A.-K.: Simulation and validation of surfactant-laden drops in two-dimensional Stokes flow. *J. Comput. Phys.* **386**, 218–247 (2019)

32. Pålsson, S., Tornberg, A.-K.: An integral equation method for closely interacting surfactant-covered droplets in wall-confined Stokes flow. *Int. J. Numer. Methods Fluids* (2020) **(to appear)**
33. Quaife, B., Biros, G.: High-volume fraction simulations of two-dimensional vesicle suspensions. *J. Comput. Phys.* **274**, 245–267 (2014)
34. Rachh, M., Serkh, K.: On the solution of Stokes equation on regions with corners. *arXiv e-prints* (2017)
35. Saad, Y., Schultz, M.H.: GMRES: a generalized minimal residual algorithm for solving nonsymmetric linear systems. *SIAM J. Sci. Stat. Comput.* **7**(3), 856–869 (1986)
36. Serkh, K., Rokhlin, V.: On the solution of elliptic partial differential equations on regions with corners. *J. Comput. Phys.* **305**, 150–171 (2016)
37. Sargentone, C., Tornberg, A.-K.: A highly accurate boundary integral equation method for surfactant-laden drops in 3D. *J. Comput. Phys.* **360**, 167–191 (2018)
38. Sargentone, C., Tornberg, A.-K., Vlahovska, P.M.: A 3D boundary integral method for the electrohydrodynamics of surfactant-covered drops. *J. Comput. Phys.* **389**, 111–127 (2019)
39. Verchota, G.: Layer potentials and regularity for the Dirichlet problem for Laplace's equation in Lipschitz domains. *J. Funct. Anal.* **59**, 572–611 (1984)
40. Wu, B., Zhu, H., Barnett, A., Veerapaneni, S.: Solution of Stokes flow in complex nonsmooth 2D geometries via a linear-scaling high-order adaptive integral equation scheme. *J. Comput. Phys.* **410** (2020) **(to appear)**
41. Zhao, H., Isfahani, A.H.G., Olson, L.N., Freund, J.B.: A spectral boundary integral method for flowing blood cells. *J. Comput. Phys.* **229**(10), 3726–3744 (2010)
42. Zinchenko, A.Z., Davis, R.H.: A boundary-integral study of a drop squeezing through interparticle constrictions. *J. Fluid Mech.* **564**, 227–266 (2006)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.