

A UNIFIED INTEGRAL EQUATION SCHEME FOR DOUBLY-PERIODIC LAPLACE AND STOKES BOUNDARY VALUE PROBLEMS IN TWO DIMENSIONS

ALEX H. BARNETT*, GARY R. MARPLE†, SHRAVAN VEERAPANENI‡, AND LIN ZHAO‡

Abstract. We present a spectrally-accurate scheme to turn a boundary integral formulation for an elliptic PDE on a single unit cell geometry into one for the fully periodic problem. The basic idea is to use a small least squares solve to enforce periodic boundary conditions without ever handling periodic Green’s functions. We describe fast solvers for the two-dimensional (2D) doubly-periodic conduction problem and Stokes non-slip fluid flow problem, where the unit cell contains many inclusions with smooth boundaries. Applications include computing the effective bulk properties of composite media (homogenization), and microfluidic chip design.

We split the infinite sum over the lattice of images into a directly-summed “near” part plus a small number of auxiliary sources which represent the (smooth) remaining “far” contribution. Applying physical boundary conditions on the unit cell walls gives an expanded linear system, which, after a rank-1 or rank-3 correction and a Schur complement, leaves a *well-conditioned* square system which can be solved iteratively using fast multipole acceleration plus a low-rank term. We are rather explicit about the consistency and nullspaces of both the continuous and discretized problems. The scheme is simple (no lattice sums, Ewald methods, nor particle meshes are required), allows adaptivity, and is essentially dimension- and PDE-independent, so generalizes without fuss to 3D and to other elliptic problems. In order to handle close-to-touching geometries accurately we incorporate recently developed spectral quadratures. We include eight numerical examples and a software implementation. We validate against high-accuracy results for the square array of discs in Laplace and Stokes cases (improving upon the latter), and show linear scaling for up to 10^4 randomly located inclusions per unit cell.

1. Introduction. Periodic boundary value problems (BVPs) arise frequently in engineering and the sciences, either when modeling the behavior of solid or fluid media with true periodic geometry, or when applying artificial periodic boundary conditions to simulate a representative domain of a random medium or particulate flow (often called a super-cell or representative volume element simulation). The macroscopic response of a given microscopic periodic composite medium can often be summarized by an effective material property (e.g. a conductivity or permeability tensor), a fact placed on a rigorous footing by the field of *homogenization* (for a review see [14]). However, with the exception of layered media that vary in only one dimension, finding this tensor requires the numerical solution of “cell problems” [62, 66], namely BVPs in which the solution is periodic up to some additive constant expressing the macroscopic driving. Application areas span all of the major elliptic PDEs, including the Laplace equation (thermal/electrical conductivity, electrostatics and magnetostatics of composites [35, 38, 27, 12]); the Stokes equations (porous flow in periodic solids [18, 50, 75, 26], sedimentation [1], mobility [69], transport by cilia carpets [16], vesicle dynamics in microfluidic flows [58]); elastostatics (microstructured periodic or random composites [29, 36, 61, 64]); and the Helmholtz and Maxwell equations (phononic and photonic crystals, bandgap materials [42, 65]). In this work we focus on the first two (non-oscillatory) PDEs above, noting that the methods that we present also apply with minor changes to the oscillatory Helmholtz and Maxwell cases, at least up to moderate frequencies [5, 13, 56].

The accurate solution of periodic BVPs has remained a challenging problem at the forefront of analytical and computational progress for well over a century [9]. Modern simulations may demand large numbers of objects per unit cell, with arbitrary geometries, that at high volume fractions may approach arbitrarily close to each other [72]. In such regimes asymptotic methods based upon expansion in the inclusion size do not apply [18, 26]. Polydisperse suspensions or nonsmooth geometries may require spatially adaptive discretizations. In microfluidics, high aspect ratio and/or skew unit cells are needed, for instance in the optimal design of particle sorters. Furthermore, to obtain accurate average material properties for random media or suspensions via the Monte Carlo method, thousands of simulation runs may be needed [38]. A variety of numerical methods are used (see [62, Sec. 2.8]), including particular solutions (starting in 1892 with Rayleigh’s method for cylinders and spheres [70], and, more recently, in Stokes flow [71]), eigenfunction expansions [75], lattice Boltzmann and finite differencing [48], and finite element methods [29, 61]. However, for general geometries, it becomes very hard to achieve high-order accuracy with any of the above apart from finite elements, and the cost of meshing renders the latter unattractive when moving geometries are involved. Integral equation methods [49, 67, 45, 55] are natural,

*Department of Mathematics, Dartmouth College, and Center for Computational Biology, Flatiron Institute

†Department of Mathematics, University of Michigan

‡INTECH, Princeton, NJ




FIG. 1.1. 2D periodic problem in the case of a single inclusion Ω with boundary $\partial\Omega$. (a) Periodic BVP in \mathbb{R}^2 , showing a possible unit cell “box” \mathcal{B} and its four walls L , R , D , U , and senses of wall normals. (b) Directly-summed “near” copies of $\partial\Omega$. (c) Circle of auxiliary sources (red dots) which represent the field in \mathcal{B} due to the infinite punctured lattice of “far” copies (dashed curves). In (b)–(c) blue arrows indicate the action (source to target) of the four matrix blocks A , B , C , and Q .

since the PDE has piecewise-constant coefficients, and are very popular [50, 36, 27, 26, 64, 12, 1, 69]. By using potential theory to represent the solution field in terms of the convolution of a free-space Green’s function with an unknown “density” function τ living only on material boundaries, one vastly reduces the number of discretized unknowns. The linear system resulting by applying boundary conditions takes the form

$$A\tau = f, \quad (1.1)$$

where the $N \times N$ matrix A is the discretization of an integral operator. The latter is often of Fredholm second-kind, hence A remains well-conditioned, independent of N , so that iterative methods converge rapidly. Further advantages include their high-order or spectral accuracy (e.g. via Nyström quadratures), the ease of adaptivity on the boundary, and the existence of fast algorithms to apply A with optimal $\mathcal{O}(N)$ cost, such as the fast multipole method (FMM) [28].

The traditional integral equation approach to periodic problems replaces the free-space Green’s function by a *periodic Green’s function*, so that the density is solved on only the geometry lying in a single unit cell [36, 27, 26, 64, 12]. There is an extensive literature on the evaluation of such periodic Green’s functions (e.g. in the Stokes case see [68, 74]); yet any such pointwise evaluation for each source-target pair leads to $\mathcal{O}(N^2)$ complexity, which is unacceptable for large problems. There are two popular approaches to addressing this in a way compatible with fast algorithms:

- Lattice sums. Noticing that the difference between the periodic and free-space Green’s function is a smooth PDE solution in the unit cell, one expands this in a particular solution basis (a cylindrical or spherical expansion); the resulting coefficients, which need to be computed only once for a given unit cell, are called *lattice sums* [7, 35, 27, 26]. They originate in the work of Rayleigh [70] and in the study of ionic crystals (both reviewed in [9, Ch. 2–3]). The FMM may then be periodized by combining the top-level multipole expansion coefficients with the lattice sums to give a correction to the local expansion coefficients (in 2D this is a discrete convolution) which may be applied fast [28, Sec. 4.1] [64]. This method has been used for the doubly-periodic Laplace BVP by Greengard–Moura [27], and Stokes BVP by Greengard–Kropinski [26].
- Particle-mesh Ewald (PME) methods. These methods exploit Ewald’s realization [20] that, although both the spatial and the spectral (Fourier) sums for the periodic Green’s function in general converge slowly, there is an analytic splitting into spatial and spectral parts such that both terms converge superalgebraically. Numerically, the spatial part is now local, while the spectral part may be applied by “smearing” onto a uniform grid, using a pair of FFTs, and (as in the non-uniform FFT [19]) correcting for the smearing. The result is a $\mathcal{O}(N \log N)$ fast algorithm; for a review see [15]. Recently this has been improved by Lindbo–Törnberg to achieve




FIG. 1.2. *Solution of the Laplace equation in a doubly-periodic domain with Neumann boundary conditions on each of $K = 10^4$ inclusions, driven by specified potential drop $\mathbf{p} = (1, 0)$ (Example 2). A single unit cell is shown, with the solution u indicated by a color scale and contours. The inset shows detail. Each inclusion has $N_k = 700$ discretization nodes on its boundary, resulting in 7 million degrees of freedom. The solution took about 10 hours on a 12-core 3.4 GHz Intel Xeon desktop. The flux is $J_1 = 0.1963795463$ with estimated absolute error of 1×10^{-10} .*




FIG. 1.3. Periodic Stokes flow using in doubly-periodic domain with no-slip boundary conditions on each of $K = 10^3$ inclusions (Example 7). Flow is driven by a specified pressure difference between left and right sides. A single unit cell is shown. Color here indicates fluid speed (red is high and blue is low). Each inclusion boundary has $N_k = 350$ discretization points, resulting in 700000 degrees of freedom (not accounting for copies). The estimated absolute error in the computed flux is 9×10^{-9} .

overall spectral accuracy in the Laplace [53] and Stokes [51] settings, with applications to 3D fluid suspensions [1].

The scheme that we present is based purely on free-space Green’s functions, and has distinct advantages over the above two periodization approaches. 1) Only physically-meaningful boundary conditions and compatibility conditions are used. Aside from conceptual and algebraic simplification, this also removes the need for (often mysterious and subtle [9]) choices of the values of various conditionally or even non-convergent lattice sums based on physical arguments [18, 28, 36, 27] (e.g., six such choices are needed in [26]). When it comes to skewed unit cells combined with applied pressure driving, such hand-picked choices in lattice sum methods become far from obvious. 2) Free-space FMM codes may be inserted without modification, and fast *direct* solvers (based on hierarchical compression of matrix inverses) [31, 23] require only slight modification [22, 58]. 3) When multi-scale spatial features are present in the problem—commonly observed in practical applications such as polydisperse suspensions and complex microfluidic chip geometries—PME codes become inefficient because of their need for uniform grids. In contrast, our algorithm retains the fast nature of the adaptive FMM in such cases. 4) In contrast to lattice sums and Ewald methods, which intrinsically rely on radially-symmetric expansions and kernels, our scheme can handle reasonably high aspect ratio unit cells with little penalty. 5) Since our Stokes formulation does not rely on complex variables (e.g. Sherman–Lauricella [26]), the whole scheme generalizes to 3D with no conceptual changes [57]. A disadvantage of our scheme is that the prefactor may be slightly worse than Ewald methods, due to the direct summation of neighboring images.

The basic idea is simple, and applies to integral representations for a variety of elliptic PDEs. Let \mathcal{B} be a (generally trapezoidal) unit cell “box”, and let $\partial\Omega$ be the material boundary which is to be periodized ($\partial\Omega$ may even intersect the walls of \mathcal{B}). Applying periodic boundary conditions is equivalent to summing the free-space representation on $\partial\Omega$ over the infinite lattice, as in Fig. 1.1(a). Our scheme sums this potential representation over only the 3×3 nearest neighbor images of $\partial\Omega$ (see Fig. 1.1(b)), but adds a small auxiliary basis of smooth PDE solutions in \mathcal{B} , with coefficient vector ξ , that efficiently represents the (distant) contribution of the rest of the infinite image lattice. For this auxiliary basis we use point sources (Fig. 1.1(c)); around 10^2 are needed, a small number which, crucially, is *independent of* N and the boundary complexity. We apply the usual (homogeneous) boundary condition on $\partial\Omega$, which forms the first block row of a linear system. We impose the desired physical periodicity as auxiliary conditions on the Cauchy data differences between wall pairs R – L and U – D (see Fig. 1.1(a)), giving a second block row. The result is a 2×2 block “extended linear system” (ELS),

$$\begin{bmatrix} A & B \\ C & Q \end{bmatrix} \begin{bmatrix} \boldsymbol{\tau} \\ \boldsymbol{\xi} \end{bmatrix} = \begin{bmatrix} \mathbf{0} \\ \mathbf{g} \end{bmatrix}. \quad (1.2)$$

Here \mathbf{g} accounts for the applied macroscopic thermal or pressure gradient in the form of prescribed jumps across one unit cell. Fig. 1.1(b)–(c) sketches the interactions described by the four operators A , B , C , and Q . This linear system is generally *rectangular and highly ill-conditioned*, but when solved in a backward-stable least-squares sense can result in accuracies close to machine precision.

Three main routes to a solution of (1.2) are clear. (a) The simplest is to use a dense direct least-squares solve (e.g., via QR), but the $O(N^3)$ cost becomes impractical for large problems with $N > 10^4$. Instead, to create fast algorithms one exploits the fact that the numbers of auxiliary rows and columns in (1.2) are both small, as follows. (b) One may attempt to eliminate ξ to get the Schur complement square linear system involving A_{per} , the periodized version of A , of the form

$$A_{\text{per}}\boldsymbol{\tau} := (A - BQ^+C)\boldsymbol{\tau} = -BQ^+\mathbf{g}, \quad (1.3)$$

where Q^+ is a pseudoinverse of Q (see Section 2.4). As we will see, (1.3) can be well-conditioned when (1.1) is, and can be solved iteratively, by using the FMM to apply $A\boldsymbol{\tau}$ while applying the second term in its *low-rank* factored form $-B((Q^+C)\boldsymbol{\tau})$. (c) One may instead eliminate $\boldsymbol{\tau}$ by forming, then applying, a compressed representation of A^{-1} via a fast direct solver; this has proven useful for the case of fixed periodic geometry with multiple boundary data vectors [22, 58]. Both methods (b) and (c) can achieve $O(N)$ complexity.

This paper explores route (b). A key contribution is overcoming the stumbling block that, for many standard PDEs and integral representations, A_{per} as written in (1.3) does not in fact exist. Intuitively,

this is due to divergence in the sum of the Green’s function over the lattice. For instance, the sum of $\log(1/r)$ clearly diverges, and thus the single-layer representation for Laplace’s equation, which we use in Section 2, cannot be periodized unless the integral of density (total charge) vanishes. This manifests itself in the 2nd block row of (1.2): the range of C contains vectors that are not in the range of Q , thus Q^+C numerically blows up. After studying in Sections 2.1 and 4.1 the consistency conditions for the linear system involving Q (the “empty unit cell” problem), we will propose rank-1 (for Laplace) and rank-3 (for no-slip Stokes) corrections to the ELS that allow the Schur complement to be taken, and, moreover that remove the nullspace associated with the physical BVP. We justify these schemes rigorously, and test them numerically.

Although the idea of imposing periodicity via extra linear conditions has aided certain electromagnetics solvers for some decades [32, 78], we believe that this idea was first combined with integral equations by the first author and Greengard in [5], where ξ controlled a Helmholtz local expansion. Since then it has become popular for a variety of PDE [5, 22, 30, 13, 12, 58, 56]. From [5] we also inherit the split into near and far images, and cancellations in C that allow a rapidly convergent scheme even when $\partial\Omega$ intersects unit cell walls. The use of point sources as a particular solution basis that is efficient for smooth solutions is known as the “method of fundamental solutions” (MFS) [8, 21], “method of auxiliary sources” [46], “charge simulation method” [44], or, in fast solvers, “equivalent source” [11] or “proxy” [59] representations. This is also used in the recent 3D periodization schemes of Gumerov–Duraishwami [30] and Yan–Shelley [76]. Finally, the low-rank perturbations which enlarge the range of Q are inspired by the low-rank perturbation methods for singular square systems of Sifuentes et al. [73].

Here is a guide to the rest of this paper. In Section 2 we present the periodic Neumann Laplace BVP, possibly the simplest application of the scheme. This first requires understanding and numerically solving the “empty unit cell” subproblem, which we do in Section 2.1. The integral formulation, discretization, and direct numerical solution of the ELS (1.2) follows in Section 2.2. A general scheme to stably take the Schur complement is given in Section 2.4, along with a scheme to remove the physical nullspace specific to our representation. The latter is tested numerically. Section 2.5 shows how the FMM and close-evaluation quadratures are incorporated, and applies them to large problems with thousands of inclusions. Section 2.6 defines the effective conductivity tensor κ and shows how to evaluate it efficiently using a pair of BVP solutions. In Section 3 we show that this same periodizing scheme solves the benchmark periodic disc array conduction problem to high accuracy. In Section 4 we move to periodic Dirichlet (no-slip) Stokes flow, and show how its periodizing scheme closely parallels the Laplace version. In fact, the only differences are new consistency conditions for the empty subproblem (Section 4.1), the use of a combined-field formulation (Section 4.2), and the need for a rank-three perturbation for a stable Schur complement (Section 4.3). Experiments on the drag of a square array of discs, and on large-scale flow problems, are performed in Section 4.4. We discuss generalizations and conclude in Section 5.

Our aim is to illustrate, with two example BVPs, a unified route to a well-conditioned periodization compatible with fast algorithms. We take care to state the *consistency conditions* and *nullspaces* of the full and empty problems, this being the main area of problem-specific variation. We believe that the approach adapts simply to other boundary conditions and PDEs, once the corresponding consistency conditions and nullspaces are laid out. Although general aspect ratios and skew unit cells are amongst our motivations, for clarity we stick to the unit square; the generalization is straightforward.

REMARK 1.1 (Software). *We maintain self-contained MATLAB codes for the methods and figures in this paper at <http://github.com/ahbarnett/DPLS-demos>*

2. The Neumann Laplace case. We now present the heat/electrical conduction problem in the exterior of a periodic lattice of insulating inclusions (corresponding to $\sigma_d = 0$ in [27]). For simplicity we first assume a single inclusion Ω per unit cell. Let \mathbf{e}_1 and \mathbf{e}_2 be vectors defining the lattice in \mathbb{R}^2 ; we work with the unit square so that $\mathbf{e}_1 = (0, 1)$ and $\mathbf{e}_2 = (1, 0)$. Let $\Omega_\Lambda := \{\mathbf{x} \in \mathbb{R}^2 : \mathbf{x} + m_1\mathbf{e}_1 + m_2\mathbf{e}_2 \in \Omega \text{ for some } m_1, m_2 \in \mathbb{Z}\}$ represent the infinite lattice of inclusions. The scalar u , representing electric

potential or temperature, solves the BVP

$$\Delta u = 0 \quad \text{in } \mathbb{R}^2 \setminus \overline{\Omega_\Lambda} \quad (2.1)$$

$$u_n = 0 \quad \text{on } \partial\Omega_\Lambda \quad (2.2)$$

$$u(\mathbf{x} + \mathbf{e}_1) - u(\mathbf{x}) = p_1 \quad \text{for all } \mathbf{x} \in \mathbb{R}^2 \setminus \overline{\Omega_\Lambda} \quad (2.3)$$

$$u(\mathbf{x} + \mathbf{e}_2) - u(\mathbf{x}) = p_2 \quad \text{for all } \mathbf{x} \in \mathbb{R}^2 \setminus \overline{\Omega_\Lambda} , \quad (2.4)$$

ie, u is harmonic, has zero flux on inclusion boundaries, and is periodic up to a given pair of constants $\mathbf{p} = (p_1, p_2)$ which encode the external (macroscopic) driving. We use the abbreviation $u_n = \frac{\partial u}{\partial n} = n \cdot \nabla u$, where n is the unit normal on the relevant evaluation curve.

PROPOSITION 2.1. *For each (p_1, p_2) the solution to (2.1)–(2.4) is unique up to an additive constant.*

Proof. As usual, one considers the homogeneous BVP arising when u is the difference of two solutions. Let \mathcal{B} be any unit cell (tiling domain) containing Ω , then using Green's first identity,

$$0 = \int_{\mathcal{B} \setminus \overline{\Omega}} u \Delta u = - \int_{\mathcal{B} \setminus \overline{\Omega}} |\nabla u|^2 + \int_{\partial \mathcal{B}} u u_n - \int_{\partial \Omega} u u_n .$$

The first boundary term cancels by periodicity, the second by (2.2), hence $\nabla u \equiv 0$. \square

To solve the periodic BVP we first re-express it as a BVP on a single unit cell \mathcal{B} with coupled boundary values on the four walls comprising its boundary $\partial \mathcal{B} := L \cup R \cup D \cup U$; see Fig. 1.1(a). For simplicity we assume for now that the square \mathcal{B} can be chosen such that Ω does not intersect any of the walls; this restriction will later be lifted (see Remark 2.7). We use the notation u_L to mean the restriction of u to the wall L , and u_{nL} for its normal derivative using the normal on L (note that this points to the right, as shown in the figure.) Consider the following reformulated BVP:

$$\Delta u = 0 \quad \text{in } \mathcal{B} \setminus \overline{\Omega} \quad (2.5)$$

$$u_n = 0 \quad \text{on } \partial \Omega \quad (2.6)$$

$$u_R - u_L = p_1 \quad (2.7)$$

$$u_{nR} - u_{nL} = 0 \quad (2.8)$$

$$u_U - u_D = p_2 \quad (2.9)$$

$$u_{nU} - u_{nD} = 0 . \quad (2.10)$$

Clearly any solution to (2.1)–(2.4) also satisfies (2.5)–(2.10). Because of the unique continuation of Cauchy data (u, u_n) as a solution to the 2nd-order PDE, the converse holds, thus the two BVPs are equivalent. We define the *discrepancy* [5] of a solution u as the stack of the four functions on the left-hand side of (2.7)–(2.10).

2.1. The empty unit cell discrepancy BVP and its numerical solution. We first analyze, then solve numerically, an important subproblem which we call the “empty unit cell BVP.” We seek a harmonic function v matching a given discrepancy $g = [g_1; g_2; g_3; g_4]$ (i.e., a stack of four functions defined on the walls L, L, D, D , respectively). That is,

$$\Delta v = 0 \quad \text{in } \mathcal{B} \quad (2.11)$$

$$v_R - v_L = g_1 \quad (2.12)$$

$$v_{nR} - v_{nL} = g_2 \quad (2.13)$$

$$v_U - v_D = g_3 \quad (2.14)$$

$$v_{nU} - v_{nD} = g_4 . \quad (2.15)$$

We now give the consistency condition and nullspace for this BVP, which shows that it behaves essentially like a square linear system with nullity 1.

PROPOSITION 2.2. *A solution v to (2.11)–(2.15) exists if and only if $\int_L g_2 ds + \int_D g_4 ds = 0$, and is then unique up to a constant.*

Proof. The zero-flux condition $\int_{\partial\mathcal{B}} v_n = 0$ holds for harmonic functions. Writing $\partial\mathcal{B}$ as the union of the four walls, with their normal senses, gives the sum of the integrals of (2.13) and (2.15). Uniqueness follows from the method of proof of Prop. 2.1. \square

We now describe a numerical solution method for this BVP that is accurate for a certain class of data g , namely those for which the solution v may be continued as a regular solution to Laplace's equation into a large neighborhood of \mathcal{B} , and hence each of g_1, \dots, g_4 is analytic. It will turn out that this class is sufficient for our periodic scheme (essentially because v will only have to represent distant image contributions, as shown in Fig. 1.1(c)). Let \mathcal{B} be centered at the origin. Recalling the fundamental solution to Laplace's equation,

$$G(\mathbf{x}, \mathbf{y}) = \frac{1}{2\pi} \log \frac{1}{r}, \quad r := \|\mathbf{x} - \mathbf{y}\|, \quad (2.16)$$

we approximate the solution in \mathcal{B} by a linear combination of such solutions with source points \mathbf{y}_j lying uniformly on a circle of radius $R_p > r_{\mathcal{B}}$, where $r_{\mathcal{B}} := 1/\sqrt{2}$ is the maximum radius of the unit cell. That is, for $\mathbf{x} \in \mathcal{B}$,

$$v(\mathbf{x}) \approx \sum_{j=1}^M \xi_j \phi_j(\mathbf{x}), \quad \phi_j(\mathbf{x}) := G(\mathbf{x}, \mathbf{y}_j), \quad \mathbf{y}_j := (R_p \cos 2\pi j/M, R_p \sin 2\pi j/M), \quad (2.17)$$

with unknown coefficient vector $\xi := \{\xi_j\}_{j=1}^M$. As discussed in the introduction, this is known as the MFS. Each basis function ϕ_j is a particular solution to (2.11) in \mathcal{B} , hence only the boundary conditions need enforcing (as in Rayleigh's original method [70]).

This MFS representation is complete for harmonic functions. More precisely, for v in a suitable class, it is capable of exponential accuracy uniformly in \mathcal{B} . In particular, since \mathcal{B} is contained within the ball $\|\mathbf{x}\| \leq r_{\mathcal{B}}$, we may apply known convergence results to get the following.

THEOREM 2.3. *Let v extend as a regular harmonic function throughout the closed ball $\|\mathbf{x}\| \leq \rho$ of radius $\rho > r_{\mathcal{B}}$. Let the fixed proxy radius $R_p \neq 1$ satisfy $\sqrt{r_{\mathcal{B}}\rho} < R_p < \rho$. For each $M \geq 1$ let $\{\phi_j^{(M)}\}_{j=1}^M$ be a proxy basis set as in (2.17). Then there is a sequence of coefficient vectors $\xi^{(M)}$, one vector for each M , and a constant C dependent only on v , such that*

$$\left\| v - \sum_{j=1}^M \xi_j^{(M)} \phi_j^{(M)} \right\|_{L^\infty(\mathcal{B})} \leq C \left(\frac{\rho}{r_{\mathcal{B}}} \right)^{-M/2}, \quad M = 1, 2, \dots$$

In addition, the vectors may be chosen so that the sequence $\|\xi^{(M)}\|_2$, $M = 1, 2, \dots$, is bounded.

This exponential convergence—with rate controlled by the distance to the nearest singularity in v —was derived by Katsurada [43, Rmk. 2.2] in the context of collocation (point matching) for a Dirichlet BVP on the boundary of the disc of radius $r_{\mathcal{B}}$, which is enough to guarantee the existence of such a coefficient vector sequence. The boundedness of the coefficient norms has been known in the context of Helmholtz scattering in the Russian literature for some time (see [47, Sec. 3.4] and references within, and [17, Thm. 2.4]). We do not know of a reference stating this for the Laplace case, but note that the proof is identical to that in [4, Thm. 6], by using equation (13) from that paper. The restriction $R_p < \rho$ is crucial for high accuracy, since when $R_p > \rho$, although convergence occurs, $\|\xi^{(M)}\|_2$ blows up, causing accuracy loss due to catastrophic cancellation [4, Thm. 7].

REMARK 2.4. *Intuitively, the restriction $R_p \neq 1$ arises because the proxy basis may be viewed as a periodic trapezoid rule quadrature approximation to a single-layer potential on the circle $\|\mathbf{x}\| = R_p$, a representation which is incomplete when $R_p = 1$: it cannot represent the constant function [77, Rmk. 1].*

To enforce boundary conditions, let $\mathbf{x}_{iL} \in L$, $\mathbf{x}_{iD} \in D$, $i = 1, \dots, m$, be two sets of m collocation points, on the left and bottom wall respectively; we use Gauss–Legendre nodes. Enforcing (2.12) between collocation points on the left and right walls then gives

$$\sum_{j=1}^M [\phi_j(\mathbf{x}_{iL} + \mathbf{e}_1) - \phi_j(\mathbf{x}_{iL})] \xi_j = g_1(\mathbf{x}_{iL}), \quad \text{for all } i = 1, \dots, m. \quad (2.18)$$




FIG. 2.1. Convergence of the proxy point numerical scheme for the empty discrepancy BVP, for the case of known solution with singularity $\mathbf{x}_0 = 1.5(\cos .7, \sin .7)$ is a distance 1.5 from the center of the square unit cell of side 1. The proxy radius is $R_p = 1.4$. (a)-(b) are for the Laplace case with g deriving from the known solution $v(\mathbf{x}) = \log \|\mathbf{x} - \mathbf{x}_0\|$; (c) is for the Stokes case with known \mathbf{v} coming from a stokeslet at \mathbf{x}_0 with force $\mathbf{f}_0 = [0.3; -0.6]$. (a) Shows convergence in maximum absolute error in v at 100 target points interior to \mathcal{B} , vs the number m of wall quadrature nodes, with fixed $M = 100$ proxy points. (b) Shows convergence in M with fixed $m = 22$ (+ signs) and its prediction via Thm. 2.3 (dotted line), the lowest five singular values of Q (grey lines), and the solution vector norm $\|\xi\|_2$ (circles). (c) is same as (b) but for Stokes. Note that in the Laplace case (b) there is one singular value smaller than the others, whereas for Stokes (c) there are three such singular values.

Continuing in this way, the full set of discrepancy conditions (2.12)–(2.15) gives the linear system

$$Q\xi = \mathbf{g}, \quad (2.19)$$

where $\mathbf{g} \in \mathbb{R}^{4m}$ stacks the four vectors of discrepancies sampled at the collocation points, while the matrix $Q = [Q_1; Q_2; Q_3; Q_4]$ consists of four block rows, each of size $m \times M$. By reading from (2.18), one sees that Q_1 has elements $(Q_1)_{ij} = \phi_j(\mathbf{x}_{iL} + \mathbf{e}_1) - \phi_j(\mathbf{x}_{iL})$. Analogously, $(Q_2)_{ij} = \frac{\partial \phi_j}{\partial n}(\mathbf{x}_{iL} + \mathbf{e}_1) - \frac{\partial \phi_j}{\partial n}(\mathbf{x}_{iL})$, $(Q_3)_{ij} = \phi_j(\mathbf{x}_{iD} + \mathbf{e}_2) - \phi_j(\mathbf{x}_{iD})$, and finally $(Q_4)_{ij} = \frac{\partial \phi_j}{\partial n}(\mathbf{x}_{iD} + \mathbf{e}_2) - \frac{\partial \phi_j}{\partial n}(\mathbf{x}_{iD})$.

We solve the (generally rectangular) system (2.19) in the least squares sense, which corresponds to minimizing a weighted L^2 -norm of the discrepancy error (we find in practice that there is no advantage to incorporating the square-root of the weights associated with the collocation nodes). As is well known in the MFS community [43, 21], Q becomes exponentially ill-conditioned as M grows; intuitively this follows from the exponential decay of the singular values of the single-layer operator from the proxy radius R_p to the circle radius $r_{\mathcal{B}}$ containing \mathcal{B} . Thus, a least squares solver is needed which can handle rank-deficiency stably, i.e. return a small-norm solution when one exists. In this case the ill-conditioning causes no loss of accuracy, and, even though there is instability in the vector ξ , for the evaluation of v errors close to machine precision ($\varepsilon_{\text{mach}}$) are achieved [4].

REMARK 2.5. For this and subsequent direct solves we use MATLAB's `linsolve` (with the option `RECT=true` to prevent the less accurate `LU` from being used in the square case), which uses column-pivoted `QR` to find the so-called “basic” solution [63, Sec. 5.7] having at most r nonzero entries, where r is the numerical rank; this is close to having minimum norm [25, Sec. 5.5].

We illustrate this with a simple numerical test, in which g is the discrepancy of a known harmonic function v of typical magnitude $\mathcal{O}(1)$ and with sufficiently distant singularity. Once Q is filled and (2.19) solved, the numerical solution is evaluated via (2.17), and the maximum error at 100 random target points in \mathcal{B} is taken (after removing an overall constant error, expected from Prop. 2.2). Fig. 2.1(a) shows exponential convergence in this error vs the number of wall nodes m . Fig. 2.1(b) shows exponential convergence with respect to M , the number of proxy points, with a rate slightly exceeding the predicted rate. It is clear that whenever $m \geq 20$ and $M \geq 70$ the norm $\|\xi\|_2$ remains $\mathcal{O}(1)$, and around 15 digits of accuracy result.

The decaying lowest few singular values of Q are also shown in panel (b): in particular there is *one* singular value decaying faster than all others. It is easy to verify that this corresponds to the null-space of

the BVP (Prop. 2.2); indeed we have tested that its right singular vector is approximately constant, and generates via (2.17) the constant function in \mathcal{B} to within $\mathcal{O}(\varepsilon_{\text{mach}})$ relative error. Likewise, the consistency condition in Prop. 2.2 manifests itself in $\text{Nul } Q^T$: let \mathbf{w} be the vector that applies the discretization of this consistency condition to a vector \mathbf{g} , namely

$$\mathbf{w} := [\mathbf{0}_m; \mathbf{w}_L; \mathbf{0}_m; \mathbf{w}_D] \in \mathbb{R}^{4m}, \quad (2.20)$$

where semicolons indicate vertical stacking, $\mathbf{w}_L, \mathbf{w}_D \in \mathbb{R}^m$ are the vectors of weights corresponding to the collocation nodes on L, D , and $\mathbf{0}_m \in \mathbb{R}^m$ is the zero vector. Then we expect that

$$\mathbf{w}^T Q \approx \mathbf{0}_M^T, \quad (2.21)$$

and indeed observe numerically that $\|\mathbf{w}^T Q\|_2 = \mathcal{O}(\varepsilon_{\text{mach}})$ once $m \geq 20$. In summary, although the matrix Q is generally rectangular and ill-conditioned, it also inherits both aspects of the unit nullity of the empty BVP that it discretizes.

REMARK 2.6. *A different scheme is possible in which Q would be square, and (modulo a nullity of one, as above) well-conditioned, based on a “tic-tac-toe” set of layer potentials (see [5, Sec. 4.2] in the Helmholtz case). However, we recommend the above proxy point version, since i) the matrix Q is so small that handling its ill-conditioning is very cheap, ii) the tic-tac-toe scheme demands close-evaluation quadratures for its layer potentials, and iii) the tic-tac-toe scheme is more complicated.*

2.2. Extended linear system for the conduction problem. We now treat the above empty BVP solution scheme as a component in a scheme for the periodic BVP (2.5)–(2.10). Simply put, we take standard potential theory for the Laplace equation [45, Ch. 6], and augment this by enforcing periodic boundary conditions. Given a density function τ on the inclusion boundary $\partial\Omega$, recall that the single-layer potential, evaluated at a target point \mathbf{x} , is defined by

$$v = (\mathcal{S}_{\partial\Omega}\tau)(\mathbf{x}) := \int_{\partial\Omega} G(\mathbf{x}, \mathbf{y})\tau(\mathbf{y})ds_{\mathbf{y}} = \frac{1}{2\pi} \int_{\partial\Omega} \log \frac{1}{\|\mathbf{x} - \mathbf{y}\|} \tau(\mathbf{y})ds_{\mathbf{y}}, \quad \mathbf{x} \in \mathbb{R}^2. \quad (2.22)$$

Using $n^{\mathbf{x}}$ to indicate the outward normal at \mathbf{x} , this potential obeys the jump relation [45, Thm. 6.18]

$$v_n^{\pm} := \lim_{h \rightarrow 0^+} n^{\mathbf{x}} \cdot \nabla(\mathcal{S}_{\partial\Omega}\tau)(\mathbf{x} \pm hn^{\mathbf{x}}) = ((\mp \tfrac{1}{2} + D_{\partial\Omega, \partial\Omega}^T)\tau)(\mathbf{x}), \quad (2.23)$$

where $D_{\Gamma', \Gamma}^T$ denotes the usual transposed double-layer operator from a general source curve Γ to a target curve Γ' , defined by

$$(D_{\Gamma', \Gamma}^T\tau)(\mathbf{x}) = \int_{\Gamma} \frac{\partial G(\mathbf{x}, \mathbf{y})}{\partial n^{\mathbf{x}}} \tau(\mathbf{y})ds_{\mathbf{y}} = \frac{1}{2\pi} \int_{\Gamma} \frac{-(\mathbf{x} - \mathbf{y}) \cdot n^{\mathbf{x}}}{\|\mathbf{x} - \mathbf{y}\|^2} \tau(\mathbf{y})ds_{\mathbf{y}}, \quad \mathbf{x} \in \Gamma'. \quad (2.24)$$

The integral implied by the self-interaction operator $D_{\partial\Omega, \partial\Omega}^T$ is to be interpreted in the principal value sense.

Our representation for the solution sums the standard single-layer potential over the 3×3 copies closest to the origin, and adds an auxiliary basis as in (2.17),

$$u = \mathcal{S}_{\partial\Omega}^{\text{near}}\tau + \sum_{j=1}^M \xi_j \phi_j, \quad \text{where } (\mathcal{S}_{\partial\Omega}^{\text{near}}\tau)(\mathbf{x}) := \sum_{m_1, m_2 \in \{-1, 0, 1\}} \int_{\partial\Omega} G(\mathbf{x}, \mathbf{y} + m_1 \mathbf{e}_1 + m_2 \mathbf{e}_2) \tau(\mathbf{y})ds_{\mathbf{y}}. \quad (2.25)$$

Our unknowns are the (physical charge) density τ and the auxiliary vector ξ . Substituting (2.25) into the Neumann boundary condition (2.6), and using the exterior jump relation on the central copy ($m_1 = m_2 = 0$) only, gives our first block row,

$$(-\tfrac{1}{2} + D_{\partial\Omega, \partial\Omega}^{\text{near}, T})\tau + \sum_{j=1}^M \xi_j \frac{\partial \phi_j}{\partial n} \Big|_{\partial\Omega} = 0, \quad (2.26)$$

where, as before, the “near” superscript denotes summation over source images as in (2.25).

The second block row arises as follows. Consider the substitution of (2.25) into the first discrepancy equation (2.7): there are nine source copies, each of which interacts with the L and R walls, giving 18 terms. However, the effect of the rightmost six sources on R is cancelled by the effect of the leftmost six sources on L , leaving only six terms, as in [5, Fig. 4(a)–(b)]. All of these surviving terms involve *distant* interactions (the distances exceed one period if Ω is contained in \mathcal{B}). Similar cancellations occur in the remaining three equations (2.8)–(2.10). The resulting four subblocks are

$$\sum_{m_2 \in \{-1, 0, 1\}} (S_{R, \partial\Omega - \mathbf{e}_1 + m_2 \mathbf{e}_2} - S_{L, \partial\Omega + \mathbf{e}_1 + m_2 \mathbf{e}_2}) \tau + \sum_{j=1}^M (\phi_j|_R - \phi_j|_L) \xi_j = p_1 \quad (2.27)$$

$$\sum_{m_2 \in \{-1, 0, 1\}} (D_{R, \partial\Omega - \mathbf{e}_1 + m_2 \mathbf{e}_2}^\top - D_{L, \partial\Omega + \mathbf{e}_1 + m_2 \mathbf{e}_2}^\top) \tau + \sum_{j=1}^M \left(\frac{\partial \phi_j}{\partial n} \Big|_R - \frac{\partial \phi_j}{\partial n} \Big|_L \right) \xi_j = 0 \quad (2.28)$$

$$\sum_{m_1 \in \{-1, 0, 1\}} (S_{U, \partial\Omega + m_1 \mathbf{e}_1 - \mathbf{e}_2} - S_{D, \partial\Omega + m_1 \mathbf{e}_1 + \mathbf{e}_2}) \tau + \sum_{j=1}^M (\phi_j|_U - \phi_j|_D) \xi_j = p_2 \quad (2.29)$$

$$\sum_{m_1 \in \{-1, 0, 1\}} (D_{U, \partial\Omega + m_1 \mathbf{e}_1 - \mathbf{e}_2}^\top - D_{D, \partial\Omega + m_1 \mathbf{e}_1 + \mathbf{e}_2}^\top) \tau + \sum_{j=1}^M \left(\frac{\partial \phi_j}{\partial n} \Big|_U - \frac{\partial \phi_j}{\partial n} \Big|_D \right) \xi_j = 0. \quad (2.30)$$

REMARK 2.7 (Wall intersection). *The cancellation of all near interactions in (2.27)–(2.30) is due to the 3×3 neighbor summation in the representation (2.25). Furthermore this cancellation allows an accurate solution even when Ω intersects $\partial\mathcal{B}$, without the need for specialized quadratures, as long as the 3×3 copies of Ω account for all of the images inside, and near to, \mathcal{B} (see Fig. 2.2(a)). Informally, the unit cell walls are “invisible” to the inclusions. If Ω is elongated such that the last condition cannot hold, one would need to include all of the (pieces of) images of Ω that fell within, and near to, \mathcal{B} , needing different bookkeeping in the above formulae; we leave this case for future work.*

(2.26)–(2.30) form a set of coupled integral-algebraic equations, where the only discrete aspect is that of the $\mathcal{O}(1)$ proxy points.¹ It is natural to ask how the unit nullity of the BVP manifests itself in the solution space for the pair (τ, ξ) . Are there pairs (τ, ξ) with no effect on u , enlarging the nullspace? It turns out that the answer is no, and that the one-dimensional nullspace (constant functions) is spanned purely by ξ , as a little potential theory now shows.

LEMMA 2.8. *In the solution to (2.26)–(2.30), τ is unique.*

Proof. Let (τ, ξ) be the difference between any two solutions to (2.26)–(2.30). Let v be the representation (2.25) using this (τ, ξ) , both in $\mathcal{B} \setminus \bar{\Omega}$ but also inside Ω . Then by construction v is a solution to the homogeneous BVP (2.5)–(2.10), i.e. with $p_1 = p_2 = 0$, thus by Prop. 2.1, v is constant in $\mathcal{B} \setminus \bar{\Omega}$. Thus by the continuity of the single-layer potential [45, Theorem 6.18], the interior limit of v on $\partial\Omega$ is constant. However, v is harmonic in Ω , so v is constant in Ω . By (2.23), $v_n^+ - v_n^- = -\tau$, but we have just shown that both v_n^+ and v_n^- vanish, so $\tau \equiv 0$. \square

2.2.1. Discretization of the ELS. We discretize (2.26), the first row of the system, using a set of quadrature nodes $\{\mathbf{x}_i\}_{i=1}^N$ on $\partial\Omega$ and weights $\{w_i\}_{i=1}^N$ such that

$$\int_{\partial\Omega} f(\mathbf{y}) ds_{\mathbf{y}} \approx \sum_{i=1}^N w_i f(\mathbf{x}_i)$$

holds to high accuracy for smooth functions f . In practice, when $\partial\Omega$ is parametrized by a 2π -periodic function $\mathbf{x}(t)$, $0 \leq t < 2\pi$, then using the periodic trapezoid rule in t gives $\mathbf{x}_i = \mathbf{x}(2\pi i/N)$ and $w_i = (2\pi/N) \|\mathbf{x}'(2\pi i/N)\|$. Since D^\top has a smooth kernel, we apply Nyström discretization [45, Sec. 12.2] to the integral equation (2.26) using these nodes, to get our first block row

$$A\boldsymbol{\tau} + B\xi = 0, \quad (2.31)$$

¹It would also be possible to write a purely continuous version by replacing the proxy circle by a single-layer potential; however, sometimes an intrinsically discrete basis $\{\phi_j\}$ is useful [5, 13].

where $\boldsymbol{\tau} = \{\tau_i\}_{i=1}^N$ is a vector of density values at the nodes, and $A \in \mathbb{R}^{N \times N}$ has entries

$$A_{ij} = -\frac{1}{2}\delta_{ij} + \sum_{m_1, m_2 \in \{-1, 0, 1\}} \frac{\partial G(\mathbf{x}_i, \mathbf{x}_j + m_1 \mathbf{e}_1 + m_2 \mathbf{e}_2)}{\partial n^{\mathbf{x}_i}} w_j, \quad (2.32)$$

where δ_{ij} is the Kronecker delta. Here, for entries $i = j$ the standard diagonal limit of the kernel $\partial G(\mathbf{x}_i, \mathbf{x}_i)/\partial n^{\mathbf{x}_i} = -\kappa(\mathbf{x}_i)/4\pi$ is needed, where $\kappa(\mathbf{x})$ is the signed curvature at $\mathbf{x} \in \partial\Omega$. The matrix $B \in \mathbb{R}^{N \times M}$ has entries $B_{ij} = \partial \phi_j(\mathbf{x}_i)/\partial n^{\mathbf{x}_i}$.

For the second block row (2.27)–(2.30) we use the above quadrature for the source locations of the operators, and enforce the four equations on the wall collocation nodes \mathbf{x}_{iL} , $\mathbf{x}_{iL} + \mathbf{e}_1$, \mathbf{x}_{iD} , $\mathbf{x}_{iD} + \mathbf{e}_2$, $i = 1, \dots, m$, to get

$$C\boldsymbol{\tau} + Q\xi = \mathbf{g}, \quad (2.33)$$

with the macroscopic driving (p_1, p_2) encoded by the right-hand side vector

$$\mathbf{g} = [p_1 \mathbf{1}_m; \mathbf{0}_m; p_2 \mathbf{1}_m; \mathbf{0}_m] \in \mathbb{R}^{4m}, \quad (2.34)$$

where $\mathbf{1}_m \in \mathbb{R}^m$ is the vector of ones. The matrix Q is precisely as in (2.19). Rather than list formulae for all four blocks in $C = [C_1; C_2; C_3; C_4]$, we have $(C_1)_{ij} = \sum_{m_2 \in \{-1, 0, 1\}} (G(\mathbf{x}_{iL} + \mathbf{e}_1, \mathbf{x}_j - \mathbf{e}_1 + m_2 \mathbf{e}_2) - G(\mathbf{x}_{iL}, \mathbf{x}_j + \mathbf{e}_1 + m_2 \mathbf{e}_2)) w_j$, with the others filled analogously. Stacking the two block rows (2.31) and (2.33) gives the $(N+4m)$ -by- $(N+M)$ extended linear system (1.2), which we emphasize is just a standard discretization of the BVP conditions (2.6)–(2.10).

For small problems (N less than a few thousand), the ELS is most simply solved by standard dense direct methods; for larger problems a Schur complement must be taken in order to solve iteratively, as presented in Section 2.4. First we perform a numerical test of the direct method.

2.3. Numerical tests using direct solution of the ELS. Example 1. We define a smooth “worm” shaped inclusion which crosses the unit cell walls L and R by $\mathbf{x}(t) = (0.7 \cos t, 0.15 \sin t + 0.3 \sin(1.4 \cos t))$. The solution of the periodic Neumann Laplace BVP with external driving $\mathbf{p} = (1, 0)$ is shown in Fig. 2.2(a). Here u is evaluated via (2.25) (using the standard far-field Nyström quadrature) both inside \mathcal{B} (where it is accurate), and, to show the nature of the representation, out to the proxy circle (where it is certainly inaccurate). Recall that the proxy sources must represent the layer potentials due to the infinite lattice of copies which excludes the 3×3 central block. Note that two tips of copies in this set slightly penetrate the proxy circle, violating the condition in Thm. 2.3 that the function the proxy sources represent be analytic in the closed ball of radius R_p . We find in practice that such slight geometric violations do not lead to problematic growth in $\|\xi\|$, but that larger violations can induce a large $\|\xi\|$ which limits achievable accuracy. Panel (c) shows the convergence of errors in u (at the pair of points shown) to their numerically-converged value, fixing converged values for M and m . Convergence to around 13 digits (for $N = 140$, $M = 70$) is apparent, and the solution time is 0.03 seconds.

As an independent verification of the method, we construct a known solution to a slightly generalized version of (2.5)–(2.10), where (2.6) is replaced by *inhomogeneous* data $u_n = f$ and a general discrepancy g is allowed. We choose the known solution

$$u_{\text{ex}}(\mathbf{x}) = \sum_{m_1, m_2 \in \{-2, -1, \dots, 2\}} n_0 \cdot \nabla G(\mathbf{x}, \mathbf{y}_0 + m_1 \mathbf{e}_1 + m_2 \mathbf{e}_2),$$

where the central dipole has direction n_0 , and location \mathbf{y}_0 chosen inside Ω and far from its boundary (so that the induced data f is smooth). The grid size of 5×5 (some of which is shown by $*$ symbols in Fig. 2.2(a)) is chosen so that g is sufficiently smooth (which requires at least 3×3), and so that the periodizing part ξ is nontrivial. From u_{ex} the right-hand side functions f and g are then evaluated at nodes, the ELS solved directly, the numerical solution (2.25) evaluated, and the difference at two points compared to its known exact value. The resulting N -convergence to 13 digits is shown in Fig. 2.2(c). The convergence rate is slower than before, due to the unavoidable closeness of the dipole source to $\partial\Omega$.

Finally, the grey lines in panel (e) show that, as with Q , there is one singular value of E which is much smaller than the others, reflecting the unit nullity of the underlying BVP (Prop. 2.1). Also apparent is the fact that, despite the exponential ill-conditioning, the solution norm remains bounded once M -convergence has occurred.

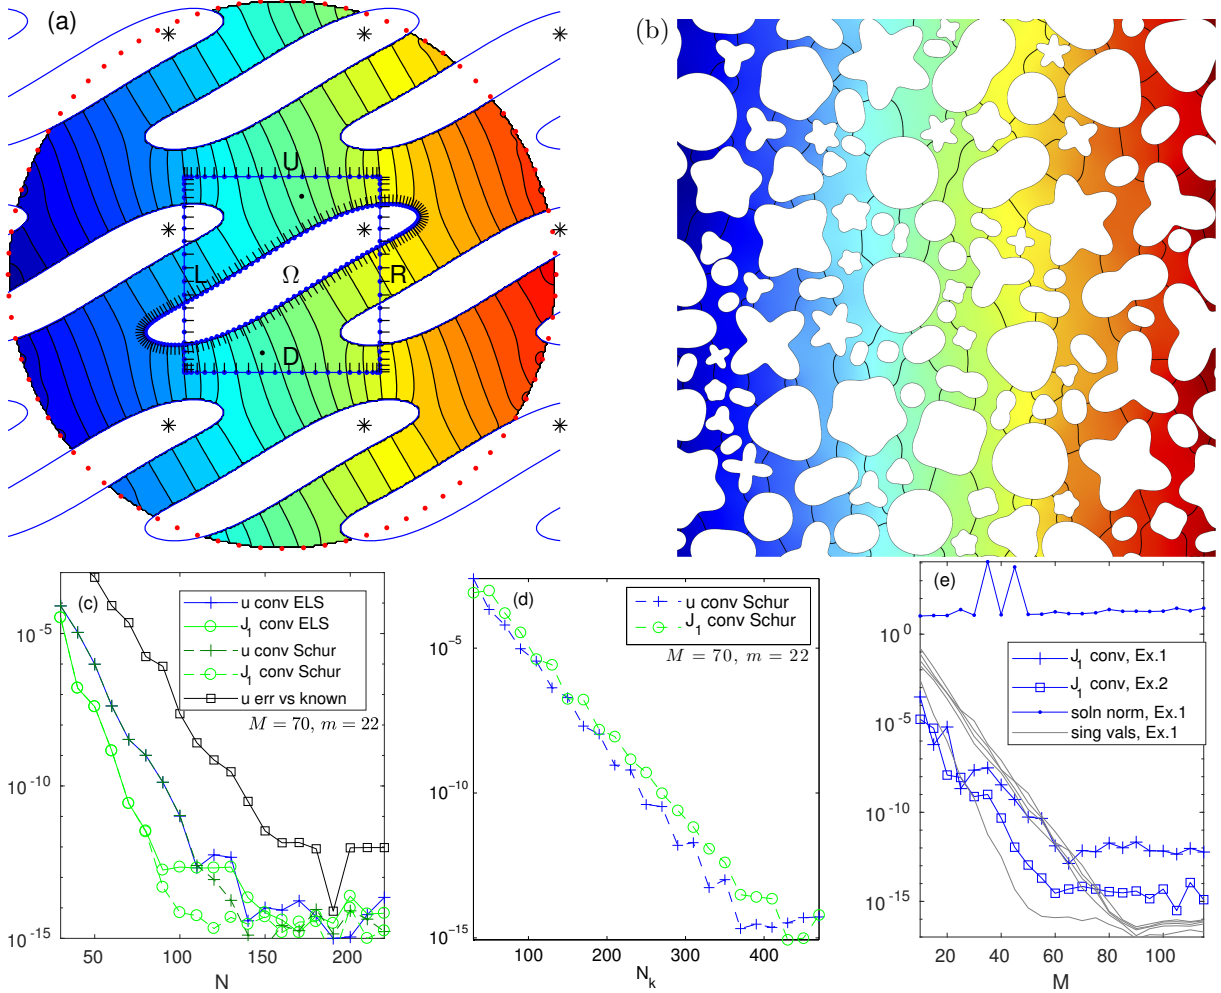


FIG. 2.2. Periodic Laplace Neumann tests driven by external driving $\mathbf{p} = (1, 0)$. (a) Solution potential u contours for “worm” inclusion (Example 1), with $N = 140$, $M = 70$, and $R_p = 1.4$. The $m = 22$ nodes per wall and nodes on $\partial\Omega$ are also shown (with normals), proxy points (red dots), test points (two black dots), and some of the 5×5 grid of dipoles generating a known solution (* symbols). The representation for u is only accurate inside the unit cell. (b) Solution with $K = 100$ inclusions (Example 2), with $N_k = 400$ unknowns per inclusion, by iterative solution of (2.38). A single unit cell is shown. (c) N -convergence of error in difference of u at the two test points, and of flux J_1 computed via (2.43), relative to their values at $N = 230$ ($M = 70$ is fixed), for Example 1. Squares show error convergence in u (difference at the test points) in the case of known u_{ex} due to the dipole grid. Solid lines are for direct solution of the ELS, dashed lines for the iterative solution of (2.38). (d) Convergence with N_k for Example 2, for pointwise u (+ symbols) and flux J_1 (circles). (e) M -convergence of flux J_1 error, for Example 1 (+ signs) and Example 2 (squares), with other parameters converged. For $K = 1$ the lowest six singular values of E are also shown (grey lines), and the solution norm $\|[\tau; \xi]\|_2$ (dots).

2.4. Schur complement system and its iterative solution. When N is large, solving the ill-conditioned rectangular ELS (1.2) is impractical. We would like to use a Schur complement in the style of (1.3) to create an equivalent $N \times N$ system, which we do in Sec. 2.4.1. Furthermore, in order to use Krylov subspace iterative methods with known convergence rates, we would like to remove the nullspace to make this well-conditioned, which we do in Sec. 2.4.2. We will do both these tasks via low-rank perturbation of certain blocks of (1.3) before applying the Schur complement.

In what follows, Q^+ is the pseudoinverse [25, Sec. 5.5] of Q , i.e. the linear map that recovers a small-norm solution (if one exists) to $Q\xi = \mathbf{g}$ via $\xi = Q^+\mathbf{g}$. The obstacle to using (1.3) as written is that Q inherits a consistency condition from the empty BVP so that Q has one smooth vector in its left null-space (\mathbf{w} ; see (2.21)). However the range of C does not respect this condition, thus Q^+C has a huge 2-norm (which we find numerically is at least 10^{16}). We first need to show that the range of a

rank-deficient matrix Q may be enlarged by a rank- k perturbation, a rectangular version of results about singular square matrices in [73].

LEMMA 2.9. *Let $Q \in \mathbb{R}^{m \times n}$ have a k -dimensional nullspace. Let $R \in \mathbb{R}^{n \times k}$ have full-rank projection onto $\text{Nul } Q$, i.e. if N has columns forming a basis for $\text{Nul } Q$ then $R^T N \in \mathbb{R}^{k \times k}$ is invertible. Let $V \in \mathbb{R}^{m \times k}$ be arbitrary. Then $\text{Ran}(Q + VR^T) \supset \text{Ran } Q \oplus \text{Ran } V$, i.e. the range now includes that of V .*

Proof. We need to check that $(Q + VR^T)x = Qx_0 + V\alpha_0$ has a solution x for all given pairs $x_0 \in \mathbb{R}^n$, $\alpha_0 \in \mathbb{R}^k$. Recalling that $R^T N$ is invertible, by substitution one may check that $x = x_0 + N(R^T N)^{-1}(\alpha_0 - R^T x_0)$ is an explicit such solution. \square

REMARK 2.10. *With additional conditions $m \geq n$, and that V has full-rank projection onto a part of $\text{Nul } Q^T$ (i.e. if W has columns forming a basis for a k -dimensional subspace of $\text{Nul } Q^T$, then $W^T V \in \mathbb{R}^{k \times k}$ is invertible), we get incidentally that $Q + VR^T$ has trivial nullspace (generalizing [73, Sec. 3]). The proof is as follows. Let $x \in \mathbb{R}^n$ solve the homogeneous equation $(Q + VR^T)x = 0$. Then $0 = W^T Qx = -(W^T V)R^T x$, but $W^T V$ is invertible, so $R^T x = 0$. Thus the homogeneous equation becomes $Qx = 0$, which means there is an $\alpha \in \mathbb{R}^k$ such that $x = N\alpha$. Thus $R^T N\alpha = 0$, but $R^T N$ is invertible, so that $\alpha = 0$, so $x = 0$.*

We also need the fact that a block-column operation allows Q to be perturbed as above while changing the ELS solution space in a known way. The proof is simple to check.

PROPOSITION 2.11. *Let A, B, C and Q be matrices, and let P be a matrix with as many rows as A has columns and as many columns as B has rows. Then the pair $(\tilde{\tau}, \xi)$ solves the block system*

$$\begin{bmatrix} A & B + AP \\ C & Q + CP \end{bmatrix} \begin{bmatrix} \tilde{\tau} \\ \xi \end{bmatrix} = \begin{bmatrix} \mathbf{0} \\ \mathbf{g} \end{bmatrix} \quad (2.35)$$

if and only if the pair (τ, ξ) , with $\tau = \tilde{\tau} + P\xi$, solves the original ELS (1.2).

2.4.1. An equivalent square system preserving the nullspace. Armed with the above, a method to “fold” the Neumann Laplace ELS into an equivalent square system is as follows:

1. Set the proxy coefficient vector $R = (1/M)\mathbf{1}_M$, and the discrete density vector $H = \mathbf{1}_N$. Create low-rank perturbed matrix blocks $\tilde{B} := B + AHR^T$ and $\tilde{Q} := Q + CHR^T$.
2. Solve for the vector $\tilde{\tau}$ in the $N \times N$ Schur complement linear system

$$(A - \tilde{B}\tilde{Q}^+C)\tilde{\tau} = -\tilde{B}\tilde{Q}^+\mathbf{g}. \quad (2.36)$$

More precisely, since (due to the numerical ill-conditioning in \tilde{Q}) multiplying by \tilde{Q}^+ would lose accuracy due to rounding error, instead solve the small systems $\tilde{Q}X = C$ for X , and $\tilde{Q}\mathbf{y} = \mathbf{g}$ for \mathbf{y} , heeding Remark 2.5. From them, build $A - \tilde{B}X$ and $-\tilde{B}\mathbf{y}$, which are respectively the system matrix and right-hand side for (2.36). This large square system may then be solved iteratively (see section 2.5).

3. Recover the proxy coefficients via $\xi = \mathbf{y} - X\tilde{\tau}$.
4. Recover the density via $\tau = \tilde{\tau} + HR^T\xi$.

Note that in Step 1 the prefactor $1/M$ leads to correct quadrature scaling, so that HR^T has similar 2-norm to the other matrices (also recommended in [73]).

THEOREM 2.12. *Let \mathbf{g} encode the driving as in (2.34). Then for all sufficiently large N, M , and m , any pair (τ, ξ) produced by the above procedure performed in exact arithmetic solves the original ELS (1.2).*

Proof. By rotational invariance, a constant single-layer density on a circle generates a constant potential inside, and inserting R into (2.17) gives the periodic trapezoid quadrature approximation to such a potential, thus generates discrepancy near zero (in fact exponentially small in M). Thus for all sufficiently large M , R is not orthogonal to $\text{Nul } Q$. Applying Lemma 2.9 with $k = 1$ gives that the range of \tilde{Q} includes the discrepancy vector $V = CH$ produced by the constant density H . Thus, to show that the range of \tilde{Q} includes all smooth vectors, i.e. that it does not have the consistency condition (2.21), one needs to check that $\mathbf{w}^T V = \mathbf{w}^T CH \neq 0$ which is done in Lemma 2.13 below. Thus the systems $\tilde{Q}X = C$ and $\tilde{Q}\mathbf{y} = \mathbf{g}$ are consistent for any C and \mathbf{g} , so that the Schur complement is well defined. Finally, Prop. 2.11, using the rank-1 matrix $P = HR^T$, insures that step 4 recovers a solution to (1.2). \square

For the missing technical lemma, we first need a form of Gauss' law, stating that for any density τ on a curve $\partial\Omega$ the single-layer potential $v = \mathcal{S}_{\partial\Omega}\tau$ generates flux equal to minus the total charge, i.e.

$$\int_{\partial\mathcal{K}} v_n = - \int_{\partial\Omega} \tau ds \quad (2.37)$$

where $\partial\mathcal{K}$ is the boundary of some open domain \mathcal{K} containing $\partial\Omega$. This may be proved by combining the jump relation $v_n^+ - v_n^- = -\tau$ (from (2.23)) with the fact that, since v is a Laplace solution in $\mathbb{R}^2 \setminus \partial\Omega$, $\int v_n = 0$ taken over the boundaries of Ω and of $\mathcal{K} \setminus \bar{\Omega}$.

LEMMA 2.13. *Let the matrix C be defined as in Section 2.2, and \mathbf{w} be as in (2.20). Then, for all N and m sufficiently large, $\mathbf{w}^T C \mathbf{1}_N \neq 0$.*

Proof. Let $d = [d_1; d_2; d_3; d_4]$ be the discrepancy of the potential $v = \mathcal{S}_{\partial\Omega}^{\text{near}} \tau$ for density $\tau \equiv 1$. As used in the proof of Prop. 2.2, the flux (left hand side of (2.37)) out of the unit cell \mathcal{B} equals $\int_L d_2 ds + \int_D d_4 ds$, which by (2.37) (and noting that only one of the nine source terms in $\mathcal{S}_{\partial\Omega}^{\text{near}}$ lies within \mathcal{B}) equals minus the perimeter $-\int_{\partial\Omega} 1 ds = -|\partial\Omega| < 0$. What we seek is the discretization of this statement about the PDE. If \mathbf{d} is the discrepancy of v sampled at the wall nodes, then its quadrature approximation is $\mathbf{d} \approx C \mathbf{1}_N$, and so (as discussed above (2.20)) the quadrature approximation to $\int_L d_2 ds + \int_D d_4 ds$ is $\mathbf{w}^T \mathbf{d} \approx \mathbf{w}^T C \mathbf{1}_N$. Thus, as N and m tend to infinity, the latter converges to $-|\partial\Omega|$. \square

Theorem 2.12 justifies rigorously one procedure to create a square system equivalent to the ELS (1.2). However, by the equivalence in Prop. 2.11, the system matrix $A - \tilde{B}X$ at the heart of the procedure is singular, as it inherits the unit nullity of the ELS, which itself derives from the unit nullity of the Laplace Neumann BVP. Since the convergence of iterative solvers for singular matrices is a subtle matter [10], this motivates the following improved variant which removes the nullspace.

2.4.2. A well-conditioned square system. The following simpler variant creates a non-singular square system from the ELS; its proof is more subtle. It is what we recommend for the iterative solution of the periodic Neumann Laplace problem, and test numerically:

1. Set the constant proxy coefficient vector $R = (1/M)\mathbf{1}_M$, the constant discrete density vector $H = \mathbf{1}_N$, and $\tilde{Q} := Q + CHR^T$.
2. Solve for the vector $\boldsymbol{\tau}$ in the $N \times N$ Schur complement linear system

$$\tilde{A}_{\text{per}} \boldsymbol{\tau} := (A - B\tilde{Q}^+ C) \boldsymbol{\tau} = -B\tilde{Q}^+ \mathbf{g}, \quad (2.38)$$

where, as before, one solves the small systems $\tilde{Q}X = C$ and $\tilde{Q}\mathbf{y} = \mathbf{g}$, to build the large system matrix $\tilde{A}_{\text{per}} = A - BX$ and right-hand side $-B\mathbf{y}$ for (2.38), which may then be solved iteratively (see section 2.5).

3. Recover the proxy coefficients via $\xi = \mathbf{y} - X\boldsymbol{\tau}$.

THEOREM 2.14. *Let \mathbf{g} encode the driving as in (2.34). Then for all N , M and m sufficiently large, the pair $(\boldsymbol{\tau}, \xi)$ produced by the above procedure is unique, and solves the ELS (1.2) with residual of order the quadrature error on boundaries.*

Proof. First note that (2.38) is the Schur complement of the perturbed ELS

$$\begin{bmatrix} A & B \\ C & Q + CHR^T \end{bmatrix} \begin{bmatrix} \boldsymbol{\tau} \\ \xi \end{bmatrix} = \begin{bmatrix} 0 \\ \mathbf{g} \end{bmatrix}. \quad (2.39)$$

In particular one may check that if $\boldsymbol{\tau}$ solves (2.38) then $(\boldsymbol{\tau}, \xi)$, with ξ as in Step 3, solves (2.39). Given such a solution $(\boldsymbol{\tau}, \xi)$, define the potential generated by the usual representation

$$v(\mathbf{x}) := \sum_{m_1, m_2 \in \{-1, 0, 1\}} \sum_{i=1}^N w_i G(\mathbf{x}, \mathbf{x}_i + m_1 \mathbf{e}_1 + m_2 \mathbf{e}_2) \tau_i + \sum_{j=1}^M \xi_j \phi_j(\mathbf{x}),$$

i.e. the quadrature approximation to (2.25). Note that $C\boldsymbol{\tau} + Q\xi$ then approximates the discrepancy of v at the nodes on the unit cell walls. The first row of (2.39) states that $v_n = 0$ on $\partial\Omega$, and since v is harmonic in $\mathcal{B} \setminus \bar{\Omega}$, the net flux of v through $\partial\mathcal{B}$ is zero. This means that the discrepancy of v obeys the same consistency condition as in Prop. 2.2, which when discretized on the walls gives $\mathbf{w}^T (C\boldsymbol{\tau} + Q\xi) = 0$, where \mathbf{w} is defined by (2.20). Subtracting this from the second row of (2.39) left-multiplied by \mathbf{w}^T leaves

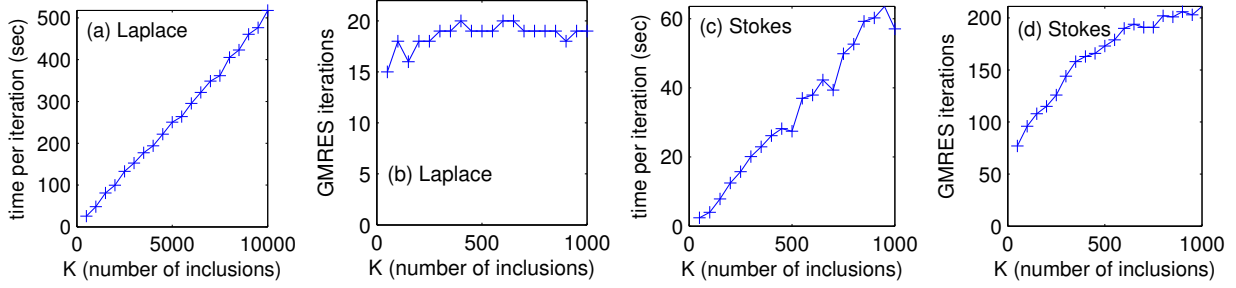


FIG. 2.3. *Scaling of computation time per iteration, and iteration counts, with K , the number of inclusions. In all four plots, we set $N_k = 256$. (a) CPU time per iteration for the Laplace problem with random star-shaped obstacles as in Figs. 1.2 and (2.2)(b). (b) Number of GMRES iterations to reach a relative residual of $\epsilon = 10^{-14}$, for circles with $f_{\text{clup}} = 10$. (c) and (d) are the same as (a) and (b), but for the Stokes problem.*

only the expression $(\mathbf{w}^T C H) R^T \xi = \mathbf{w}^T \mathbf{g}$. By Lemma 2.13, $\mathbf{w}^T C H \neq 0$, and by computation $\mathbf{w}^T \mathbf{g} = 0$. Thus $R^T \xi = 0$, so the pair also solves the original ELS (1.2). Thus Lemma 2.8 (strictly, its quadrature approximation) holds, so $\boldsymbol{\tau}$ is unique. \square

In contrast to the previous section, the system (2.38) to be solved is *well conditioned* if the non-periodic BIE matrix A is; the unit nullity of (1.2) has been removed by imposing one extra condition $R^T \xi = 0$.

We now test the above procedure for the Laplace Neumann periodic BVP of Example 1 (Fig. 2.2(a)). We solve (2.38) iteratively via GMRES with a relative stopping residual tolerance of 10^{-14} . In Step 2 we use `linsolve` as in Remark 2.5, and verify that the resulting norm $\|X\|_2 \approx 9$ is not large. Fig. 2.2(c) includes (as dashed lines) the resulting self-convergence of u , and of the flux J_1 (computed as in Sec. 2.6), with other parameters converged as before. The converged values agree to around 10^{-13} . Above 10^{-13} , the errors are identical to those of the full ELS. The condition number of \tilde{A}_{per} is 8.3, and the number of GMRES iterations required was 12, both independent of N and M .

2.5. Multi-inclusion examples and close-to-touching geometries. Generalizing the above to $K > 1$ disjoint inclusions $\{\Omega_k\}_{k=1}^K$ in the unit cell is largely a matter of bookkeeping. The representation (2.25) becomes a 3×3 image sum over single-layer potentials on each inclusion boundary,

$$u = \sum_{k=1}^K \mathcal{S}_{\partial\Omega_k}^{\text{near}} \tau_k + \sum_{j=1}^M \xi_j \phi_j, \quad \text{with } \tau := \{\tau_k\}. \quad (2.40)$$

In particular, the proxy representation is unchanged. Upon discretization using N_k nodes on the k th inclusion boundary, with a total number of unknowns $N := \sum_{k=1}^K N_k$, the A matrix now has a K -by- K block structure where the $-\frac{1}{2}$ identity only appears in the diagonal blocks. For large N , to solve the resulting linear system (2.38) iteratively via GMRES, one needs to apply $\tilde{A}_{\text{per}} = A - BX$ to any given vector $\boldsymbol{\tau}$. We can perform this matrix-vector multiply in $\mathcal{O}(N)$ time, as follows. We apply the off-diagonal of A using the FMM with source charges $\tau_j w_j$, and the diagonal of A as discussed below (2.32). Then, having pre-stored B and X , which needs $\mathcal{O}(MN)$ memory, we compute the correction $-B(X\boldsymbol{\tau})$ using two standard BLAS2 matrix-vector multiplies.

When curves come close (in practice, closer than $5h$, where h is the local node spacing [3, Rmk. 6]), high accuracy demands replacing the native Nyström quadrature (2.32) with special quadrature formulae. Note that this does not add extra unknowns. For this we use a variant of recently developed close-evaluation quadratures for the periodic trapezoid rule with the Laplace single-layer potential [6]. In particular, for barycentric evaluation of Cauchy integrals we use Helsing's exterior formula, and interpolation of the derivative, following Remarks 2 and 3 in [6]; the relevant code is `Cau_closeglobal.m` (see Remark 1.1). Exterior close evaluations build upon this; the single-layer case needs an additional choice of interior point far from the source curve $\partial\Omega_k$, for which we use its centroid. For each of the source curves for which a given target is within $5h$, we subtract off the native contribution from the above FMM evaluation and add in the special close evaluation for that curve.




FIG. 2.4. Efficient evaluation of fluxes (J_1, J_2) using far field interactions alone. (a) Nine terms in L wall integral in J_1 from the 3×3 layer potential image sum in (2.25). (b) Re-interpretation as a sum over targets (nine copies of L) for the potential v . The red dotted line shows the closure of the boundary where flux conservation is applied. (c) Resulting weights of the flux integrals of v on nine wall segments; note all are distant from $\partial\Omega$. (d) The wall weights for J_2 .

REMARK 2.15 (Geometry generation). In all our remaining numerical examples except Examples 4 and 6, we create random geometries with a large number K of inclusions as follows. We generate polar curves of the form $r(\theta) = s(1 + a \cos(w\theta + \phi))$, with ϕ random, a uniform random in $[0, 0.5]$, w randomly chosen from the set $\{2, 3, 4, 5\}$, and s varying over a size ratio of 4. Starting with the largest s , we add in such curves (translated to uniformly random locations in the unit cell), then discard any that intersect. This is repeated in a sequence of decreasing s values until a total of K inclusions are generated. Finally, the s (size) of all inclusions were multiplied by 0.97. This has the effect of making a random geometry with a minimum relative closeness of around 3% of the radius (this is only approximate since it depends on local slope $r'(\theta)$). Helsing–Ojala [40] defines for circles a closeness parameter f_{clup} as the upper bound on the the circumference of the larger circle divided by the minimum distance between the curves. If in their definition one replaces circumference by 2π times the largest radius of a non-circular curve, then in our geometry $f_{\text{clup}} \approx 200$.

Example 2. With $K = 100$ inclusions generated as just described, and an external driving $\mathbf{p} = (1, 0)$, we use the well-conditioned iterative method from Sec. 2.4.2, and the Laplace FMM of Gimbutas–Greengard [24]. Fig. 2.2(b) shows the solution potential, and Fig. 2.2(d) the self-convergence of the potential at a point and of the flux, with respect to N_k . Both achieve at least 13 digits of absolute accuracy (the flux is of size around 0.2, so the relative accuracy is similar). At a fully-converged $N_k = 400$, the solution time was 445 seconds. Fig. 2.2(e) shows the convergence with respect to M , the number of proxy points: this confirms that this convergence rate is at least as good as it is for $K = 1$. In other words, as least for a square unit cell, the M required for close to machine accuracy is around 70, and, as expected, is independent of the complexity of the geometry.

For this example, we verify linear complexity of the scheme in Fig. 2.3(a), for up to $K = 10^4$ inclusions ($N = 2.56 \times 10^6$ total unknowns). In Fig. 1.2, we plot the solution for $K = 10^4$ inclusions and $N_k = 700$ (ie $N = 7 \times 10^6$ total unknowns). The solution requires 69 GMRES iterations and around 10 CPU hours. The absolute flux error of 10^{-10} is estimated by convergence, and comparison to the solution with $N_k = 1000$. (The flux is again around 0.2, so relative error is similar.) The code is found in the directory `multiinclusion_Laplace` (see Remark 1.1).

Example 3. A natural question is how the complexity of the geometry affects the number of GMRES iterations. To address this, we generate simpler random geometries using K circles, again with random sizes of ratio up to 4, but with $f_{\text{clup}} = 10$, which means that curves are not very close to each other. In Fig. 2.3(b) shows that the number of iterations grows very weakly, if at all, with K . The interesting question of the impact of f_{clup} on iteration count we leave as an open question. We note that this has been studied in the non-periodic case [40]; we would expect similar results.

2.6. Computing the effective conductivity tensor. An important task is to compute the effective conductivity $\kappa \in \mathbb{R}^{2 \times 2}$, which expresses how the mean flux depends on the driving. Let $\mathbf{J} := (J_1, J_2)$ be the mean flux, with components

$$J_1 := \int_L u_n ds, \quad \text{and} \quad J_2 := \int_D u_n ds, \quad (2.41)$$

and, recalling the pressure vector $\mathbf{p} = (p_1, p_2)$, the conductivity tensor is defined by Darcy's law

$$\mathbf{J} = \kappa \mathbf{p} . \quad (2.42)$$

As is well known [62, 27], to extract the four elements of κ , it is sufficient to solve two BVPs ("cell problems"), one with $p_1 = 1, p_2 = 0$ (from which one may read off $\kappa_{11} = J_1$ and $\kappa_{21} = J_2$), and the other with $p_1 = 0, p_2 = 1$ (and read off $\kappa_{12} = J_1$ and $\kappa_{22} = J_2$). Note that κ is symmetric [14, Cor. 6.10], hence $|\kappa_{12} - \kappa_{21}|$ provides an independent gauge of numerical accuracy. Also note that in the trivial case of no inclusions (i.e. uniform medium), the tensor takes the value of the background conductivity $\kappa = I$.

For large-scale problems, approximating the integrals (2.41) directly by quadrature on the walls L and D is inconvenient, because, when inclusions intersect the walls, this forces the integral to be broken into intervals and forces close-evaluation quadratures to be used. One could deform the integration paths to avoid inclusions, but finding such a smooth path is complicated, and needs many quadrature nodes, due to having to pass near inclusions. Instead, we propose the following method which pushes all interactions to the far field, and thus requires only a fixed $m \approx 20$ target nodes per wall and no special quadratures.

PROPOSITION 2.16. *Let u be represented by (2.25) and solve the BVP (2.5)–(2.10) in \mathcal{B} . Define $v = \mathcal{S}_{\partial\Omega}\tau$. Then the horizontal flux in (2.41) can be written*

$$J_1 = \sum_{m_1 \in \{-1, 0, 1\}} (m_1 + 2) \left(\int_{U+m_1\mathbf{e}_1+\mathbf{e}_2} v_n ds - \int_{D+m_1\mathbf{e}_1-\mathbf{e}_2} v_n ds \right) + 3 \sum_{m_2 \in \{-1, 0, 1\}} \int_{R+\mathbf{e}_1+m_2\mathbf{e}_2} v_n ds + \int_L \sum_{j=1}^M \frac{\partial \phi_j}{\partial n} \xi_j ds . \quad (2.43)$$

The flux integrals involving v are on the *distant* walls of the 3×3 "super-cell", with locations and weights shown in Fig. 2.4(c). The final term involves a smooth integrand on the original wall L . A similar far-field formula for J_2 is achieved by reflection through the line $x_1 = x_2$, with weights shown in Fig. 2.4(d).

Proof. Substituting (2.25) into J_1 in (2.41) involves a 3×3 sum over density sources (Fig. 2.4(a)) which by translational invariance we reinterpret as a sum over displaced target copies as in panel (b). These nine copies of L form three continuous vertical walls. Since v is harmonic in $\mathbb{R}^2 \setminus \overline{\Omega}$, then $\int_{\Gamma} v_n = 0$ for any closed curve Γ that does not enclose nor touch $\partial\Omega$. However, since $u_n = 0$ on $\partial\Omega$ and $u - v$ is harmonic in a neighborhood of Ω , this also holds if Γ encloses $\partial\Omega$. Thus the flux through each length-3 vertical wall is equal to the flux through its closure to the right along the dotted contour shown in panel (b). Summing these three contour closures, with appropriate normal senses, gives the weights in panel (c), i.e. (2.43). One may check that the result is unaffected by intersections of $\partial\Omega$ with the original unit cell walls. \square

We have tested that this formula matches the naive quadrature of (2.41) when $\partial\Omega$ is far from L . For Example 1, Fig. 2.2(c) and (d) include convergence plots for the flux J_1 as computed by (2.43), showing that it converges at least as fast as do pointwise potential values. For the parameters of panel (a) we find that $|\kappa_{12} - \kappa_{21}| = 6 \times 10^{-14}$, indicating high accuracy of the computed tensor.

3. The case of conducting inclusions. The above section generalizes easily to the case where the inclusions have nonzero conductivity σ ; as before, without loss of generality, the background medium is assigned unit conductivity. This replaces the pair (2.5)–(2.6) by

$$\Delta u = 0 \quad \text{in } \mathcal{B} \setminus \partial\Omega \quad (3.1)$$

$$u^+ - u^- = 0 \quad \text{on } \partial\Omega \quad (3.2)$$

$$u_n^+ - \sigma u_n^- = 0 \quad \text{on } \partial\Omega \quad (3.3)$$

where the superscripts $+$ ($-$) denote exterior (interior) limits as in (2.23). Following [27] we keep the standard single-layer representation (2.25); however, this representation now acquires physical meaning both inside and outside the inclusion. This, plus the continuity of the single-layer potential across $\partial\Omega$ [45, Ch. 6], means that (3.1)–(3.2) are already satisfied, leaving (3.3) as the only remaining equation. Inserting into this the jump relation (2.23) and simplifying gives

$$\left(\frac{1}{2\lambda} + D_{\partial\Omega, \partial\Omega}^{\text{near}, T} \right) \tau + \sum_{j=1}^M \xi_j \frac{\partial \phi_j}{\partial n} \Big|_{\partial\Omega} = 0 , \quad (3.4)$$

c	σ	N	# iters.	CPU time	κ_{calc}	κ_{H94}	κ_{H17}	rel. diff.
100	100	200	56	0.8 s	35.95764629703	35.957646297	35.95764629705	3×10^{-13}
100	1000	200	93	0.8 s	164.1473079002	164.14730790	164.1473079006	2×10^{-12}
1000	100	200	81	0.8 s	36.67257761067	36.672577611	36.67257761078	3×10^{-12}
1000	1000	300	187	2.6 s	243.00597811	243.005978	243.0059781329	8×10^{-11}

TABLE 3.1

Validation of the effective conductivity of an infinite unit square array of discs with radius $r = \frac{1}{2}\sqrt{1-c^{-2}}$ and conductivity σ within a background conductivity of 1 (Example 4). The well-conditioned Schur system of (2.38) was used, with the matrix A given by (2.32) with the factor $-1/2$ replaced by $1/\lambda$. GMRES with a relative tolerance of 10^{-14} was used, although stagnation at up to 4 digits worse than this occurred. κ_{calc} shows the results. κ_{H94} shows published results in [35, Table 2]; κ_{H17} shows results using a slight variant of the code `demo14b.m` due to Helsing [37, Sec. 21.2]. The last column shows the relative differences $|\kappa_{\text{calc}} - \kappa_{\text{H17}}|/\kappa_{\text{H17}}$.

where $\lambda := (\sigma - 1)/(\sigma + 1)$ is a contrast parameter. This equation differs from (2.26) only in the prefactor of the identity term, and becomes identical in the insulating case $\sigma = 0$. Thus, in the discretized matrix A of (2.32) the factor $-1/2$ becomes $1/\lambda$, but no other aspect of the scheme of Sec. 2 changes.

Example 4. The effective conductivity of an infinite square array of conducting discs of radius r is a standard test case. Only a single driving $\mathbf{p} = (1, 0)$ need be applied, because κ is a multiple of the identity. The more challenging cases are when the gap $1 - 2r$ is very small and σ is large [60]. In Table 3.1 we validate against four such results published by Helsing [35, Table 2], also against recomputation with Helsing’s recent RCIP method [37]. Both of these methods use lattice sums to periodize. Radius is controlled by a parameter $c = (1 - 4r^2)^{-1/2}$; for $c = 1000$ the gap is 5×10^{-7} (i.e. $f_{\text{clup}} \approx 10^7$). We are still able to efficiently exploit the periodic trapezoid rule by combining the close-evaluation quadratures discussed in Section 2.5 with a geometrically-graded parameterization of the circle,

$$\mathbf{x}(t) = (r \cos \theta(t), r \sin \theta(t)) , \quad \text{where} \quad \theta'(t) = \alpha \cosh(\beta \sin 2t) , \quad (3.5)$$

which bunches quadrature points by a factor of order e^β around the four close-to-touching regions. We fix $\beta = -\frac{1}{2} \log(1 - 2r)$, based on theory [60] implying that τ is smooth on scales as large as $(1 - 2r)^{1/2}$. The constant α in (3.5) is evaluated numerically by enforcing $\int_0^{2\pi} \theta'(t) dt = 2\pi$. This simple reparametrization achieves a similar effect as adaptively refined panel quadratures (which would be preferable in more general geometries). As evident in the table, a very small N may thus be used, while still matching all published digits of the lattice-sum based results in [35, Table 2], and matching to 10-12 digits values found with the more involved panel-based RCIP method.

Since the main point of Example 4 is to validate the periodizing scheme, we do not push to more extreme closeness (higher c). Nor do we use integral formulations [39, Eq. (9)] [37, Eq. (90)] better suited to large σ ; perhaps doing so would reduce our iteration count to near the 10 or 11 needed by Helsing’s RCIP method. CPU times are quoted for a laptop with an i7-7700HQ processor at 2.8 GHz, running the code `tbl_discarray_effcond.m` (see Remark 1.1). They are dominated by a rather crude $\mathcal{O}(N^3)$ MATLAB implementation of close-evaluation quadratures for filling A ; they could be much reduced.

4. The no-slip Stokes flow case. We move to our second BVP, that of viscous flow through a periodic lattice of inclusions with no-slip boundary conditions. We follow the normalization and some of the notation of [41, Sec. 2.2, 2.3]. Let the constant $\mu > 0$ be the fluid viscosity. The periodic BVP is to solve for a velocity field \mathbf{u} and pressure p function satisfying

$$-\mu \Delta \mathbf{u} + \nabla p = 0 \quad \text{in } \mathbb{R}^2 \setminus \overline{\Omega_\Lambda} \quad (4.1)$$

$$\nabla \cdot \mathbf{u} = 0 \quad \text{in } \mathbb{R}^2 \setminus \overline{\Omega_\Lambda} \quad (4.2)$$

$$\mathbf{u} = \mathbf{0} \quad \text{on } \partial\Omega_\Lambda \quad (4.3)$$

$$\mathbf{u}(\mathbf{x} + \mathbf{e}_1) = \mathbf{u}(\mathbf{x} + \mathbf{e}_2) = \mathbf{u}(\mathbf{x}) \quad \text{for all } \mathbf{x} \in \mathbb{R}^2 \setminus \overline{\Omega_\Lambda} \quad (4.4)$$

$$p(\mathbf{x} + \mathbf{e}_1) - p(\mathbf{x}) = p_1 \quad \text{for all } \mathbf{x} \in \mathbb{R}^2 \setminus \overline{\Omega_\Lambda} \quad (4.5)$$

$$p(\mathbf{x} + \mathbf{e}_2) - p(\mathbf{x}) = p_2 \quad \text{for all } \mathbf{x} \in \mathbb{R}^2 \setminus \overline{\Omega_\Lambda} . \quad (4.6)$$

The first two are Stokes' equations, expressing local force balance and incompressibility, respectively. The third is the no-slip condition, and the remainder express that the flow is periodic and the pressure periodic up to the given macroscopic pressure driving (p_1, p_2) .

We recall some basic definitions. Given a pair (\mathbf{u}, p) the Cauchy stress tensor field has entries

$$\sigma_{ij}(\mathbf{u}, p) := -\delta_{ij}p + \mu(\partial_i u_j + \partial_j u_i), \quad i, j = 1, 2. \quad (4.7)$$

The hydrodynamic traction \mathbf{T} (force vector per unit length that a boundary surface with outwards unit normal n applies to the fluid), also known as the Neumann data, has components

$$T_i(\mathbf{u}, p) := \sigma_{ij}(\mathbf{u}, p)n_j = -pn_i + \mu(\partial_i u_j + \partial_j u_i)n_j, \quad (4.8)$$

where here and below the Einstein convention of summation over repeated indices is used. We first show that the BVP has a one-dimensional nullspace.

PROPOSITION 4.1. *For each (p_1, p_2) the solution (\mathbf{u}, p) to (4.1)–(4.6) is unique up to an additive constant in p .*

Proof. The proof parallels that of Prop. 2.1. Given any function pairs (\mathbf{u}, p) and (\mathbf{v}, q) , Green's first identity on a domain \mathcal{K} is [49, p. 53]

$$\int_{\mathcal{K}} (\mu \Delta u_i - \partial_i p) v_i = -\frac{\mu}{2} \int_{\mathcal{K}} (\partial_i u_j + \partial_j u_i)(\partial_i v_j + \partial_j v_i) + \int_{\partial \mathcal{K}} T_i(\mathbf{u}, p) v_i. \quad (4.9)$$

Now let (\mathbf{u}, p) be the difference between two BVP solutions, and set $\mathbf{v} = \mathbf{u}$ and $\mathcal{K} = \mathcal{B} \setminus \overline{\Omega}$ in (4.9). The left-hand side vanishes due to (4.1), the $\partial \Omega$ boundary term vanishes due to (4.3), and the unit cell wall terms vanish by \mathbf{u} periodicity, leaving only $\int_{\mathcal{B} \setminus \overline{\Omega}} \sum_{i,j=1}^2 (\partial_i u_j + \partial_j u_i)^2 = 0$. Thus \mathbf{u} has zero stress, i.e. is a rigid motion, so, by (4.3), $\mathbf{u} \equiv \mathbf{0}$. Thus, by (4.1), and because p has no pressure drop, p is constant. \square

Since, for Stokes, the Cauchy data is (\mathbf{u}, \mathbf{T}) [41, Sec. 2.3], the BVP (4.1)–(4.6) is equivalent to the BVP on a single unit cell,

$$-\mu \Delta \mathbf{u} + \nabla p = 0 \quad \text{in } \mathcal{B} \setminus \overline{\Omega} \quad (4.10)$$

$$\nabla \cdot \mathbf{u} = 0 \quad \text{in } \mathcal{B} \setminus \overline{\Omega} \quad (4.11)$$

$$\mathbf{u} = \mathbf{0} \quad \text{on } \partial \Omega \quad (4.12)$$

$$\mathbf{u}_R - \mathbf{u}_L = \mathbf{0} \quad (4.13)$$

$$\mathbf{T}(\mathbf{u}, p)_R - \mathbf{T}(\mathbf{u}, p)_L = p_1 n \quad (4.14)$$

$$\mathbf{u}_U - \mathbf{u}_D = \mathbf{0} \quad (4.15)$$

$$\mathbf{T}(\mathbf{u}, p)_U - \mathbf{T}(\mathbf{u}, p)_D = p_2 n, \quad (4.16)$$

where the normal n has the direction and sense for the appropriate wall as in Fig. 1.1(a). Discrepancy will refer to the stack of the four vector functions on the left-hand side of (4.13)–(4.16).

4.1. The Stokes empty unit cell discrepancy BVP. Proceeding as with Laplace, one must first understand the empty unit cell BVP with given discrepancy data $\mathbf{g} := [\mathbf{g}_1; \mathbf{g}_2; \mathbf{g}_3; \mathbf{g}_4]$, which is to find a pair (\mathbf{v}, q) solving

$$-\mu \Delta \mathbf{v} + \nabla q = 0 \quad \text{in } \mathcal{B} \quad (4.17)$$

$$\nabla \cdot \mathbf{v} = 0 \quad \text{in } \mathcal{B} \quad (4.18)$$

$$\mathbf{v}_R - \mathbf{v}_L = \mathbf{g}_1 \quad (4.19)$$

$$\mathbf{T}(\mathbf{v}, q)_R - \mathbf{T}(\mathbf{v}, q)_L = \mathbf{g}_2 \quad (4.20)$$

$$\mathbf{v}_U - \mathbf{v}_D = \mathbf{g}_3 \quad (4.21)$$

$$\mathbf{T}(\mathbf{v}, q)_U - \mathbf{T}(\mathbf{v}, q)_D = \mathbf{g}_4, \quad (4.22)$$

This BVP has three consistency conditions and three nullspace dimensions, as follows.²

²Note that, although three is also the nullity of the 2D Stokes interior Neumann BVP [41, Table 2.3.3], both nullspace and consistency conditions differ from that case.

PROPOSITION 4.2. A solution (\mathbf{v}, q) to (4.17)–(4.22) exists if and only if

$$\int_L \mathbf{g}_2 ds + \int_D \mathbf{g}_4 ds = \mathbf{0} \quad (\text{zero net force}), \text{ and} \quad (4.23)$$

$$\int_L n \cdot \mathbf{g}_1 ds + \int_D n \cdot \mathbf{g}_3 ds = 0 \quad (\text{volume conservation}), \quad (4.24)$$

and then is unique up to translational flow and additive pressure constants. I.e. the solution space is $(\mathbf{v} + \mathbf{c}, q + c)$ for $(\mathbf{c}, c) \in \mathbb{R}^3$.

Proof. Noting that (4.9), with (\mathbf{u}, p) and (\mathbf{v}, q) swapped, holds for all constant flows \mathbf{u} shows that $\int_{\partial\mathcal{K}} \mathbf{T}(\mathbf{v}, q) = \mathbf{0}$; setting $\mathcal{K} = \mathcal{B}$ gives (4.23). (4.24) follows from the divergence theorem and (4.18). The proof of Prop. 4.1 shows that the nullspace is no larger than constant p and rigid motions for \mathbf{v} , but it is easy to check that rotation is excluded due to its effect on \mathbf{g}_1 and \mathbf{g}_3 . \square

We solve this Stokes empty BVP in an entirely analogous fashion to Laplace, namely via an MFS representation, but now with vector-valued coefficients. For $\mathbf{x} \in \mathcal{B}$,

$$\mathbf{v}(\mathbf{x}) \approx \sum_{j=1}^M \phi_j(\mathbf{x}) \boldsymbol{\xi}_j, \quad \phi_j(\mathbf{x}) := G(\mathbf{x}, \mathbf{y}_j), \quad \mathbf{y}_j := (R_p \cos 2\pi j/M, R_p \sin 2\pi j/M), \quad (4.25)$$

$$q(\mathbf{x}) \approx \sum_{j=1}^M \phi_j^p(\mathbf{x}) \cdot \boldsymbol{\xi}_j, \quad \phi_j^p(\mathbf{x}) := G^p(\mathbf{x}, \mathbf{y}_j), \quad (4.26)$$

where the velocity fundamental solution (stokeslet or single-layer kernel) is the tensor $G(\mathbf{x}, \mathbf{y})$ with components

$$G_{ij}(\mathbf{x}, \mathbf{y}) = \frac{1}{4\pi\mu} \left(\delta_{ij} \log \frac{1}{r} + \frac{r_i r_j}{r^2} \right), \quad i, j = 1, 2, \quad \mathbf{r} := \mathbf{x} - \mathbf{y}, \quad r := \|\mathbf{r}\|. \quad (4.27)$$

and the single-layer pressure kernel is the vector G^p with components

$$G_j^p(\mathbf{x}, \mathbf{y}) = \frac{1}{2\pi} \frac{r_j}{r^2}, \quad j = 1, 2. \quad (4.28)$$

We will also need the single-layer traction kernel G^t with components (applying (4.8) to the above),

$$G_{ik}^t(\mathbf{x}, \mathbf{y}) = \sigma_{ij}(G_{\cdot,k}(\cdot, \mathbf{y}), G_k^p(\cdot, \mathbf{y}))(\mathbf{x}) n_j^{\mathbf{x}} = -\frac{1}{\pi} \frac{r_i r_k}{r^2} \frac{\mathbf{r} \cdot \mathbf{n}^{\mathbf{x}}}{r^2}, \quad i, k = 1, 2, \quad (4.29)$$

where the target \mathbf{x} is assumed to be on a surface with normal $\mathbf{n}^{\mathbf{x}}$.

The generalization of the linear system from (scalar) Laplace to (vector) Stokes is routine bookkeeping, which we now outline. The MFS coefficient vector $\boldsymbol{\xi} := \{\boldsymbol{\xi}_j\}_{j=1}^M$ has $2M$ unknowns, which we order with the M 1-components followed by the M 2-components. Discretizing (4.19)–(4.22) on m collocation nodes per wall, as in Sec. 2.1, gives $Q\boldsymbol{\xi} = \mathbf{g}$ as in (2.19), with discrepancy vector $\mathbf{g} \in \mathbb{R}^{8m}$. We choose to order each of the four blocks in \mathbf{g} with all m 1-components followed by m 2-components. As before, $Q = [Q_1; Q_2; Q_3; Q_4]$ with each Q_k split into 2×2 sub-blocks based on the 1- and 2-components. For instance, writing ϕ_j^{kl} , $k, l = 1, 2$, for the four components of the basis function ϕ_j in (4.25), then the R - L velocity block $Q_1 = \begin{bmatrix} Q_1^{11} & Q_1^{21} \\ Q_1^{12} & Q_1^{22} \end{bmatrix}$, with each sub-block having entries $(Q_1^{kl})_{ij} = \phi_j^{kl}(\mathbf{x}_{iL} + \mathbf{e}_1) - \phi_j^{kl}(\mathbf{x}_{iL})$. The R - L traction block Q_2 has sub-blocks $(Q_2^{kl})_{ij} = G_{kl}^t(\mathbf{x}_{iL} + \mathbf{e}_1, \mathbf{y}_j) - G_{kl}^t(\mathbf{x}_{iL}, \mathbf{y}_j)$, where it is implied that the target normal is on L . The other two blocks are similar.

In Fig. 2.1(c) we show the convergence of the solution error produced by numerical solution of the Stokes version of (2.19) (again, see Remark 2.5). The results are almost identical to the Laplace case in Fig. 2.1(b), converging to 15 digit accuracy with a similar exponential rate; keep in mind that now there are $2M$ unknowns rather than M . Note that we do not know of a Stokes version of Theorem 2.3. Fig. 2.1(c) also shows that three singular values decay faster than the others, as expected from Prop. 4.1.

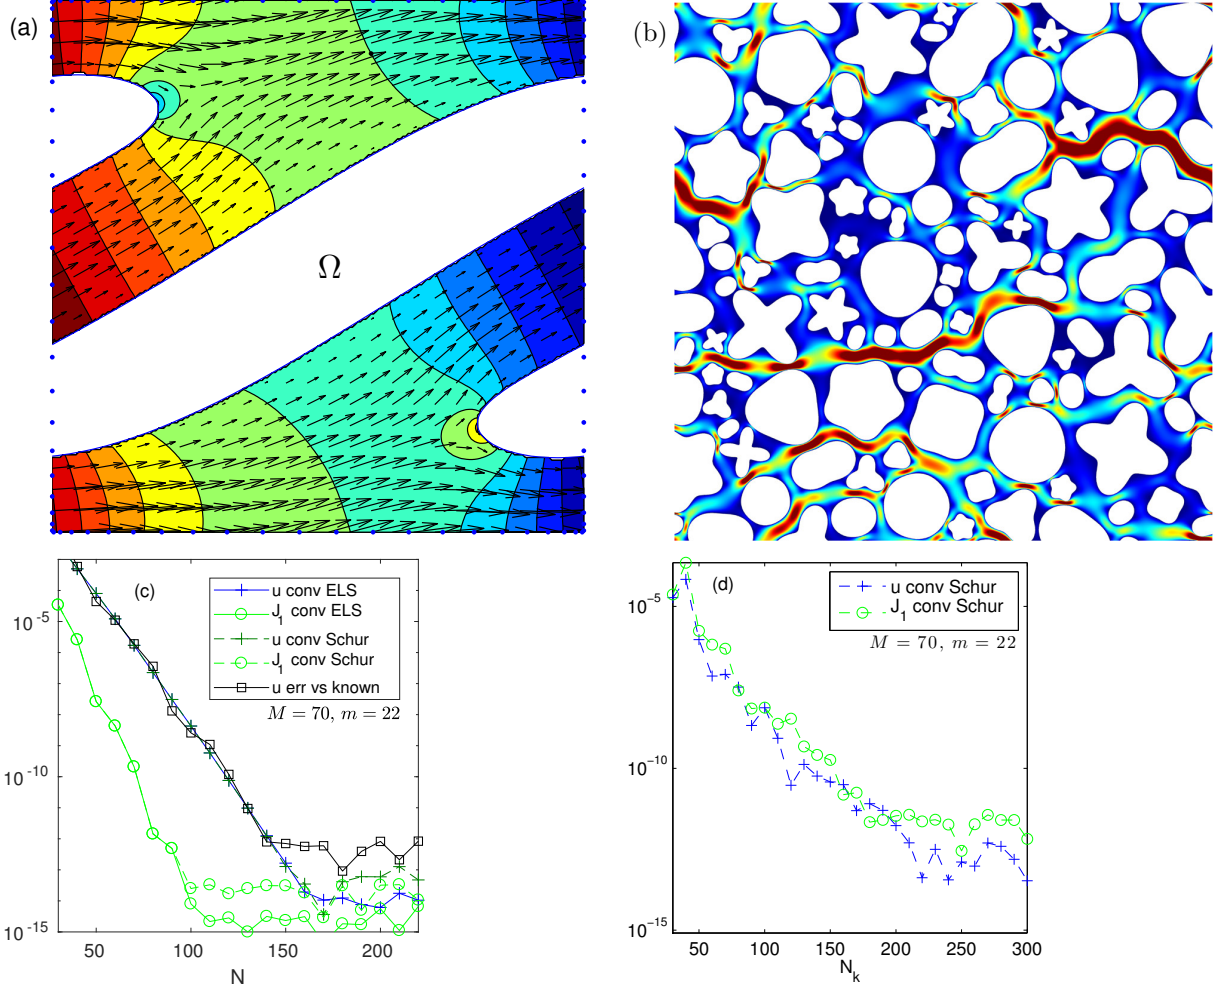


FIG. 4.1. Periodic Stokes flow tests with applied pressure drop $\mathbf{p} = (1, 0)$. (a) Solution velocity \mathbf{u} (arrows) and pressure p (contours) for Example 5 (same geometry as Fig. 2.2(a)), with $N = 150$, $M = 70$, and $R_p = 1.4$. A single unit cell is shown, and the $m = 22$ nodes per wall are shown as dots. (b) Flow speed $|\mathbf{u}|$ for $K = 100$ inclusions per unit cell, using iterative solution of (4.44) (Example 7; geometry identical to Fig. 2.2(b)). (c) N -convergence of \mathbf{u} error (with $M = 70$ fixed) at the same test points as in Fig. 2.2(a), and of flux J_1 (see end of Sec. 4.2), relative to their values at $N = 230$ (Example 5). Squares show error convergence in \mathbf{u} in the case of known \mathbf{u}_{ex} due to a Stokeslet grid. Solid lines are for direct solution of the ELS, dashed lines for the iterative solution of (4.44). (d) Convergence with N_k for $K = 100$ inclusions per unit cell (Example 7). Errors in velocity (+ signs) and flux J_1 (circles) are shown, estimated relative to their values at $N_k = 400$.

Finally, the discretization of the consistency conditions in Prop. 4.1 is the statement that

$$W^T Q \approx \mathbf{0}_{3 \times 2M}, \quad (4.30)$$

where $\mathbf{0}_{m \times n}$ is the m -by- n zero matrix, and the weight matrix W (the analog of (2.20)) is

$$W = \begin{bmatrix} 0 & 0 & \mathbf{w}_L^T & 0 & 0 & 0 & \mathbf{w}_D^T & 0 \\ 0 & 0 & 0 & \mathbf{w}_L^T & 0 & 0 & 0 & \mathbf{w}_L^T \\ \mathbf{w}_L^T & 0 & 0 & 0 & 0 & \mathbf{w}_D^T & 0 & 0 \end{bmatrix}^T \in \mathbb{R}^{8m \times 3}. \quad (4.31)$$

4.2. Integral representation and the Stokes ELS. In contrast to the Laplace problem, we are now solving a Dirichlet BVP, which suggests a pure double-layer representation on inclusion boundaries. However, this would lead to a 3D nullspace for each inclusion, associated with its complementary interior

Neumann BVP [41, Table 2.3.3]. We remove this nullspace via a “compound” or mixed double-layer formulation, namely an admixture of double- and single-layers [34, Thm. 2.1] [67, p.128]; this avoids the use of extra interior stokeslet degrees of freedom which can add $O(K^3)$ work for K inclusions [26].

For convenience we define the 2D Stokes layer potentials. Given a vector-valued density $\boldsymbol{\tau} = (\tau_1, \tau_2)$ on $\partial\Omega$, the Stokes single-layer potential generates (via (4.27) and (4.28)) velocity and pressure

$$\mathbf{u}(\mathbf{x}) = (\mathcal{S}_\Gamma \boldsymbol{\tau})(\mathbf{x}) := \int_\Gamma G(\mathbf{x}, \mathbf{y}) \boldsymbol{\tau}(\mathbf{y}) ds_{\mathbf{y}} , \quad p(\mathbf{x}) = (\mathcal{S}_\Gamma^p \boldsymbol{\tau})(\mathbf{x}) := \int_\Gamma G^p(\mathbf{x}, \mathbf{y}) \boldsymbol{\tau}(\mathbf{y}) ds_{\mathbf{y}} . \quad (4.32)$$

Using $\mathbf{n}^{\mathbf{y}}$ to denote the source normal, the Stokes double-layer velocity potential is

$$\mathbf{u}(\mathbf{x}) = (\mathcal{D}_\Gamma \boldsymbol{\tau})(\mathbf{x}) := \int_\Gamma D(\mathbf{x}, \mathbf{y}) \boldsymbol{\tau}(\mathbf{y}) ds_{\mathbf{y}} \quad \text{where } D_{ij}(\mathbf{x}, \mathbf{y}) = \frac{1}{\pi} \frac{r_i r_j}{r^2} \frac{\mathbf{r} \cdot \mathbf{n}^{\mathbf{y}}}{r^2} , \quad i, j = 1, 2 , \quad (4.33)$$

a kernel which is the negative transpose of (4.29), and the associated pressure

$$p(\mathbf{x}) = (\mathcal{D}_\Gamma^p \boldsymbol{\tau})(\mathbf{x}) := \int_\Gamma D^p(\mathbf{x}, \mathbf{y}) \boldsymbol{\tau}(\mathbf{y}) ds_{\mathbf{y}} \quad \text{where } D_j^p(\mathbf{x}, \mathbf{y}) = \frac{\mu}{\pi} \left(-\frac{n_j^{\mathbf{y}}}{r^2} + 2\mathbf{r} \cdot \mathbf{n}^{\mathbf{y}} \frac{r_j}{r^4} \right) , \quad j = 1, 2. \quad (4.34)$$

The Stokes analog of the boundary integral operator (2.24) is the single-layer traction

$$(D_{\Gamma', \Gamma}^T \boldsymbol{\tau})(\mathbf{x}) := \int_\Gamma G^t(\mathbf{x}, \mathbf{y}) \boldsymbol{\tau}(\mathbf{y}) ds_{\mathbf{y}} , \quad \mathbf{x} \in \Gamma' , \quad (4.35)$$

which uses (4.29). Finally, the hypersingular double-layer traction operator is needed,

$$(T_{\Gamma', \Gamma} \boldsymbol{\tau})(\mathbf{x}) := \int_\Gamma D^t(\mathbf{x}, \mathbf{y}) \boldsymbol{\tau}(\mathbf{y}) ds_{\mathbf{y}} , \quad \mathbf{x} \in \Gamma' , \quad (4.36)$$

whose kernel is computed by inserting (4.33) and (4.34) into (4.8) and simplifying to get (e.g. [55, (5.27)]),

$$D_{ik}^t = \frac{\mu}{\pi} \left[\left(\frac{n^{\mathbf{y}} \cdot n^{\mathbf{x}}}{r^2} - 8d_{\mathbf{x}} d_{\mathbf{y}} \right) \frac{r_i r_k}{r^2} + d_{\mathbf{x}} d_{\mathbf{y}} \delta_{ik} + \frac{n_i^{\mathbf{x}} n_k^{\mathbf{y}}}{r^2} + d_{\mathbf{x}} \frac{r_k n_i^{\mathbf{y}}}{r^2} + d_{\mathbf{y}} \frac{r_i n_k^{\mathbf{x}}}{r^2} \right] \quad (4.37)$$

where for conciseness we defined the target and source “dipole functions”, respectively

$$d_{\mathbf{x}} = d_{\mathbf{x}}(\mathbf{x}, \mathbf{y}) := (\mathbf{r} \cdot \mathbf{n}^{\mathbf{x}})/r^2 , \quad d_{\mathbf{y}} = d_{\mathbf{y}}(\mathbf{x}, \mathbf{y}) := (\mathbf{r} \cdot \mathbf{n}^{\mathbf{y}})/r^2 .$$

The Stokes jump relations [49, Sec. 3.2] are identical to the usual Laplace ones (2.23) with the potential taken to be velocity potential, and the normal derivative replaced by the traction. In short, the single-layer velocity and double-layer tractions are continuous, whereas the single-layer traction jump is minus the density, and the double-layer velocity jump is the density itself.

Armed with the above, our representation for the solution (\mathbf{u}, p) in $\mathcal{B} \setminus \overline{\Omega}$ is the mixed formulation,

$$\mathbf{u} = (\mathcal{D}_{\partial\Omega}^{\text{near}} + \mathcal{S}_{\partial\Omega}^{\text{near}}) \boldsymbol{\tau} + \sum_{j=1}^M \phi_j \boldsymbol{\xi}_j , \quad p = (\mathcal{D}_{\partial\Omega}^{\text{p, near}} + \mathcal{S}_{\partial\Omega}^{\text{p, near}}) \boldsymbol{\tau} + \sum_{j=1}^M \phi_j^p \cdot \boldsymbol{\xi}_j , \quad (4.38)$$

where “near” denotes 3×3 image sums as in (2.25). With this choice made, the continuous form of the ELS is similar to (2.26)–(2.30). The first block is, applying the exterior jump relation to (4.12),

$$(-\tfrac{1}{2} + D_{\partial\Omega, \partial\Omega}^{\text{near}} + S_{\partial\Omega, \partial\Omega}^{\text{near}}) \boldsymbol{\tau} + \sum_{j=1}^M \phi_j|_{\partial\Omega} \boldsymbol{\xi}_j = \mathbf{0} . \quad (4.39)$$

Rather than list all second subblocks, we note that they are simply (2.27)–(2.30) with S replaced by $D + S$, D^T replaced by $T + D^T$, and $\partial\phi_j(\mathbf{x})/\partial n$ replaced by the basis traction, which has kernel $G_{lk}^t(\mathbf{x}, \mathbf{y}_j)$. Finally, the pressure driving (p_1, p_2) is encoded by the discrete right hand side

$$\mathbf{g} = [0; 0; p_1 \mathbf{1}_m; 0; 0; 0; 0; p_2 \mathbf{1}_m] \in \mathbb{R}^{8m} , \quad (4.40)$$

so that the ELS has, as for Laplace, the form (1.2).

This mixed formulation allows an analog of Lemma 2.8: the Stokes ELS nullspace is spanned by ξ .

LEMMA 4.3. *In the solution to the Stokes ELS (i.e. (4.39) plus the Stokes version of (2.27)–(2.30) described above), τ is unique.*

Proof. The proof parallels that of Lemma 2.8, but adapts the 2D version of the proof of [34, Thm. 2.1] for the interior uniqueness step. Let \mathbf{v}^+ and \mathbf{v}^- denote the exterior and interior limits respectively of \mathbf{v} on $\partial\Omega$, and \mathbf{T}^+ and \mathbf{T}^- be the same limits of $\mathbf{T}(\mathbf{v}, q)$. If (τ, ξ) is the difference between any two ELS solutions, let (\mathbf{v}, q) be given by the representation (4.38) both in $B \setminus \bar{\Omega}$ and inside Ω . Thus, by Prop. 4.1, $\mathbf{v} \equiv \mathbf{0}$ and $q = c$, some constant, and $\mathbf{T}^+ = -cn$. The jump relations for $\mathcal{D} + \mathcal{S}$ then give $\mathbf{v}^- = \tau$ and $\mathbf{T}^- = -cn - \tau$. Thus in Ω , (\mathbf{v}, q) is a Stokes solution with Robin (impedance) data, $\mathbf{v}^- + \mathbf{T}^- = -cn$. Applying (4.9) with $\mathbf{u} = \mathbf{v}$ and $\mathcal{K} = \Omega$ gives $0 = -(\mu/2) \int_{\Omega} (\partial_i v_j + \partial_j v_i)^2 + c \int_{\partial\Omega} n \cdot \mathbf{v}^- - \int_{\partial\Omega} \|\mathbf{v}^-\|^2$. However by incompressibility $\int_{\partial\Omega} n \cdot \mathbf{v}^- = 0$, and the two remaining non-positive terms have the same sign, so must both vanish. Hence $\mathbf{v} \equiv \mathbf{0}$ in Ω . By the double-layer jump relation, $\tau = \mathbf{v}^+ - \mathbf{v}^-$, both of which have been shown to vanish, thus $\tau \equiv \mathbf{0}$. \square

The filling of the discrete $(2N + 8m)$ -by- $(2N + 2M)$ ELS is now routine, being as in Sec. 2.2, apart from vector bookkeeping, and the following two details. i) For the discretization of $D_{\partial\Omega, \partial\Omega}$ with kernel (4.33), the diagonal limit at node \mathbf{x}_k is $D_{ij}(\mathbf{x}_k, \mathbf{x}_k) = (-\kappa(\mathbf{x}_k)/2\pi) t_i(\mathbf{x}_k) t_j(\mathbf{x}_k)$ where (t_1, t_2) are the components of the unit tangent vector on $\partial\Omega$. ii) The kernel of $S_{\partial\Omega, \partial\Omega}$ has a logarithmic singularity on the diagonal, so the plain Nyström rule cannot be used. However, there exist several high-order discretizations for such a kernel [45, 2, 33]; when the N per inclusion is not large we prefer the spectral product quadrature due to Martensen, Kussmaul and Kress [45, Ch. 12.3] [33, Sec. 6.2].

We provide an implementation of the above kernel quadratures in `StoSLP.m` and `StoDLP.m`, and of filling the ELS in `fig_stoconvK1.m`; see Remark 1.1.

Example 5. We illustrate the above with results obtained via a simple dense direct solve of the ELS, for the periodic “worm” geometry from Example 1, with no-slip boundary conditions, pressure driving $\mathbf{p} = (1, 0)$, and $\mu = 0.7$. Fig. 4.1(a) shows the resulting flow \mathbf{u} and pressure p . Here we evaluated \mathbf{u} using the close-evaluation quadratures from [6], and p using similar formulae, to provide spectral accuracy even up to the curve. As expected from Prop. 4.1, the ELS exhibits unit numerical nullity: the Stokes version of Fig. 2.2(e) is very similar, and we do not show it. Fig. 4.1(c) shows the pointwise convergence in \mathbf{u} , and is consistent with exponential convergence down to 13 digit accuracy. For a known solution produced by stokeslets on the same grid as in Sec. 2.3, convergence stops at around 12 digits.

Rapid convergence of the flux J_1 is also shown in Fig. 4.1(c), reaching 13 digit accuracy at only $N = 100$. Here, flux is evaluated using the Stokes version of (2.43), which replaces the representation \mathcal{S} by $\mathcal{D} + \mathcal{S}$, replaces $\partial/\partial n$ by $n \cdot$, and in the proof replaces the zero-flux condition by incompressibility.

4.3. Schur complement systems for Stokes. Having filled the Stokes version of the ELS (1.2), we would like to eliminate the periodizing unknowns ξ_j to leave a $2N \times 2N$ system that can be solved iteratively.

4.3.1. An equivalent square system preserving the nullspace. It is possible to build such a system that preserves the rank-3 nullity using exactly the recipe in Section 2.4.1. For this we need only two ingredients: i) H must be a stack of three densities so that CH enlarges the range of \tilde{Q} to include all smooth vectors, and ii) R must be a stack of three proxy coefficient vectors with full rank projection onto $\text{Nul } Q$.

We first need the Stokes version of the technical lemma. For Stokes there are *two* types of Gauss’ law. Let $\mathbf{v} = \mathcal{S}_{\partial\Omega} \sigma + \mathcal{D}_{\partial\Omega} \tau$ be a velocity potential with arbitrary densities σ and τ on a curve $\partial\Omega$, with associated pressure $q = \mathcal{S}_{\partial\Omega}^p \sigma + \mathcal{D}_{\partial\Omega}^p \tau$. Then,

$$-\int_{\partial\mathcal{K}} \mathbf{T}(\mathbf{v}, q) = \int_{\partial\Omega} \sigma ds \quad (\text{net force}), \quad \text{and} \quad \int_{\partial\mathcal{K}} \mathbf{v} \cdot n = \int_{\partial\Omega} \tau \cdot n ds \quad (\text{net flux}), \quad (4.41)$$

where $\partial\mathcal{K}$ is the boundary of some open domain \mathcal{K} containing $\partial\Omega$. The proofs of both are similar to that of (2.37), summing over the boundaries of Ω and of $\mathcal{K} \setminus \bar{\Omega}$ to get the left side integral. For the first law one uses the single-layer traction jump relation, and, for the second, the double-layer velocity jump relation [49, p. 57–58].

REMARK 4.4. The second law of (4.41) implies that the net fluid flux through the boundary of \mathcal{K} is generally nonzero, i.e. that the Stokes double-layer velocity potential on the enclosed curve does not generally conserve fluid volume. This surprising fact we did not find stated in standard literature.

Now, for H we can choose densities that, when used with our $\mathcal{D} + \mathcal{S}$ representation, generate net force in the 1-direction, in the 2-direction, and break volume conservation, respectively:

$$H = \begin{bmatrix} \mathbf{h}^{(1)} & \mathbf{h}^{(2)} & \mathbf{h}^{(3)} \end{bmatrix} = \begin{bmatrix} \mathbf{1}_N & \mathbf{0}_N & \{n_1^{\mathbf{x}_i}\}_{i=1}^N \\ \mathbf{0}_N & \mathbf{1}_N & \{n_2^{\mathbf{x}_i}\}_{i=1}^N \end{bmatrix}, \quad (4.42)$$

where in the columns of the matrix we use “nodes fast, components slow” ordering, and recall that the normals live on the nodes of $\partial\Omega$. Next we show that CH has full-rank projection onto $\text{Ran } W$.

LEMMA 4.5. Given H in (4.42), W in (4.31), and C the Stokes discrepancy matrix filled as in Sec. 4.2, then the 3×3 matrix $W^T CH$ is diagonal, and, for all N and m sufficiently large, invertible.

Proof. The proof parallels that of Lemma 2.13. For instance, setting $\mathbf{v} = (\mathcal{S}_{\partial\Omega}^{\text{near}} + \mathcal{D}_{\partial\Omega}^{\text{near}})\mathbf{h}^{(1)}$ and using (4.41) we see the discrepancy has net force $|\partial\Omega|$ in the 1-direction only, and since $\int_{\partial\Omega} n = \mathbf{0}$ the discrepancy is volume-conserving. Discretizing shows that the first column of $W^T CH$ converges (in m and N) to $[-|\partial\Omega| \ 0 \ 0]^T$. Similar steps apply to the other two columns, showing that the matrix in fact converges, up to signs, to $|\partial\Omega|$ times the 3×3 identity. \square

For R we choose the three analogous scaled coefficient vectors on the proxy circle:

$$R = \begin{bmatrix} \mathbf{r}^{(1)} & \mathbf{r}^{(2)} & \mathbf{r}^{(3)} \end{bmatrix} = \frac{1}{M} \begin{bmatrix} \mathbf{1}_M & \mathbf{0}_M & \{\cos 2\pi j/M\}_{j=1}^M \\ \mathbf{0}_M & \mathbf{1}_M & \{\sin 2\pi j/M\}_{j=1}^M \end{bmatrix}. \quad (4.43)$$

Reasoning as in the start of the proof of Theorem 2.12, the three (linearly independent) columns of R are exponentially close to generating constant flow in the 1-direction, the 2-direction, and constant pressure, respectively, which, according to Prop. 4.2, fall into $\text{Nul } Q$.

Given the above statements about H and R , Theorem 2.12 holds for the Stokes version of the procedure of Section 2.4.1; in its proof one need only change \mathbf{w} to the W of (4.31). Because of the unit nullity of the Stokes ELS, this would result in a rank-1 deficient square system matrix.

4.3.2. A well-conditioned square Stokes system. To remove this nullspace, it is tempting to take the obvious Stokes generalization of Section 2.4.2, leaving the B block unperturbed as in (2.39). However, this fails (the solution having constant, but incorrect, pressure drops) because $W^T \mathbf{g} = [p_1; p_2; 0] \neq \mathbf{0}_3$. However, since only one nullspace dimension needs to be removed, and this nullspace is common to the empty unit cell BVP, this suggests a hybrid where the perturbation of B involves *only the remaining two* dimensions. This indeed leads to a well-conditioned square system:

1. Set H as in (4.42) and R as in (4.43), then set $\tilde{Q} := Q + CHR^T$ and $\hat{B} := B + AP_{12}$, where $P_{12} := [\mathbf{h}^{(1)} \ \mathbf{h}^{(2)}][\mathbf{r}^{(1)} \ \mathbf{r}^{(2)}]^T$ involves the two force (but not the one pressure) perturbations.
2. Solve for the vector $\hat{\tau}$ in the $2N \times 2N$ Schur complement linear system

$$\hat{A}_{\text{per}} \hat{\tau} := (A - \hat{B}\tilde{Q}^+ C) \hat{\tau} = -\hat{B}\tilde{Q}^+ \mathbf{g}, \quad (4.44)$$

where, as before, one solves the small systems $\tilde{Q}X = C$ and $\tilde{Q}\mathbf{y} = \mathbf{g}$, to build the large system matrix $\hat{A}_{\text{per}} = A - \hat{B}X$ and right-hand side $-\hat{B}\mathbf{y}$ for (4.44), which may then be solved iteratively.

3. Recover the proxy coefficients via $\xi = \mathbf{y} - X\hat{\tau}$.
4. Recover the density via $\tau = \hat{\tau} + P_{12}\xi$.

THEOREM 4.6. Let \mathbf{g} encode the pressure driving as in (4.40). Then for all N , M and m sufficiently large, the pair (τ, ξ) produced by the above procedure is unique, and solves the Stokes version of the ELS (1.2) with residual of order the quadrature error on boundaries.

This is a Stokes analog of Theorem 2.14, and its proof structure is similar.

Proof. (4.44) is the Schur complement of the perturbed ELS

$$\begin{bmatrix} A & B + AP_{12} \\ C & Q + CHR^T \end{bmatrix} \begin{bmatrix} \hat{\tau} \\ \xi \end{bmatrix} = \begin{bmatrix} 0 \\ \mathbf{g} \end{bmatrix}, \quad (4.45)$$

c (vol. frac.)	N	# iters.	CPU time	est. rel. err.	D_{calc}	D_{GK}
0.05	60	5	0.02 s	8×10^{-15}	15.5578223189024	15.5578
0.1	60	7	0.03 s	3×10^{-14}	24.831675248849	24.8317
0.2	60	8	0.02 s	3×10^{-14}	51.526948961799	51.5269
0.3	100	10	0.02 s	1×10^{-13}	102.88130367005	102.881
0.4	150	10	0.03 s	1×10^{-13}	217.89431472214	217.894
0.5	150	12	0.03 s	9×10^{-13}	$5.325481184633 \times 10^2$	5.32548×10^2
0.6	300	13	0.08 s	1×10^{-12}	$1.763573112520 \times 10^3$	1.76357×10^3
0.7	400	20	0.14 s	7×10^{-11}	$1.35191501298 \times 10^4$	1.35193×10^4
0.75	500	28	0.22 s	2×10^{-10}	1.2753159105×10^5	1.27543×10^5
0.76	600	32	0.31 s	1×10^{-9}	2.948165200×10^5	2.94878×10^5
0.77	700	44	0.42 s	7×10^{-9}	1.03826971×10^6	1.03903×10^6
0.78	1300	64	1.5 s	6×10^{-8}	1.4379206×10^7	—

TABLE 4.1

Computation of the 2D viscous Stokes drag of an infinite square array of no-slip discs of various volume fractions c (Example 6). The well-conditioned Schur system of (4.44) was used. D_{calc} shows the dimensionless drag results. Relative errors were estimated by convergence (comparing each against three N values larger than shown, and increasing M). The last column D_{GK} shows for comparison the numerical findings of [26, Table 1, 5th column].

so that if $\hat{\tau}$ solves (4.44) then $(\hat{\tau}, \xi)$, with ξ as in Step 3, solves (4.45). Applying Step 4, one may check that the pair (τ, ξ) then solves a rank-1 perturbed system similar to (2.39),

$$\begin{bmatrix} A & B \\ C & Q + C\mathbf{h}^{(3)}\mathbf{r}^{(3)\top} \end{bmatrix} \begin{bmatrix} \tau \\ \xi \end{bmatrix} = \begin{bmatrix} 0 \\ \mathbf{g} \end{bmatrix}. \quad (4.46)$$

From this pair, define \mathbf{v} as the resulting quadrature approximation to the velocity potential (4.38), namely

$$\mathbf{v}(\mathbf{x}) := \sum_{m_1, m_2 \in \{-1, 0, 1\}} \sum_{i=1}^N w_i [D(\mathbf{x}, \mathbf{x}_i + m_1 \mathbf{e}_1 + m_2 \mathbf{e}_2) + G(\mathbf{x}, \mathbf{x}_i + m_1 \mathbf{e}_1 + m_2 \mathbf{e}_2)] \tau_i + \sum_{j=1}^M \phi_j(\mathbf{x}) \xi_j.$$

Since it matches that of the ELS, the first row of (4.46) implies that $\mathbf{v}|_{\partial\Omega} = \mathbf{0}$. Thus $\int_{\partial\Omega} \mathbf{v} \cdot \mathbf{n} = 0$ and, by volume conservation in $\mathcal{B} \setminus \overline{\Omega}$, the discrepancy of (\mathbf{v}, q) obeys (4.24), whose discretization is $\mathbf{w}^{(3)\top}(C\tau + Q\xi) = 0$, where $\mathbf{w}^{(3)}$ is the third column of (4.31). Subtracting this from the 2nd row of (4.46) left-multiplied by $\mathbf{w}^{(3)\top}$ leaves $(\mathbf{w}^{(3)\top} C\mathbf{h}^{(3)})\mathbf{r}^{(3)\top} \xi = \mathbf{w}^{(3)\top} \mathbf{g}$. Lemma 4.5 implies that $\mathbf{w}^{(3)\top} C\mathbf{h}^{(3)} \neq 0$, and by computation $\mathbf{w}^{(3)\top} \mathbf{g} = 0$. Thus $\mathbf{r}^{(3)\top} \xi = 0$, so the pair also solves the ELS. Thus Lemma 4.3 holds, so τ is unique. \square

We expect (4.44) to be well-conditioned if A is, since the unit nullity has been removed by enforcing one extra condition $\mathbf{r}^{(3)\top} \xi = 0$. In contrast to Laplace, here the rank-3 nullity of the empty BVP demands a hybrid scheme which “segregates” the other two conditions.

We test this scheme returning to Example 5 (Fig. 4.1). Panel (c) includes (as dotted lines) the convergence using the above recipe, using GMRES with tolerance 10^{-14} . The convergence is almost identical to that of the direct ELS solution, and the converged values differ by around 2×10^{-14} . The condition number of \hat{A}_{per} is 82.7, independent of N and M , and the number of GMRES iterations varied between 25 and 30.

4.4. More challenging numerical tests for Stokes. We now validate the above periodic Stokes scheme against a known test case, and in geometries with large numbers of inclusions.

Example 6. The effective permeability of the no-slip flow around an infinite square lattice of discs is a standard test case. Only one pressure drop $\mathbf{p} = (1, 0)$ need be solved, since the permeability tensor is a multiple of the identity. The dimensionless drag is then related to the permeability κ_{11} by $D = 1/(\mu\kappa_{11})$. We computed this for a list of volume fractions $c = \pi r^2$ including those of Greengard–Kropinski [26] (note that to match the literature, the definition of c has changed from Example 4). We used the graded circle parameterization (3.5) with N chosen by convergence studies. As c approaches $\pi/4$, the discs come

closer, the solution density becomes more peaked, and the required N grows. At the most extreme case of $c = 0.78$, the gap is 3×10^{-3} (i.e. $f_{\text{clap}} \approx 1800$). Since the gap is much larger than in Example 4, we were able to use the plain far-field Nyström rule for filling A . (We found that, in this setting, close-evaluation quadratures [6] were slower, and around one digit less accurate.) At large c the N values needed are still smaller than in [26, Table 1]. We used GMRES with $\epsilon = 10^{-14}$ and dense matrix-vector multiplication, and find similar numbers of iterations as in [26]. CPU times are quoted for a laptop with an i7-7700HQ processor at 2.8 GHz, running the code `tbl_discarray_drag.m` (see Remark 1.1).

Our results are given in Table 4.1. For validation we include in the last column the numerical findings from [26, Table 1], which were quoted to 6 digits, and in turn validated against prior less accurate results [71]. For $c \leq 0.6$, our results match the 6 digits quoted in [26, Table 1], but above this the match deteriorates, reaching only 3 digits at their most challenging case $c = 0.77$ (we believe that their non-matching digits are incorrect). For N values larger than shown, our results fluctuate only at the estimated error levels shown; thus we are confident that our scheme reaches around 14 digits of accuracy for the smaller c values, dropping to 7 digits at $c = 0.78$. This loss of accuracy cannot be explained by condition number alone, since the latter reaches only 10^4 at $c = 0.78$. Further work is needed to ascertain whether this loss is controlled by an underlying condition number of the BVP (and hence impossible to improve upon in double-precision arithmetic), or if a different representation or quadrature, analogous to [37], could reduce the error.

Example 7. We solve the geometry of Example 2, with $K = 100$ inclusions, but with Stokes no-slip conditions, and the well-conditioned iterative scheme of Sec. 4.3.2. We apply Stokes potentials using a few calls to the Laplace FMM, as presented in [6, Sec. 2.1]. Close-evaluation quadratures for S and D are used, as in Example 5. One period of the resulting solution flow speed is plotted in Fig. 4.1(b). Convergence is shown by Fig. 4.1(d); 11–12 digits of accuracy in flux is achieved, which is around 2 digits less than for Laplace. At the converged value $N_k = 200$, the solution takes 1720 seconds.

Linear complexity as a function of K is demonstrated by Fig. 2.3(c). The time per iteration is similar to that for Laplace; this is because we pre-store the close-evaluation matrix blocks whereas for Laplace we evaluated them on the fly. We show the solution for $K = 10^3$ in Fig. 1.3, using $N_k = 350$. This required 1289 GMRES iterations and took 21 hours using the same machine as in Fig. 1.2. By comparison to the solution at $N_k = 450$, we estimate that the absolute flux error is 9×10^{-9} . However, due to the narrow channels in this porous medium, the size of the flux is only around 10^{-5} , meaning that *relative* accuracy is only 3 digits. Further study is needed to determine if it is an underlying condition number of the flux problem that prevents more digits from being achieved. The code is found in the directory `multiinclusion_Stokes` (see Remark 1.1).

Example 8. In the previous example large numbers of GMRES iterations were needed, in contrast to the Laplace case (cf. Example 2). This motivates studying the growth of iteration count with geometric complexity. We use exactly the same circular geometries as in Example 3. Fig. 2.3(d) shows that the numbers of iterations are around ten times those for Laplace, and that they grow with K in a more aggressive manner possibly consistent with $K^{1/2}$.

5. Conclusions. We have presented a unified framework for the 2nd-kind integral equation solution of large-scale periodic Laplace and Stokes BVPs, that we believe will be useful in related settings such as multiphase composites, photonic crystals, and particulate flows. The philosophy is quite simple, and we encourage readers to try implementing it in their applications:

- Use a free-space potential theory representation for the near neighbor images, plus $\mathcal{O}(1)$ auxiliary “proxy” degrees of freedom to represent the rest of the infinite lattice.
- Augment the linear system by applying *physical* boundary conditions on a unit cell.
- Use a low-rank perturbation of the empty BVP system to remove the consistency condition, allowing elimination of the auxiliary unknowns to give a well-conditioned system compatible with fast algorithms.

This cures two ills in one fell swoop: it removes zero eigenvalue(s) that arise in the physical BVP, and circumvents the issue of non-existence of the periodic Green’s function for general densities.

In Section 2.4.1 we provided a simple unified procedure for periodizing Laplace and Stokes integral representations when the physical BVP nullspace is numerically tolerable—we believe that this would be a useful starting point for related elliptic BVPs. However, when a well-conditioned system is desired, we

have presented and analyzed procedures for Laplace (Section 2.4.2) and Stokes (Section 4.3.2) systems, where the latter exploits the fact that the physical BVP consistency conditions form a *subset* of those of the empty BVP. The relevant theorems (2.12, 2.14 and 4.6) generalize simply to the case of any consistent right hand side. We showcased the scheme with high-accuracy computations of the effective conductivity, and the viscous drag, of a square lattice of discs and of random composites with a large number of inclusions needing several million unknowns.

REMARK 5.1 (Randomization). *In this work we have been explicit about choosing low-rank matrices R and H proven to have the required full-rank projections, in both Laplace and Stokes cases (Lemmata 2.9, 2.13 and 4.5). For other BVPs, this may not be easy or convenient. However, in practice, there is much flexibility in their choice. Following [73], one may even pick random matrices for both R and H , exploiting the fact that with probability one they have the required full-rank projections. In our experiments, the only penalty is a small loss of GMRES convergence rate, and a larger, fluctuating, condition number.*

We finish with directions for future work. We believe that an elastostatic version [36, 12] would be very similar to what we present for Stokes. While we restricted our attention to static geometry BVPs, the scheme can handle moving geometry problems, such as the flow of bubbles, vesicles or bacteria, in a straightforward manner (e.g., see [58]). The scheme generalizes easily to 3D (requiring $M \sim 10^3$ auxiliary unknowns) [57]. Note that in (2.25) one need not sum all 3×3 copies of source quadrature nodes when applying A , but just the ones falling inside the proxy circle (or sphere) [30]. Exploiting this can lower the overall constant, especially in 3D, but requires extra bookkeeping in the C matrix block.

By further augmenting the linear system to include decay/radiation conditions, the scheme has already proven useful for cases when the *periodicity is less than the space dimension*, such as singly-periodic in 2D [22, 58, 13] or doubly-periodic in 3D [56]. The latter case has applications in electrostatics [54] and Stokes flow [52, 16]. The case of singly-periodic in 3D, and of non-periodic solutions in periodic geometries, await development. The generalization to arbitrary skew and/or high aspect unit cells has applications in shearing suspensions and in microfluidics. This is easy in principle, with the proxy circle becoming an oval in order to “shield” the far images, as in Fig. 1.1(c). However, since M must grow with the aspect ratio, this will require new ideas beyond an aspect ratio of, say, 10^3 in 2D.

More theoretical work is required, since our analysis does not attempt precise bounds on how quadrature errors propagate to solution errors for the scheme. Work is also needed to understand if observed limitations in Stokes relative flux accuracy (such as 7 digits at the end of Table 4.1, and 3 relative digits in Fig. 1.3) are *intrinsic* (unavoidable in double precision due to ill-conditioning of the BVP), or are instead due to lack of optimality in the scheme. Finally, to reduce the large GMRES iteration counts for Stokes flows in complicated close-to-touching geometries, adaptive schemes such as that of Helsing [37], or the use of fast direct solvers or local preconditioners, will be important.

6. Acknowledgments. We are grateful for useful discussions with Bob Kohn, Leslie Greengard, Jun Lai, Mark Tygert, Manas Rachh, and the anonymous referees. The work of AB and LZ was supported in part by NSF grant DMS-1216656. The work of LZ was largely done while at Dartmouth College. The work of GM and SV was supported in part by NSF grants DMS-1418964 and DMS-1454010. The Flatiron Institute is a division of the Simons Foundation.

REFERENCES

- [1] L. af Klinteberg and A.-K. Tornberg. Fast Ewald summation for Stokesian particle suspensions. *Int. J. Num. Meth. Fluids*, 76:669–698, 2014.
- [2] B. K. Alpert. Hybrid Gauss-trapezoidal quadrature rules. *SIAM J. Sci. Comput.*, 20:1551–1584, 1999.
- [3] 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.
- [4] A. H. Barnett and T. Betcke. Stability and convergence of the Method of Fundamental Solutions for Helmholtz problems on analytic domains. *J. Comput. Phys.*, 227(14):7003–7026, 2008.
- [5] A. H. Barnett and L. Greengard. A new integral representation for quasi-periodic fields and its application to two-dimensional band structure calculations. *J. Comput. Phys.*, 229:6898–6914, 2010.
- [6] A. H. 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.
- [7] C. L. Berman and L. Greengard. A renormalization method for the evaluation of lattice sums. *J. Math. Phys.*, 35:6036–6048, 1994.

- [8] A. Bogomolny. Fundamental solutions method for elliptic boundary value problems. *SIAM J. Numer. Anal.*, 22(4):644–669, 1985.
- [9] J. M. Borwein, L. Glasser, R. McPhedran, J. G. Wan, and I. J. Zucker. *Lattice Sums: Then and Now*. Cambridge University Press, Cambridge, 2013.
- [10] P. N. Brown and H. F. Walker. GMRES on (nearly) singular systems. *SIAM J. Matrix Anal. Appl.*, 18:37–51, 1997.
- [11] O. P. Bruno and L. A. Kunyansky. Surface scattering in three dimensions: an accelerated high-order solver. *Proc. R. Soc. Lond. A*, 457:2921–2934, 2001.
- [12] P. Cazeaux and O. Zahm. A fast boundary element method for the solution of periodic many-inclusion problems via hierarchical matrix techniques. *ESAIM Proceedings and Surveys*, 48:156–168, 2015.
- [13] M. H. Cho and A. H. Barnett. Robust fast direct integral equation solver for quasi-periodic scattering problems with a large number of layers. *Opt. Express*, 23(2):1775–1799, 2015.
- [14] D. Cioranescu and P. Donato. *An introduction to homogenization*. Oxford University Press, 1999.
- [15] 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):7678–7693, 1998.
- [16] Y. Ding, J. C. Nawroth, M. J. McFall-Ngai, and E. Kanso. Mixing and transport by ciliary carpets: a numerical study. *J. Fluid Mech.*, 743:124–140, 2014.
- [17] A. Doicu, Y. A. Eremin, and T. Wriedt. *Acoustic and Electromagnetic Scattering Analysis Using Discrete Sources*. Academic Press, San Diego, CA, 2000.
- [18] J. E. Drummond and M. I. Tahir. Laminar viscous flow through regular arrays of parallel solid cylinders. *Int. J. Multiphase Flow*, 10:515–540, 1984.
- [19] A. Dutt and V. Rokhlin. Fast Fourier transforms for nonequispaced data. *SIAM J. Sci. Comput.*, 14:1369–1393, 1993.
- [20] P. P. Ewald. Die berechnung optischer und elektrostatischer gitterpotentiale. *Ann. Phys.*, 64:253–287, 1921.
- [21] G. Fairweather and A. Karageorghis. The method of fundamental solutions for elliptic boundary value problems. *Adv. Comput. Math.*, 9(1-2):69–95, 1998.
- [22] A. Gillman and A. Barnett. A fast direct solver for quasiperiodic scattering problems. *J. Comput. Phys.*, 248:309–322, 2013.
- [23] A. Gillman, P. Young, and P. Martinsson. A direct solver with $O(N)$ complexity for integral equations on one-dimensional domains. *Frontiers of Mathematics in China*, 7(2):217–247, 2012.
- [24] Z. Gimbutas and L. Greengard. FMMLIB2D, Fortran libraries for fast multipole methods in two dimensions. <https://github.com/zgimbutas/fmmlib2d>.
- [25] G. H. Golub and C. F. Van Loan. *Matrix computations*. Johns Hopkins Studies in the Mathematical Sciences. Johns Hopkins University Press, Baltimore, MD, third edition, 1996.
- [26] L. Greengard and M. C. Kropinski. Integral equation methods for stokes flow in doubly-periodic domains. *J. Eng. Math.*, 48:157–170, 2004.
- [27] L. Greengard and M. Moura. On the numerical evaluation of electrostatic fields in composite materials. *Acta Numerica*, 3:379–410, 1994.
- [28] L. Greengard and V. Rokhlin. A fast algorithm for particle simulations. *J. Comput. Phys.*, 73:325–348, 1987.
- [29] J. M. Guedes and N. Kikuchi. Preprocessing and postprocessing for materials based on the homogenization method with adaptive finite element methods. *Comput. Methods in Appl. Mech. Eng.*, 83:143–198, 1990.
- [30] N. A. Gumerov and R. Duraiswami. A method to compute periodic sums. *J. Comput. Phys.*, 272:307–326, 2014.
- [31] W. Hackbusch. A sparse matrix arithmetic based on H-matrices; Part I: Introduction to H-matrices. *Computing*, 62:89–108, 1999.
- [32] C. Hafner. *The Generalized Multipole Technique for Computational Electromagnetics*. Artech House Books, Boston, 1990.
- [33] S. Hao, A. H. Barnett, P. G. Martinsson, and P. Young. High-order accurate Nyström discretization of integral equations with weakly singular kernels on smooth curves in the plane. *Adv. Comput. Math.*, 40(1):245–272, 2014.
- [34] F.-K. Hebeker. Efficient boundary element methods for three-dimensional exterior viscous flows. *Numer. Methods Partial Differential Equations*, 2:273–297, 1986.
- [35] J. Helsing. Fast and accurate calculations of structural parameters for suspensions. *Proc. R. Soc. Lond. A*, 445:127–140, 1994.
- [36] J. Helsing. An integral equation method for elastostatics of periodic composites. *J. Mech. Phys. Solids*, 43(6):815–828, 1995.
- [37] J. Helsing. Solving integral equations on piecewise smooth boundaries using the RCIP method: a tutorial, 2017. preprint, 51 pages, [arXiv:1207.6737v7](https://arxiv.org/abs/1207.6737).
- [38] J. Helsing and L. Greengard. A numerical study of the ζ_2 parameter for random suspensions of disks. *J. Appl. Phys.*, 77(5):2015–2019, 1995.
- [39] J. Helsing and R. Ojala. Corner singularities for elliptic problems: Integral equations, graded meshes, quadrature, and compressed inverse preconditioning. *J. Comput. Phys.*, 227:8820–8840, 2008.
- [40] J. Helsing and R. Ojala. On the evaluation of layer potentials close to their sources. *J. Comput. Phys.*, 227:2899–2921, 2008.
- [41] G. Hsiao and W. L. Wendland. *Boundary Integral Equations*. Applied Mathematical Sciences, Vol. 164. Springer, 2008.
- [42] J. D. Joannopoulos, S. G. Johnson, R. D. Meade, and J. N. Winn. *Photonic Crystals: Molding the Flow of Light*. Princeton Univ. Press, Princeton, NJ, 2nd edition, 2008.
- [43] M. Katsurada. A mathematical study of the charge simulation method. II. *J. Fac. Sci. Univ. Tokyo Sect. IA Math.*, 36(1):135–162, 1989.
- [44] M. Katsurada and H. Okamoto. A mathematical study of the charge simulation method. I. *J. Fac. Sci. Univ. Tokyo Sect. IA Math.*, 35(3):507–518, 1988.

- [45] R. Kress. *Linear Integral Equations*, volume 82 of *Appl. Math. Sci.* Springer, second edition, 1999.
- [46] V. D. Kupradze. On the approximate solution of problems in mathematical physics. *Uspekhi Mat. Nauk*, 22(2(134)):59–107, 1967.
- [47] A. G. Kyurkchan, B. Y. Sternin, and V. E. Shatalov. Singularities of continuation of wave fields. *Physics - Uspekhi*, 12:1221–1242, 1996.
- [48] A. J. C. Ladd. Numerical simulations of particulate suspensions via a discretized Boltzmann equation part ii. Numerical results. *J. Fluid Mech.*, 271:311–339, 1994.
- [49] O. A. Ladyzhenskaya. *The Mathematical Theory of Viscous Incompressible Flow, revised 2nd edition*. Mathematics and Its Applications 2. Gordon and Breach, 1969.
- [50] R. E. Larson and J. J. L. Higdon. Microscopic flow near the surface of two-dimensional porous media. part 2. transverse flow. *J. Fluid Mech.*, 178:119–136, 1987.
- [51] D. Lindbo and A.-K. Tornberg. Spectrally accurate fast summation for periodic Stokes potentials. *J. Comput. Phys.*, 229:8994–9010, 2010.
- [52] D. Lindbo and A.-K. Tornberg. Fast and spectrally accurate summation of 2-periodic Stokes potentials, 2011. [arXiv:1111.1815](https://arxiv.org/abs/1111.1815).
- [53] D. Lindbo and A.-K. Tornberg. Spectral accuracy in fast Ewald-based methods for particle simulations. *J. Comput. Phys.*, 230:8744–8761, 2011.
- [54] D. Lindbo and A.-K. Tornberg. Fast and spectrally accurate Ewald summation for 2-periodic electrostatic systems. *J. Chem. Phys.*, 136(16):164111, 2012.
- [55] Y. Liu. *Fast Multipole Boundary Element Method: Theory and Applications in Engineering*. Cambridge University Press, 2009.
- [56] Y. Liu and A. H. Barnett. Efficient numerical solution of acoustic scattering from doubly-periodic arrays of axisymmetric objects. *J. Comput. Phys.*, 324:226–245, 2016.
- [57] G. Marple. Fast, high-order algorithms for simulating vesicle flows through periodic geometries, 2016. Ph.D thesis, (University of Michigan).
- [58] G. R. Marple, A. H. Barnett, A. Gillman, and S. Veerapaneni. A fast algorithm for simulating multiphase flows through periodic geometries of arbitrary shape, 2016. in press, SIAM J. Sci. Comput.
- [59] P. G. Martinsson and V. Rokhlin. A fast direct solver for scattering problems involving elongated structures. *J. Comput. Phys.*, 221:288–302, 2007.
- [60] R. C. McPhedran, L. Poladian, and G. W. Milton. Asymptotic studies of closely spaces, highly conducting cylinders. *Proc. R. Soc. Lond. A*, 415:185–196, 1988.
- [61] C. Miehe and A. Koch. Computational micro-to-macro transitions of discretized microstructures undergoing small strains. *Arch. Appl. Mech.*, 72:300–317, 2002.
- [62] G. W. Milton. *The theory of composites*. Cambridge University Press, 2004.
- [63] C. B. Moler. *Numerical computing with MATLAB*. SIAM, revised reprint edition, 2008.
- [64] Y. Otani and N. Nishimura. A fast multipole boundary integral equation method for periodic boundary value problems in three-dimensional elastostatics and its application to homogenization. *Int. J. Multiscale Comput. Eng.*, 4(4):487–500, 2006.
- [65] Y. Otani and N. Nishimura. A periodic FMM for Maxwell’s equations in 3D and its applications to problems related to photonic crystals. *J. Comput. Phys.*, 227:4630–52, 2008.
- [66] G. A. Pavliotis and A. Stuart. *Multiscale Methods. Averaging and Homogenization*. Springer, 2008.
- [67] C. Pozrikidis. *Boundary Integral and Singularity Methods for Linearized Viscous Flow*. Cambridge Texts in Applied Mathematics. Cambridge University Press, 1992.
- [68] C. Pozrikidis. Computation of periodic Green’s functions of Stokes flow. *J. Eng. Math.*, 30:79–96, 1995.
- [69] M. Rachh and L. Greengard. Integral equation methods for elastance and mobility problems in two dimensions, 2015. [arXiv:1507.05925](https://arxiv.org/abs/1507.05925).
- [70] J. W. Rayleigh. On the influence of obstacles arranged in rectangular order on the properties of the medium. *Philos. Mag.*, 43:481–502, 1892.
- [71] A. S. Sangani and A. Acrivos. Slow flow past periodic arrays of cylinders with application to heat transfer. *Int. J. Multiphase Flow*, 8(3):193–206, 1982.
- [72] A. S. Sangani and G. Mo. Inclusion of lubrication forces in dynamic simulations. *Phys. Fluids*, 6(5):1653–1662, 1994.
- [73] J. Sifuentes, Z. Gimbutas, and L. Greengard. Randomized methods for rank-deficient linear systems. *Electron. Trans. Numer. Anal.*, 44:177–188, 2015.
- [74] G. A. L. van de Vorst. Integral formulation to simulate the viscous sintering of a two-dimensional lattice of periodic unit cells. *J. Eng. Math.*, 30:97–118, 1996.
- [75] C. Y. Wang. Stokes flow through an array of rectangular fibers. *Int. J. Multiphase Flow*, 22(1):185–194, 1996.
- [76] W. Yan and M. Shelley. Flexibly imposing periodicity in kernel independent FMM: A multipole-to-local operator approach, 2017. preprint, [http://arxiv.org/abs/1705.02043](https://arxiv.org/abs/1705.02043).
- [77] Y. Yan and I. H. Sloan. On integral equations of the first kind with logarithmic kernels. *J. Integral Eqns. Appl.*, 1(4):549–579, 1988.
- [78] J. Yuan, Y. Y. Lu, and X. Antoine. Modeling photonic crystals by boundary integral equations and Dirichlet-to-Neumann maps. *J. Comput. Phys.*, 227:4617–4629, 2008.