

FINITE ELEMENTS FOR HELMHOLTZ EQUATIONS WITH A NONLOCAL BOUNDARY CONDITION*

ROBERT C. KIRBY[†], ANDREAS KLÖCKNER[‡], AND BEN SEPANSKI[§]

Abstract. Numerical resolution of exterior Helmholtz problems requires some approach to domain truncation. As an alternative to approximate nonreflecting boundary conditions and invocation of the Dirichlet-to-Neumann map, we introduce a new, nonlocal boundary condition. This condition is exact and requires the evaluation of layer potentials involving the free space Green's function. However, it seems to work in general unstructured geometry, and Galerkin finite element discretization leads to convergence under the usual mesh constraints imposed by Gårding-type inequalities. The nonlocal boundary conditions are readily approximated by fast multipole methods, and the resulting linear system can be preconditioned by the purely local operator involving transmission boundary conditions.

Key words. finite element method, Helmholtz, boundary conditions, layer potentials

AMS subject classifications. 65N30, 65N80, 65F08

1. Introduction. The exterior Helmholtz problem plays an essential role in scattering problems and also serves as a starting point to consider exterior problems in electromagnetics and problems in other unbounded domains such as waveguides. The literature contains several techniques to address the challenge that unbounded domains pose for numerical methods. Essentially, these techniques include some combination of truncating the domain to a bounded one, posing boundary conditions that enforce (or approximate) the Sommerfeld condition on the newly-introduced boundary, and/or modifying the PDE near the computational boundary to absorb any reflected waves.

An early paper on finite elements for the exterior problem is [20], where the domain is truncated at radius R and an approximate radiation condition is posed at R . Although the error estimates contain a factor of R^{-2} , it is also possible to carefully increase the mesh spacing near the boundary, somewhat mitigating the cost of a large domain. *Perfectly matched layers* [3] modify the PDE near the boundary of the computational domain, changing the coefficient of the elliptic term to ‘absorb’ outgoing waves. While such methods allow small effective computational domains, the resulting linear systems do not yield readily to standard iterative techniques like multigrid, although we refer to recent work [44] that poses a domain decomposition strategy to use a direct solver only near the boundary and standard iterative techniques inside.

There is also considerable literature on *nonlocal* boundary conditions for domain truncation. Following early work [19, 25, 28], one can use a Dirichlet-to-Neumann (DtN) operator on the artificial boundary to enforce proper far-field behavior. Givoli [15] provides a survey of similar techniques and local conditions as well, and [16, 22] give techniques for the time rather than frequency domain case. The DtN

*Submitted to the editors DATE.

Funding: This work was funded by NSF awards SHF-1909176 and SHF-1911019. This material is based upon work supported by the U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing Research, Department of Energy Computational Science Graduate Fellowship under Award Number DE-SC0021110.

[†]Baylor University, Department of Mathematics, Waco, TX (robert.kirby@baylor.edu, <https://sites.baylor.edu/robert.kirby/>).

[‡]Department Computer Science, University of Illinois at Urbana-Champaign, Urbana, IL (andreas@illinois.edu)

[§]University of Texas, Department of Computer Science, Austin, TX (ben.sepanski@utexas.edu)

is typically given as an infinite series (truncated in computation) obtained by separating variables. This limits the shape of the domain boundary, although perturbations of such domains and use of high-order methods are possible [6, 35, 36]. Careful error analysis for finite element discretizations can include the effect of truncating the infinite series as well as polynomial approximation error [33]. Lastly, it is worth noting that boundary integral equation methods solve exterior Helmholtz problems with optimal complexity (linear in the number of boundary degrees of freedom), however they are somewhat more difficult to adapt than (volume) PDE discretizing-methods to specific (and potentially nonlinear) near-surface physics.

In this paper, we propose an alternative nonlocal boundary condition based on Green's Theorem [46] that has several important features. Like the DtN approach, we have an (in principle) exact boundary condition, incurring no error in our domain truncation. However, because we rely on the free-space Green's function, there is (again, in principle) no restriction on the shape of our computational domain. The layer potentials appearing in our nonlocal boundary condition can be efficiently computed by appropriate fast algorithms such as variants of the Fast Multipole Method [8]. So, although Galerkin's method would give matrices with dense sub-blocks, we can quickly compute the matrix-vector product required in a Krylov method. Because our nonlocal operator involves double integrals over distinct boundaries, we avoid the need to evaluate any singular integrals. We note that “two-boundary” approaches have been explored in the time-domain literature [17, 24, 47]. Finally, the local part of the operator (a standard finite element matrix) serves as an excellent preconditioner for the system, so that an optimal solver for the local part would give $\mathcal{O}(n \log n)$ solution time in unstructured geometry. Our method works equally well in two and three space dimensions.

Another method combining boundary integral and volumetric discretizations is due to Johnson and Nédélec [27, 45]. This technique encloses a compactly-supported volume source in a truncating boundary. Finite elements are used to compute the solution inside the boundary and a boundary integral method is used on the boundary to handle the exterior. Our present method bears some similarity, employing the same kind of operators. However, we require only a single finite element space and do not introduce additional unknowns on the domain boundary.

In the rest of the paper, we pose the model and its finite element discretization in Section 2. We describe a preconditioned Krylov system for this system in Section 3. Our implementation, which relies on the high-level codes Firedrake [40] and PYTENTIAL [32], warrants some discussion, which is given in Section 4. Finally, we give numerical results in Section 5.

2. Model and discretization. Let $\Omega^c \subset \mathbb{R}^d$ with $d = 2, 3$ be a bounded domain with boundary Γ , and $\Omega = \mathbb{R}^d \setminus \Omega^c$ its exterior. We consider the classic Helmholtz exterior problem on Ω

$$(2.1) \quad -\Delta u - \kappa^2 u = 0,$$

where κ^2 is nonzero wave number. It may be complex (typically with positive real part), and may take different forms in various application fields such as acoustics or electromagnetics. We also pose Neumann boundary conditions

$$(2.2) \quad \frac{\partial u}{\partial n} = f$$

on the interior boundary Γ . The Sommerfeld radiation condition

$$(2.3) \quad \lim_{r \rightarrow \infty} r^{\frac{d-1}{2}} \left(\frac{\partial u}{\partial r} - i\kappa u \right) = 0,$$

where r is the outward radial direction, must also hold. For computational purposes, one typically poses the problem only on a truncated domain Ω' . Hence, we impose an artificial boundary Σ , and let Ω' denote that subset of Ω enclosed between Γ and Σ . We assume that these boundaries are such that Ω' is a Lipschitz domain. An example is shown in Figure 1:

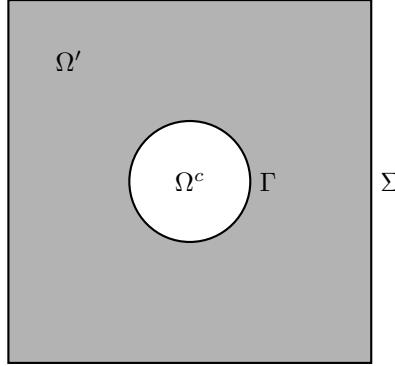


FIG. 1. A 2D example of a truncated domain

A major challenge for volume-discretizing numerical methods is imposing a suitable boundary condition on Σ . For example, a simple approach is to impose the Robin-type condition

$$(2.4) \quad i\kappa u - \frac{\partial u}{\partial n} = 0$$

on Σ rather than at infinity. Frequently called “transmission” boundary conditions, this changes the boundary value problem, incurring errors that do not vanish under mesh refinement, and can create artificial wave reflections at the boundary.

We propose a new approach to the problem that, for constant-coefficient problems at least, allows highly effective iterative solvers to be combined with effective domain truncation. Let

$$\mathcal{K}(x) := \begin{cases} u(x) = \frac{i}{4} H_0^{(1)}(\kappa|x|) & d = 2, \\ u(x) = \frac{i}{4\pi|x|} e^{i\kappa|x|} & d = 3. \end{cases}$$

be the free-space Green’s function for the Helmholtz equation, where $H_0^{(1)}(x)$ be the first-kind Hankel function of index 0. Recall Green’s formula in the exterior [9, Thm. 2.5] for the solution $u(x)$ to (2.1):

$$(2.5) \quad u(x) = \int_{\Gamma} \left(\frac{\partial}{\partial n} \mathcal{K}(x-y) \right) u(y) - \left(\frac{\partial}{\partial n} u(y) \right) \mathcal{K}(x-y) dy,$$

for $x \in \Omega'$. It is known that Green’s theorem holds in general Lipschitz domains [41].

Substituting in the Neumann boundary condition (2.2), we obtain

$$(2.6) \quad \begin{aligned} u(x) &= \int_{\Gamma} \left(\frac{\partial}{\partial n} \mathcal{K}(|x-y|) \right) u(y) - f(y) \mathcal{K}(|x-y|) dy \\ &\equiv D(u)(x) - S(f)(x), \end{aligned}$$

for $x \in \Omega'$, where

$$(2.7) \quad D(u)(x) = \text{PV} \int_{\Gamma} \left(\frac{\partial}{\partial n} \mathcal{K}(x - y) \right) u(y) dy$$

is the double layer potential and

$$(2.8) \quad S(u)(x) = \int_{\Gamma} \mathcal{K}(x - y) u(y) dy$$

is the single layer potential [9].

Now, we can pose an exact nonlocal Robin-type boundary condition as follows. We use the representation (2.5) to write (suppressing the argument x):

$$(2.9) \quad i\kappa u - \frac{\partial u}{\partial n} = i\kappa (D(u) - S(f)) - \frac{\partial}{\partial n} (D(u) - S(f)),$$

so that over Σ ,

$$(2.10) \quad \frac{\partial u}{\partial n} = i\kappa u - \left(i\kappa - \frac{\partial}{\partial n} \right) (D(u) - S(f)).$$

2.1. Variational Setting. We let $(f, g) = \int_{\Omega'} f(x) \overline{g(x)} dx$ be the standard L^2 inner product over the computational domain, and $\langle f, g \rangle_S$ that over some portion $S \subseteq \partial\Omega'$ of its boundary. We also let $H^s(\Omega')$ be the standard Sobolev spaces consisting of L^2 functions with weak derivatives of order up to and including s in L^2 .

When X is some Banach space, $\|\cdot\|_X$ refers to its norm. As we use several different norms throughout our analysis, we explicitly label each such norm to limit confusion.

We give a variational formulation of the PDE and hence a standard Galerkin finite element discretization as follows. We take the inner product of (2.1) with any $v \in H^1(\Omega')$. Integration by parts and the Neumann boundary condition on Γ give

$$(2.11) \quad (\nabla u, \nabla v) - \kappa^2 (u, v) - \langle \frac{\partial u}{\partial n}, v \rangle_{\Sigma} = \langle f, v \rangle_{\Gamma},$$

and substituting (2.10) in for $\frac{\partial u}{\partial n}$ on Σ gives

$$(2.12) \quad \begin{aligned} (\nabla u, \nabla v) - \kappa^2 (u, v) - i\kappa \langle u, v \rangle_{\Sigma} + \langle \left(i\kappa - \frac{\partial}{\partial n} \right) D(u), v \rangle_{\Sigma} \\ = \langle f, v \rangle_{\Gamma} + \langle \left(i\kappa - \frac{\partial}{\partial n} \right) S(f), v \rangle_{\Sigma}. \end{aligned}$$

Hence, the solution to the Helmholtz equation (2.1) on Ω together with (2.2) and (2.3) satisfies the variational problem of finding $u \in H^1(\Omega')$ such that

$$(2.13) \quad a(u, v) = F(v)$$

for all $v \in H^1(\Omega')$. Here, the bilinear form

$$(2.14) \quad a(u, v) = (\nabla u, \nabla v) - \kappa^2 (u, v) - i\kappa \langle u, v \rangle_{\Sigma} + \langle \left(i\kappa - \frac{\partial}{\partial n} \right) D(u), v \rangle_{\Sigma}$$

consists of the standard bilinear form using transmission boundary conditions (2.4) augmented by nonlocal terms involving a convolution-type integral with a Green's function kernel. We write $a(u, v) = a_L(u, v) + a_{NL}(u, v)$, where

$$(2.15) \quad \begin{aligned} a_L(u, v) &= (\nabla u, \nabla v) - \kappa^2 (u, v) - i\kappa \langle u, v \rangle_{\Sigma}, \\ a_{NL}(u, v) &= \langle \left(i\kappa - \frac{\partial}{\partial n} \right) D(u), v \rangle_{\Sigma}. \end{aligned}$$

Similarly, the linear form

$$(2.16) \quad F(v) = \langle f, v \rangle_\Gamma + \langle (i\kappa - \frac{\partial}{\partial n}) S(f), v \rangle_\Sigma$$

involves the Neumann data on the scatterer together with its appearance in the single layer potential.

By taking $V_h \subset H^1(\Omega')$ as any suitable finite element space, we can introduce a Galerkin finite element method of finding $u_h \in V_h$ such that

$$(2.17) \quad a(u_h, v_h) = F(v_h)$$

for all $v_h \in V_h$.

At this point, we pause compare our method to the Dirichlet-to-Neumann map P . (In the literature, the same operator is sometimes called the Steklov-Poincaré operator. Generically, S-P operators convert one type of boundary data into another.) Replacing $\frac{\partial u}{\partial n}$ on Σ in (2.11) with P acting on u would give

$$(\nabla u, \nabla v) - \kappa^2 (u, v) + \langle Pu, v \rangle_\Sigma = \langle f, v \rangle_\Gamma.$$

Compared to (2.14), this appears to only have a single nonlocal term. Moreover, P is a symmetric elliptic operator from $H^{1/2}(\Sigma)$ into $H^{-1/2}(\Sigma)$, so that $\langle Pu, u \rangle \geq 0$ and a Gårding estimate readily holds for the bilinear form. Unfortunately, the Steklov-Poincaré operator is not typically explicitly available, and thus its application requires the solution of a linear system at additional computational cost, e.g. in the form of a boundary integral equation solve. Approximating Pu with a Fourier series is possible, however doing so requires separable geometry.

2.2. Convergence theory. Our argument will rely on showing the boundedness of the bilinear form a and establishing a Gårding-type inequality. Using standard techniques [7], this leads to discrete solvability and optimal *a priori* error estimates under a constraint on the maximal mesh size.

We will rely on the trace estimates [7, 21] that since Ω' is Lipschitz, there exists a constant C such that

$$(2.18) \quad \|v\|_{L^2(\partial\Omega')} \leq C \|v\|_{L^2(\Omega')}^{1/2} \|v\|_{H^1(\Omega')}^{1/2} \leq C \|v\|_{H^1(\Omega')}$$

for all $v \in H^1(\Omega')$.

PROPOSITION 2.1. *If the Neumann data satisfies $f \in H^{-\frac{1}{2}}(\Gamma)$, the functional F defined in (2.16) is a bounded linear functional on H^1 .*

Proof. Linearity is clear from the linearity of integration and differentiation. To see that it is bounded, let $v \in H^1(\Omega)$ be given. The local portion of F is bounded thanks to Cauchy-Schwarz and the second trace estimate in (2.18). For the nonlocal portion, it is known [46] that $S(f) \in H^1(\Omega')$ and so it has a normal derivative on Σ in $H^{-1/2}(\Sigma)$. \square

The following result implies both the boundedness of a_{NL} on $H^1 \times H^1$ and is critical to establishing the Gårding inequality:

LEMMA 2.2. *There exists a $C_{NL} > 0$ such that for all $u, v \in H^1(\Omega')$,*

$$(2.19) \quad |a_{NL}(u, v)| \leq C_{NL} \|u\|_{L^2(\Gamma)} \|v\|_{L^2(\Sigma)}.$$

Proof. First, we simplify the notation by writing the first argument in a_{NL} as

$$(2.20) \quad (i\kappa - \frac{\partial}{\partial n}) D(u) = (i\kappa - \frac{\partial}{\partial n}) \int_{\Gamma} \mathcal{K}(x - y) u(y) dy,$$

From the properties of the kernel, $\mathcal{K}(x - y)$ is smooth and bounded provided that $\|x - y\|$ is bounded below away from zero. Since we have $x \in \Sigma$ and $y \in \Gamma$, this is the case as long as the truncating boundary stays away from the scatterer. By writing the normal derivative in (2.20) as the limit of a difference quotient, passing under the integral, and appealing to the Lebesgue Dominated Convergence Theorem in the usual way, we can then write

$$(2.21) \quad (i\kappa - \frac{\partial}{\partial n}) D(u) = \int_{\Gamma} (i\kappa - \frac{\partial}{\partial n}) \mathcal{K}(x - y) dy = \int_{\Gamma} \tilde{\mathcal{K}}(x - y) dy,$$

where $\tilde{\mathcal{K}}(x - y)$ is also smooth and bounded for x and y separated. We can write the nonlocal bilinear form now as

$$(2.22) \quad |a_{NL}(u, v)| = \left| \int_{\Gamma} \int_{\Sigma} \tilde{\mathcal{K}}(x, y) u(y) \overline{v(x)} dx dy \right| \leq K_0 \int_{\Gamma} u(y) dy \int_{\Sigma} \overline{v(x)} dx,$$

and the result holds with $C_{NL} = K_0 |\Gamma|^{1/2} |\Sigma|^{1/2}$ by the Cauchy-Schwarz inequality. \square

PROPOSITION 2.3. *There exists $C_B > 0$ such that for all $u, v \in H^1(\Omega')$,*

$$(2.23) \quad |a(u, v)| \leq C_B \|u\|_{H^1(\Omega')} \|v\|_{H^1(\Omega')}.$$

Proof. Let $u, v \in H^1(\Omega')$. Then

$$(2.24) \quad \begin{aligned} |a(u, v)| &\leq \|\nabla u\|_{L^2(\Omega')} \|\nabla v\|_{L^2(\Omega')} \\ &\quad + \kappa^2 \|u\|_{L^2(\Omega')} \|v\|_{L^2(\Omega')} + \kappa \|u\|_{L^2(\Sigma)} \|v\|_{L^2(\Sigma)} + |a_{NL}(u, v)|, \end{aligned}$$

and the proof is finished by applying the previous Lemma and trace theorem. \square

The bilinear form a satisfies a Gårding inequality. That is, shifting a by a multiple of the L^2 inner product renders a coercive bilinear form. For complex Hilbert spaces, it is sufficient to demonstrate that the real part itself is coercive.

PROPOSITION 2.4. *There exists a real number M and an $\alpha > 0$ such that*

$$(2.25) \quad \operatorname{Re}(a(u, u)) + M \|u\|_{L^2(\Omega')}^2 \geq \alpha \|u\|_{H^1(\Omega')}^2.$$

Proof. We calculate

$$(2.26) \quad a(u, u) = \|\nabla u\|_{L^2(\Omega')}^2 - \kappa^2 \|u\|_{L^2(\Omega')}^2 - i\kappa \|u\|_{L^2(\Sigma)}^2 + a_{NL}(u, u),$$

and note that the real part of this is just

$$(2.27) \quad \begin{aligned} \operatorname{Re}(a(u, u)) &= \|\nabla u\|_{L^2(\Omega')}^2 - \kappa^2 \|u\|_{L^2(\Omega')}^2 + \operatorname{Re}(a_{NL}(u, u)) \\ &= \|u\|_{H^1(\Omega')}^2 - (\kappa^2 + 1) \|u\|_{L^2(\Omega')}^2 + \operatorname{Re}(a_{NL}(u, u)). \end{aligned}$$

Using Lemma 2.2, the trace inequality (2.18), and a weighted Young's inequality, we can bound this below by

$$(2.28) \quad \begin{aligned} \operatorname{Re}(a(u, u)) &\geq \|u\|_{H^1(\Omega')}^2 - (\kappa^2 + 1) \|u\|_{L^2(\Omega')}^2 - C_{NL} \|u\|_{L^2(\Sigma)} \|u\|_{L^2(\Gamma)} \\ &\geq \|u\|_{H^1(\Omega')}^2 - (\kappa^2 + 1) \|u\|_{L^2(\Omega')}^2 - C_{NL} \|u\|_{H^1(\Omega')} \|u\|_{L^2(\Omega')} \\ &\geq \frac{1}{2} \|u\|_{H^1(\Omega')}^2 - \left(\kappa^2 + 1 + \frac{C_{NL}^2}{2} \right) \|u\|_{L^2(\Omega')}^2, \end{aligned}$$

so (2.26) holds with $\alpha = \frac{1}{2}$ provided $M \geq \kappa^2 + \frac{3+C_{NL}^2}{2}$. \square

Now, following standard techniques for general elliptic (but possibly not coercive) problems [7], suitably adapted for the complex-valued case, we have a general solvability and approximation result. We suppose that the standard abstract approximation result

$$(2.29) \quad \inf_{v \in V_h} \|u - v\|_{H^1(\Omega')} \leq C_A h |u|_{H^2(\Omega')}$$

holds and that the solution to (2.1) is in $H^2(\Omega')$. We also require that the adjoint problem of finding $w \in H^1(\Omega')$ such that

$$(2.30) \quad a(v, w) = (f, v)$$

for all $v \in H^1(\Omega')$ has a unique solution with regularity estimate

$$(2.31) \quad |u|_{H^2(\Omega')} \leq C_R \|f\|_{L^2(\Omega')}.$$

With these assumptions, the arguments leading to Theorem 5.7.6 of [7] give this result:

THEOREM 2.5. *Under the above conditions, there exists h_0 such that for $h \leq h_0$, the discrete variational problem (2.32) has a unique solution u_h satisfying the error estimate*

$$(2.32) \quad \|u - u_h\|_{H^1(\Omega')} \leq C \inf_{v \in V_h} \|u - v\|_{H^1(\Omega')}.$$

Moreover, with the same assumptions, there exists another $C > 0$ such that

$$(2.33) \quad \|u - u_h\|_{L^2(\Omega')} \leq Ch \|u - u_h\|_{H^1(\Omega')}.$$

Note that this is a quasi-optimal result, independent of the particular choice of polynomial spaces. So, it gives error estimates for higher-order approximations as well as for standard P^1 elements.

Remark 2.6. In particular, following [7], one can show

$$(2.34) \quad h_0 = \frac{\left(\frac{\alpha}{2M}\right)^{1/2}}{C_B C_A C_R},$$

and the constant C in (2.32) can be taken as $\frac{2C_B}{\alpha}$ and that in (2.33) as $C_B C_A C_R$.

Remark 2.7. This convergence theory assumes the layer potentials and boundary integrals are evaluated exactly. These results can be extended to account for approximation to the layer potential along the lines of [34, Thm. 13.6/7] and quadrature in the bilinear forms using the standard theory of variational crimes [7].

3. Linear algebra. We can effectively solve our variational formulation using preconditioned GMRES, which is a parameter-free algorithm approximating the solution of a linear system $A\mathbf{x} = \mathbf{b}$ as the element of the Krylov subspace $\text{span}\{A^i\mathbf{b}\}_{i=0}^m$ minimizing the equation residual. Building the subspace does not require the entries of A , just the action of A on vectors. Unlike conjugate gradients, GMRES is not restricted to operators that are symmetric and positive definite.

For most problems arising in the discretization of PDE, the condition number of A degrades quickly under mesh refinement, and GMRES is most frequently used

in conjunction with a (left) *preconditioner*. Mathematically, we multiply the linear system through by some matrix \widehat{P}^{-1} :

$$(3.1) \quad \widehat{P}^{-1} A \mathbf{x} = \widehat{P}^{-1} \mathbf{b},$$

and so the Krylov space then is $\text{span}\{\left(\widehat{P}^{-1} A\right)^i \widehat{P}^{-1} \mathbf{b}\}_{i=0}^m$.

The overall performance of GMRES typically is determined by two factors – the cost of building and applying the operators \widehat{P}^{-1} and A , and the total number of iterations. One hopes to obtain a per-application cost that scales linearly (or log-linearly) with respect to the number of unknowns in the linear system, and a total number of GMRES iterations that is bounded independently of the number of unknowns. We think of \widehat{P}^{-1} being an approximation to the inverse of some matrix P that somehow approximates A . In our case, we will find it useful to let $P = A^L$, the local part of the operator. Then, applying \widehat{P}^{-1} might correspond to applying the inverse of P by a sparse direct method or perhaps just some sweeps of a multigrid algorithm.

3.1. Structure of the discrete problem. By taking a standard finite element basis $\{\phi_i\}_{i=1}^N$ for V_h , the stiffness matrix is

$$(3.2) \quad A_{ij} = a(\phi_j, \phi_i) = a_L(\phi_j, \phi_i) + a_{NL}(\phi_j, \phi_i) = A_{ij}^L + A_{ij}^{NL}.$$

The portion A^L is the standard sparse matrix one obtains for discretization of the Helmholtz operator with transmission boundary conditions, while A^{NL} contains the contributions for the nonlocal terms. To further consider the sparsity of this system, supposing we use standard P^1 basis functions and have about $\mathcal{O}(N^d)$ total vertices and hence basis functions. Then A^L has nonzero entries corresponding to vertices sharing a common mesh element – typically about 6-7 nonzeros per row on two-dimensional triangulations and 20-30 for three-dimensional tetrahedral meshes when using linear basis functions. The total storage required for A^L will be proportional to the number of vertices in the mesh.

Explicit sparse storage of A^{NL} , however, can be quite different. Since K involves a convolution-type integral over Γ ,

$$(3.3) \quad A_{ij}^{NL} = \langle (i\kappa - \frac{\partial}{\partial n}) K(\phi_j), \phi_i \rangle_{\Sigma}$$

will be nonzero whenever ϕ_j is supported on Γ and ϕ_i is supported on Σ . Suppose that we have $\mathcal{O}(N^{d-1})$ basis functions supported on Σ and the same order on Γ . Then, A^{NL} will be nonzero except for a dense logical subblock. However, each basis function associated with Σ will interact with each basis functions associated with Γ , so that the dense subblock will contain about $\mathcal{O}(N^{2d-2})$ nonzero entries. When $d = 2$, A^{NL} has the same order of nonzeros as A^L and so conceivably could be stored explicitly. On the other hand, when $d = 3$, A^{NL} has $\mathcal{O}(N^4)$ nonzero entries and so its storage dominates that of the local part A^L . Consequently, a matrix-free application of A^{NL} that bypasses the storage may be preferred, as described in Section 4.2.

3.2. Operator application. From (3.2), the system matrix $A = A^L + A^{NL}$ is the sum of two matrices corresponding to the local and nonlocal terms in the bilinear form. Although we could implement a matrix-free action of A^L , we opt to assemble a standard sparse matrix and only apply A^{NL} in a matrix-free fashion as follows.

Recall that $A^{NL} = a_{NL}(\phi_j, \phi_i) = \langle (i\kappa - \frac{\partial}{\partial n}) D(\phi_j), \phi_i \rangle$. Any vector $\mathbf{x} \in \mathbb{R}^{\dim V_h}$ can be identified uniquely with some v_h in the finite element space so that

$$(3.4) \quad (A^{NL} \mathbf{x})_i = a_{NL}(v_h, \phi_i).$$

Note that $(A^{NL}\mathbf{x})_i$ will be nonzero exactly when ϕ_i has support on the exterior boundary Σ .

In a startup phase, we prepare the boundary geometry according to the algorithm of [50] to construct a GIGAQBX tree structure [39, 49, 50, 51] for the approximation of the layer potentials D and S . These allow us to efficiently approximate $(i\kappa - \frac{\partial}{\partial n}) D(v_h)$ at a collection of ‘target’ points. In particular, we can evaluate on quadrature points on each facet of Σ . Hence, we can loop over the facets on Σ to integrate against the basis functions supported on that facet and sum the contributions in the usual way. This gives the action of A^{NL} onto some \mathbf{x} , and the full action of A onto \mathbf{x} is computed by summing this with $A^L\mathbf{x}$ computed by a standard sparse matrix-vector product.

3.3. Preconditioners. Rather than letting the preconditioning matrix P equal A itself, we opt for $P = A^L$. If we were to exactly invert A^L , then the resulting system becomes

$$(3.5) \quad \left(I + (A^L)^{-1} A^{NL} \right) x = (A^L)^{-1} b.$$

Since A^L discretizes an elliptic equation and A^{NL} a bounded operator, this has the form of a discretization of a compact perturbation of the identity. In [18], GMRES convergence for a similar situation was shown to be very favorable. We also comment that preconditioning a system to obtain a compact perturbation of the identity was used heuristically to good effect for Bénard convection [26].

It is possible to replace the inverse of A^L with an approximation, and it is likewise possible to use a suitable, spectrally equivalent preconditioner, such as algebraic multigrid [1, 43, 37]. This gives a preconditioner that scales well with mesh refinement, but can degrade as the wave number increases [12, 13].

4. Implementation. Our implementation rests on combining the capabilities of Firedrake [40] for the finite element part of our problem and PYTENTIAL [32] for the evaluation of layer potentials K and S . Krylov solvers and preconditioners are accessed using PETSc.

4.1. Firedrake. Firedrake [40] is an automated system for the solution of partial differential equations using the finite element method. It allows users to describe the variational form of a PDE using the Unified Form Language [2], from which it generates effective lower-level numerical code. We make use of Firedrake for loading computational meshes, defining the local part of a , and integrating evaluated layer potentials against test functions. Firedrake can be built supporting complex arithmetic at every level (definition of bilinear forms down to a complex-enabled PETSc build).

Firedrake also makes it possible to compare our new method against domain truncation by means of a perfectly matched layer (PML). We implement the technique of [4] which uses an unbounded integral as the absorbing function on the PML. This approach is parameter-free and simple to implement in UFL.

4.2. Pytential. PYTENTIAL [32] is an open-source, MIT licensed software system that allows the evaluation of layer potentials from source geometry represented by unstructured meshes in two and three dimensions with near-optimal complexity and at a high order of accuracy. The main aspects of functionality provided by PYTENTIAL are the discretization of a source surface using discretization tools (through its use of a sister tool, MESHMODE [30]) for high-order accurate nonsingular quadrature [52],

its refinement according to accuracy requirements [50], and, finally, the evaluation of weakly singular, singular, and hypersingular integral operators via quadrature by expansion (QBX) [29] and the associated GIGAQBX fast algorithm [49], with rigorous accuracy guarantees in two and three dimensions [51]. This fast algorithm, can, in turn make use of FMMLIB [14, 31] for the evaluation of translation operators in the moderate-frequency regime for the Helmholtz equation.

While the layer potential evaluations in a_{NL} are nonsingular, we nonetheless benefit from the use of the QBX machinery in the event that source and target surfaces are chosen to lie in close proximity for increased efficiency of the finite element method. See [50] for estimates of the error incurred in the evaluation of the layer potential. Our 2D experiments employ higher order discretizations and fine meshes, so we set FMM order adaptively at each level of the FMM tree to provide a precision stricter than machine epsilon for double precision. Our 3D experiments are limited to piecewise linear elements on coarse meshes, so it is sufficient to use an FMM order of 12.

4.3. Representing the linear system in PETSc. At the top level, our code builds and solves the linear system (3.2). To do this, we have implemented a PYTHON matrix type in PETSC4PY [10]. Its Python context builds the bilinear form a_L in Firedrake and assembles A^L . It also sets up the layer potential evaluation in PYTENTIAL. The class also provides a multiplication method that multiplies by A^L (itself just a PETSC call) and A^{NL} (which requires more code) and sums the results. It also provides a handle to A^L so that it can be used as a preconditioning matrix. Setting up a KSP context in PETSc, we can then select from any available Krylov method and apply any preconditioning technique to A^L in the standard ways.

The application of A^{NL} to a vector requires some low-level interaction of Firedrake and PYTENTIAL beneath their public interfaces, and warrants some explanation. Data transfer between PYTENTIAL and Firedrake occurs in two directions. The transfer of density information from Firedrake to PYTENTIAL occurs through (exact) interpolation within P^N from the C^0 finite element space used for a_L to the discontinuous finite element space on Vioreanu-Rokhlin nodes [48] used for the density in a_{NL} . This requires some attention to details regarding ordering of degrees of freedom, vertices [42], and data formats. The transfer of layer potential information back to Firedrake meanwhile is straightforward by comparison. PYTENTIAL is able to evaluate the layer potential with guaranteed accuracy *anywhere* in the target domain, even close to the source surface, where this might otherwise require special treatment such as near-singular quadrature, e.g. by adaptive techniques. Thus we merely evaluate the layer potential at a set of quadrature points supplied by Firedrake to obtain an approximate projection of the (analytically) C^∞ potential back into the C^0 finite element space.

5. Numerical results. Now, we present some empirical investigation of our method. We establish the accuracy obtained using finite element approximation using our nonlocal boundary condition and also consider preconditioning the nonlocal boundary system. We find that the accuracy obtained using the nonlocal boundary condition compares favorably with that rendered by PML and the transmission boundary condition. Moreover, when methods are available to accurately approximate the inverse of A^L , we find that it is an excellent preconditioner for the overall system. However, as the wave number increases, the difficulty of attaining an accurate approximation increases.

To verify the accuracy of our method in two and three space dimensions, we chose the unit disc/sphere as a scatterer and use a manufactured solution based on the free-

space Helmholtz Green's function. That is, the true solution outside of the scatterer is taken as

$$\begin{cases} u(x) = \frac{i}{4} H_0^{(1)}(\kappa|x|) & d = 2, \\ u(x) = \frac{i}{4\pi|x|} e^{i\kappa|x|} & d = 3. \end{cases}$$

In two dimensions, we truncated the domain as the 6×6 square centered at the origin, shown in Figure 2 with the PML region separately highlighted. In three dimensions, we created an analogous mesh, the unit sphere embedded in $[-3, 3]^3$, with the PML sponge region taking up $[-3, 3]^3 \setminus [-2, 2]^3$, shown in Figure 3.

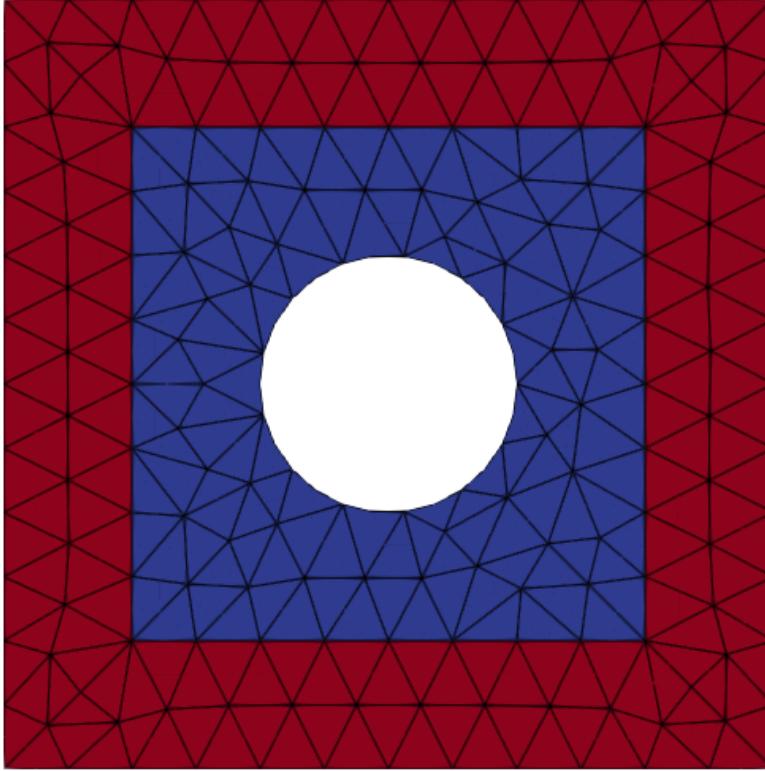


FIG. 2. Example 2d mesh (with PML region colored in red)

We compare our approach with transmission boundary conditions and with PML. Since the PML-based simulation is only accurate in the non-PML region (the blue region in Figure 2), we evaluate the error only over the non-PML region for all methods, even though we solve over the entire computational domain. For computations with piecewise linear approximating spaces, we use affine geometry, although we use quadratic geometry for computations with higher-order spaces.

We implement both PML and transmission boundary conditions in Firedrake. For PML, we use unbounded absorbing functions as described in [4]. This form of PML has no parameters to be fitted, is simple to implement in Firedrake, and has been shown to recover the exact solution (up to discretization errors) on annular domains [5]. However, other implementations of PML (or equivalent boundary conditions) can obtain higher accuracy at lower wavelengths [11, 23].

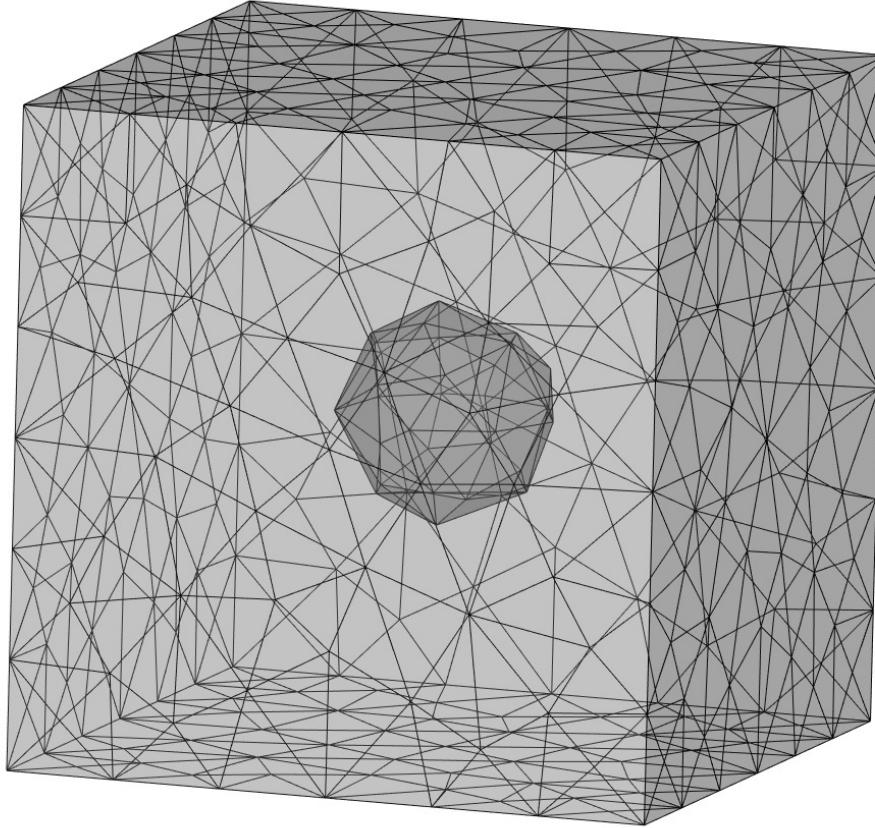


FIG. 3. Sample coarse 3D mesh with a spherical exclusion at the center of a cube.

We compare the accuracy of transmission boundary conditions, PML, transmission, and our new approach in 2D for degree 1 approximations in Figure 4. We observe that the transmission boundary conditions, which incur a perturbation of the PDE, lead to convergence to a slightly incorrect solution. Both PML and our new boundary conditions, however, seem to be converging to the true solution at the proper rate of $\mathcal{O}(h^2)$ predicted in Theorem 2.5. For small κ , the nonlocal condition seems quite a bit more accurate, although they give nearly the same error for larger κ . Accuracy for the 2D case is reported for higher degrees in Figures 5, 6, and 7 respectively. In these case, we observe the theoretically-predicted convergence rates for our nonlocal method, although for higher degree our (suboptimal) PML does not obtain full accuracy. Our theoretical results apply to 3D as well as 2D, although the computations are considerably more expensive. As a simple test, we have used linear polynomials on rather coarse meshes, presenting the results in Figure 8. We see comparable behavior to that obtained in 2D.

In Figure 9, we study the error obtained versus the number of degrees of freedom using various orders of approximation. In these computations, we use the domain Ω' as $[-2, 2]^2$ minus the unit square and pose the boundary condition (2.10) along the outer boundary. Then, for each $\kappa \in \{0.1, 1, 5, 10\}$ and polynomial degrees 1 through 4, we computed the L^2 error in the numerical approximation. In addition to giving faster convergence rates, we also see that higher-order polynomials provide a

lower error per degree of freedom used. Moreover, no modifications to our method or boundary condition were required to obtain this higher accuracy.

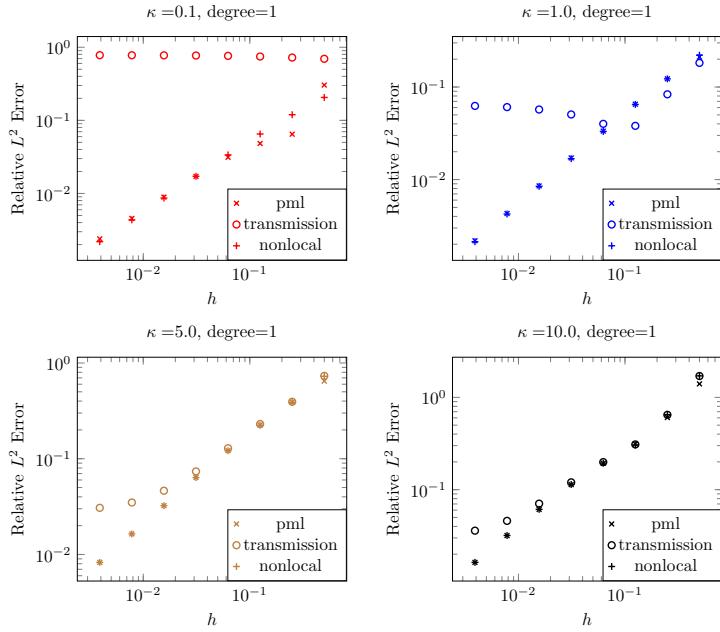


FIG. 4. L^2 Relative error using degree 1 polynomials with respect to refinement in 2D. A KSP relative tolerance of 10^{-12} was used. Since the transmission BC are a perturbation of the actual boundary value problem, the convergence levels off as the method converges to a slightly incorrect solution. The PML and nonlocal methods give comparable accuracy.

We recall that our theoretical results apply to 3D as well as 2D, although the computations are considerably more expensive. As a simple test, we have used linear polynomials on rather coarse meshes, presenting the results in Figure 8. We see comparable behavior to that obtained in 2D.

In Figure 9, we study the error obtained versus the number of degrees of freedom using various orders of approximation. In these computations, we use the domain Ω' as $[-2, 2]^2$ minus the unit square and pose the boundary condition (2.10) along the outer boundary. Then, for each $\kappa \in \{0.1, 1, 5, 10\}$ and polynomial degrees 1 through 4, we computed the L^2 error in the numerical approximation. In addition to giving faster convergence rates, we also see that higher-order polynomials provide a lower error per degree of freedom used. Moreover, no modifications to our method or boundary condition were required to obtain this higher accuracy.

Now, we turn to efficient solution of the linear system, focusing on the two-dimensional case. In table 1 we see that, even for low wave numbers on coarse meshes with a piecewise linear discretization, solving the system without a preconditioner is not scalable.

We want to demonstrate that the local part of our operator (2.4) provides an effective preconditioner, so that our new method can be seen as comparably difficult to solve as the local problem. As a first approach, we can compute a sparse LU factorization of A^L to apply the inverse as a preconditioner for A . The GMRES iteration counts are shown in Figure 10. For a fixed κ , we see mild *decrease* in the iteration

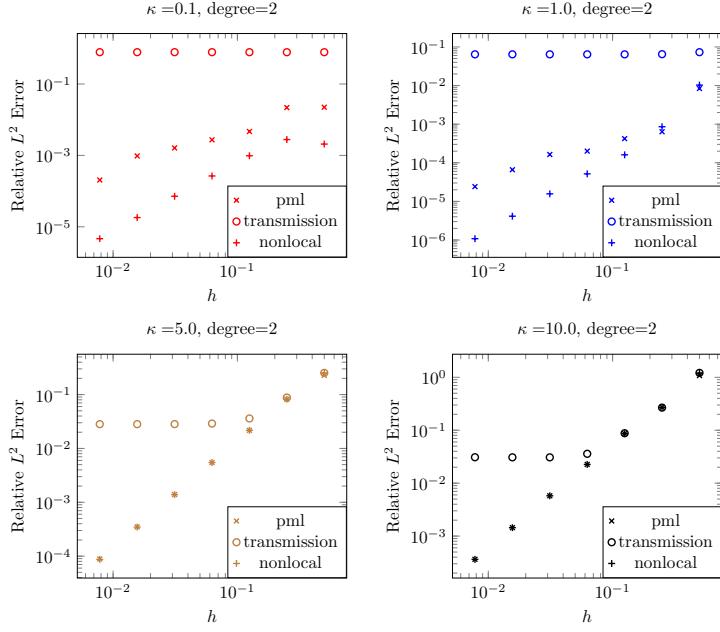


FIG. 5. L^2 Relative error using degree 2 polynomials with respect to refinement in 2D. A KSP relative tolerance of 10^{-12} was used.

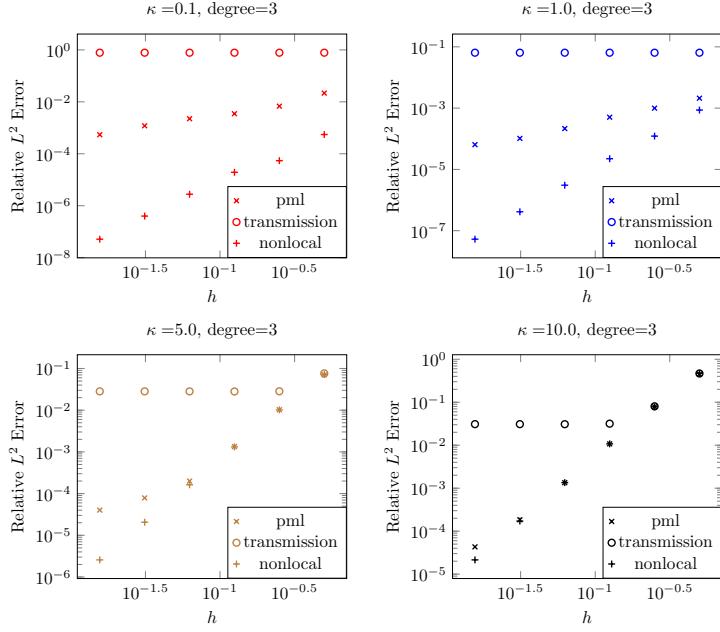


FIG. 6. L^2 Relative error using degree 3 polynomials with respect to refinement in 2D. A KSP relative tolerance of 10^{-12} was used.

count under mesh refinement. Moreover, for a fixed mesh, increasing κ corresponds only to a slight increase in iteration count. So, if the underlying transmission operator

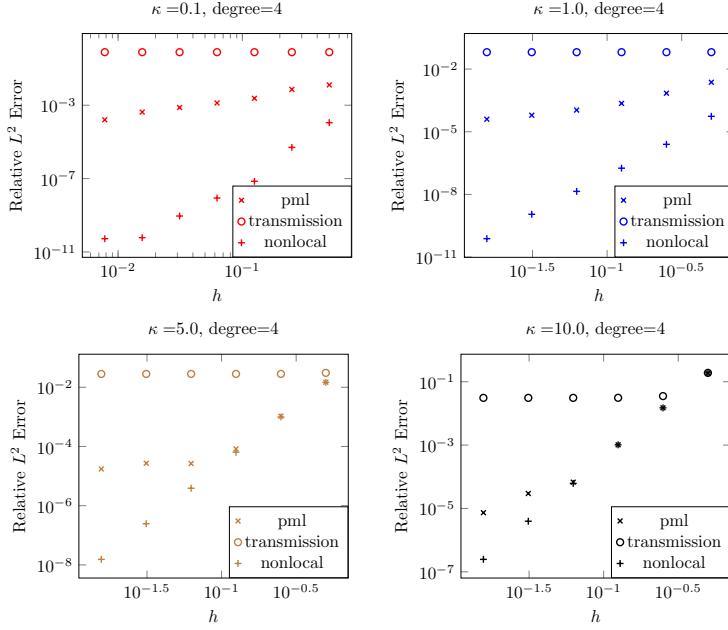


FIG. 7. L^2 Relative error using degree 4 polynomials with respect to refinement in 2D. A KSP relative tolerance of 10^{-12} was used.

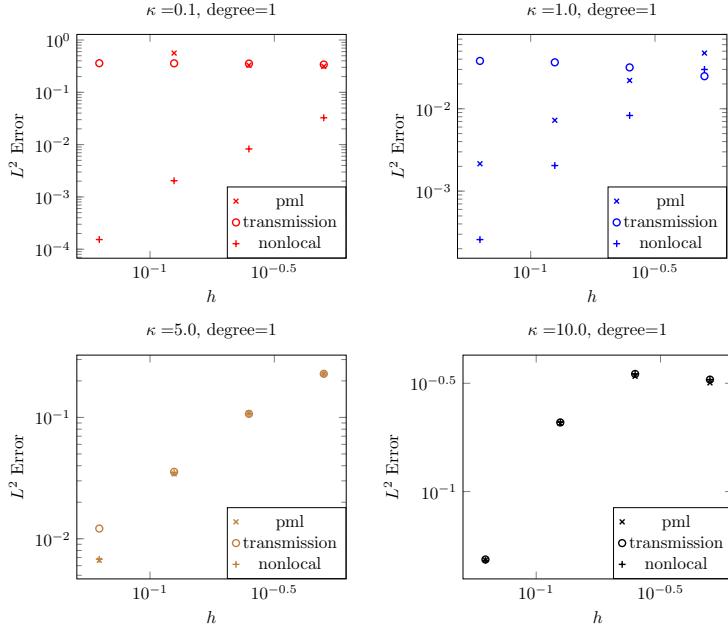


FIG. 8. L^2 Error with respect to refinement in 3D using degree 1 polynomials. A KSP relative tolerance of 10^{-7} was used. Comparable results to Figure 4 are obtained, although we have not been able to attain the same mesh resolutions as in 2D.

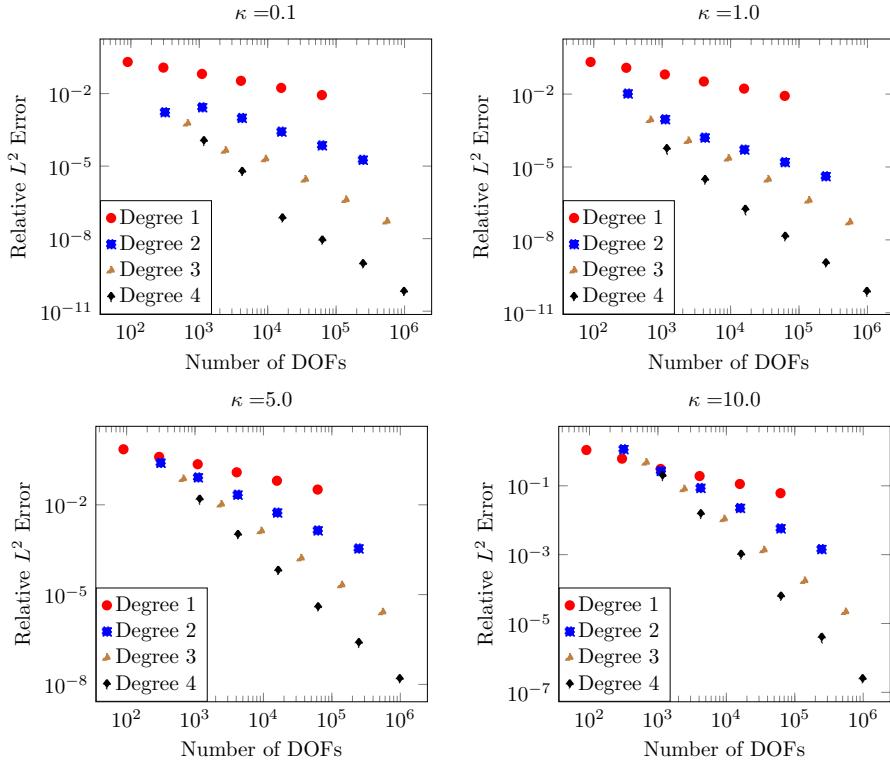


FIG. 9. Relative L^2 error of our nonlocal method for degrees 1-4 as the number of degrees of freedom increase. A KSP relative tolerance of 10^{-12} was used. We used a mesh with no PML sponge region and $\Sigma = [-2, 2]^2$.

can be effectively inverted, this will in turn serve as an excellent preconditioner for the system with nonlocal boundary conditions. Using a direct method on A^L , sparse factorization is typically the dominant cost. Then, each Krylov iteration requires a sparse matrix-vector product, an FMM evaluation, some quadrature, and solution with the sparse factors.

At large enough scale, one might wish to (approximately) invert A^L with an iterative method rather than factorization. To move in this direction, we used `gamg` [1], a PETSc-accessible algebraic multigrid scheme that supports complex arithmetic. This performed admirably at low wave number ($\kappa \lesssim 1$), but not beyond this. We were able to tackle higher wave numbers using the approach in [37]. The Laplacian has eigenmodes that become increasingly oscillatory as the eigenvalues increase. The indefinite Helmholtz operator shifts the eigenvalues (with the same eigenmodes) leftward in the complex plane. Hence, the eigenvalues closest to zero correspond to certain higher-frequency modes for Helmholtz. The technique in [37] approximates this oscillatory near null space with plane waves. To apply this method, we wrapped PyAMG [38] as a PETSc4PY preconditioner. We applied a fixed number of W-multicycles, augmented with plane waves in the same way as in [38], to A^L as a preconditioner for the overall system. Figure 11 shows the results we obtained. The preconditioner is very effective at low κ but requires more iterations for larger ones. However, we see that applying more W-cycles within the preconditioner typically leads to a lower outer

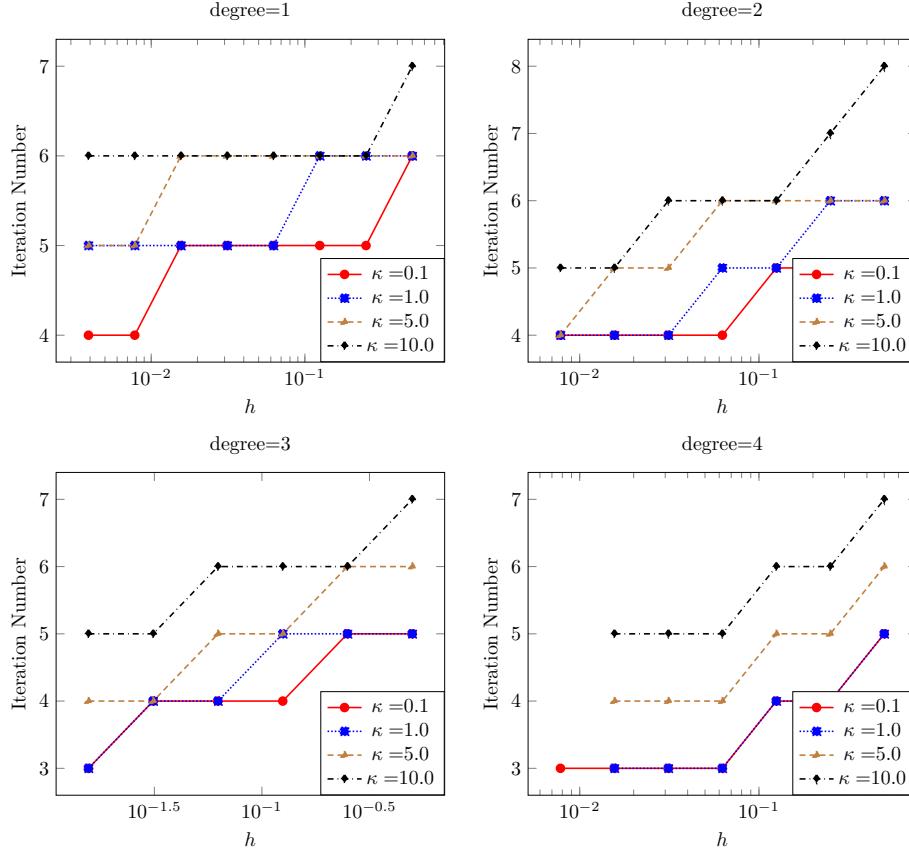


FIG. 10. GMRES iteration counts for a two-dimensional mesh using LU factorization for A^L as a preconditioner for various values of κ . Fixing κ and refining the mesh (right to left) leads to a slight decrease in iteration count, while fixing a mesh and increasing κ leads to a mild increase in iteration count.

iteration count. Comparing Figure 11 to Figure 10 suggests that the difference in iteration counts follows from the difficulty in obtaining an effective iterative method for the regular Helmholtz operator rather than new difficulties presented by our nonlocal boundary condition.

Finally, we devise an experiment to demonstrate our method's robustness with respect to the distance between the scatterer boundary Γ and the truncated boundary Σ . We continue to use the circle of radius 1 centered at the origin as the scatterer Γ . For side length $s \in \{2.25, 2.5, 3.0, 4.0, 5.0, 6.0\}$, we truncated the domain as the square $[-\frac{s}{2}, \frac{s}{2}]^2$ take out the circle of radius 1. In Figure 12, we measure the relative L^2 error of our nonlocal method for each domain. We see no pathologies emerging as the computational domain becomes smaller.

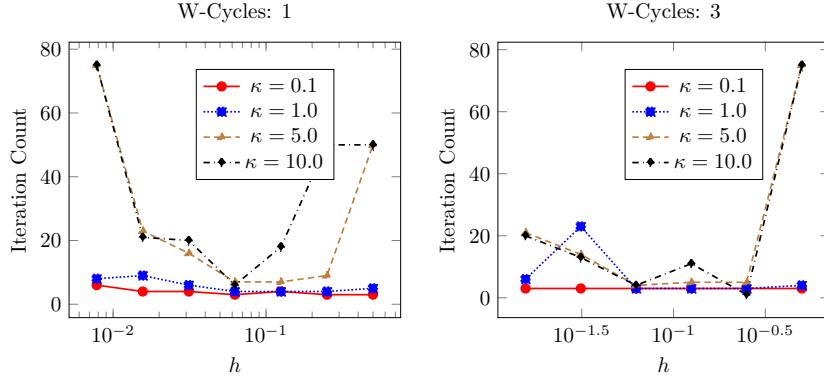


FIG. 11. Outer GMRES iteration count for various meshes and κ values. We apply PyAMG's smoothed aggregation augmented with plane waves to A^L as a preconditioner for the total system. At small wave numbers, we see relatively low iteration counts that are fairly flat under mesh refinement. As κ increases, however, far more outer iterations are required to obtain convergence.

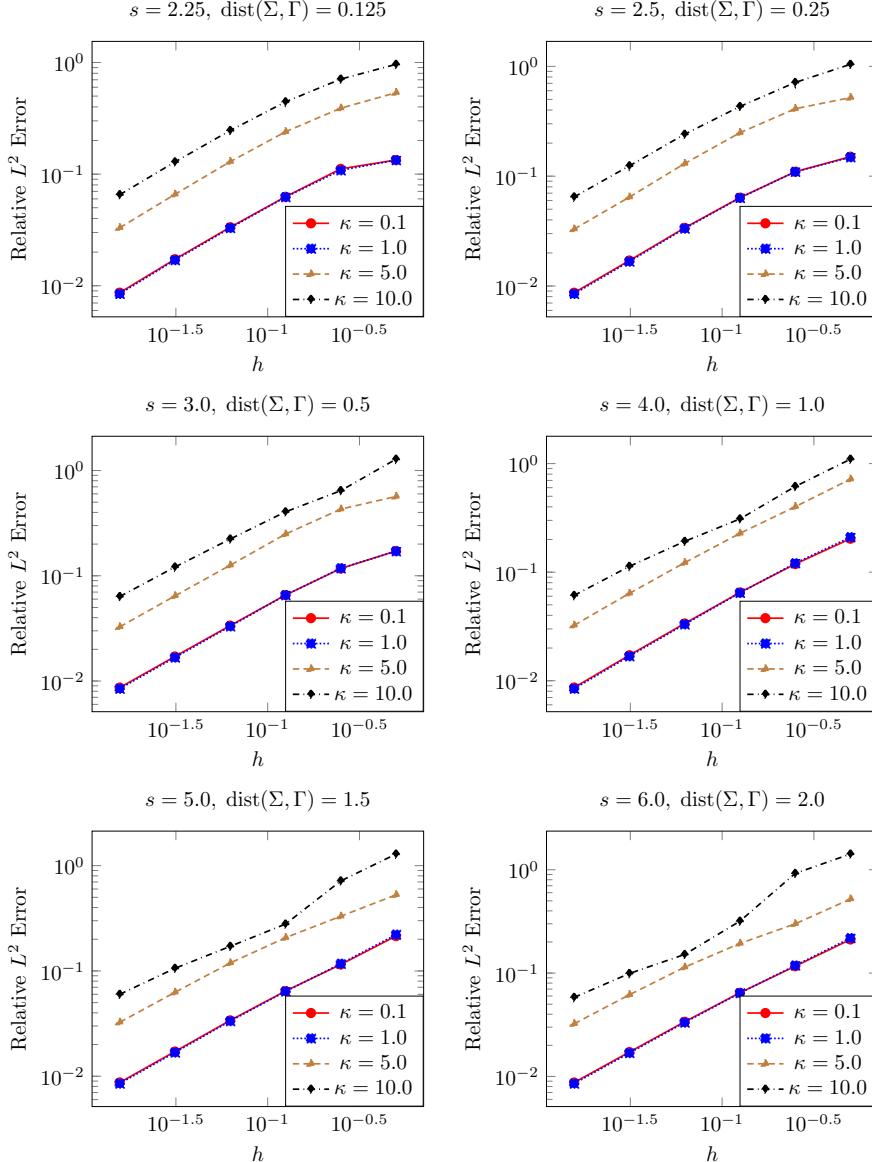


FIG. 12. Relative L^2 error of the nonlocal method using degree 1 polynomials as κ and the truncated domain size vary. A KSP relative tolerance of 10^{-12} was used.

6. Conclusions and future work. We have proposed a new nonlocal boundary condition for exterior Helmholtz problems. This condition, based on Green's formula and expressed in terms of layer potentials, works in general unstructured geometry in two and three dimensions. Thanks to a Gårding inequality, we have optimal finite element error estimates under standard conditions. The nonlocal terms are amenable to approximation by fast multipole expansions, and the discrete system can be readily preconditioned by its local part. In the future, it should be possible to extend the analysis to handle inexactness in evaluating the boundary terms. Moreover, we anticipate being able to apply this technique to a much broader class of problems such as exterior curl-curl problems. Additionally, we are working to integrate layer potentials with Firedrake's top-level language to make it easier to apply the method.

REFERENCES

- [1] M. F. ADAMS, *Algebraic multigrid methods for constrained linear systems with applications to contact problems in solid mechanics*, Numerical linear algebra with applications, 11 (2004), pp. 141–153.
- [2] M. S. ALNÆS, A. LOGG, K. B. ØLGAARD, M. E. ROGNES, AND G. N. WELLS, *Unified Form Language: a domain-specific language for weak formulations of partial differential equations*, ACM Transactions on Mathematical Software, 40 (2014), pp. 9:1–9:37, <https://doi.org/10.1145/2566630>, <https://arxiv.org/abs/1211.4047>.
- [3] J.-P. BERENGER, *A perfectly matched layer for the absorption of electromagnetic waves*, Journal of Computational Physics, 114 (1994), pp. 185–200.
- [4] A. BERMUDEZ, L. HERVELLA-NIETO, A. PRIETO, AND R. RODRIGUEZ, *An optimal finite-element/PML method for the simulation of acoustic wave propagation phenomena*, Variational Formulations in Mechanics: Theory and Applications, (2006).
- [5] A. BERMÚDEZ, L. HERVELLA-NIETO, A. PRIETO, AND R. RODRIGUEZ, *An optimal perfectly matched layer with unbounded absorbing function for time-harmonic acoustic scattering problems*, J. Comput. Phys., 223 (2007), pp. 469–488, <https://doi.org/10.1016/j.jcp.2006.09.018>, <https://doi.org/10.1016/j.jcp.2006.09.018>.
- [6] T. L. BINFORD, D. P. NICHOLLS, N. NIGAM, AND T. WARBURTON, *Exact non-reflecting boundary conditions on perturbed domains and hp-finite elements*, Journal of Scientific Computing, 39 (2009), pp. 265–292.
- [7] S. BRENNER AND R. SCOTT, *The mathematical theory of finite element methods*, vol. 15, Springer Science & Business Media, 2007.
- [8] J. CARRIER, L. GREENGARD, AND V. ROKHLIN, *A Fast Adaptive Multipole Algorithm for Particle Simulations*, SIAM Journal on Scientific and Statistical Computing, 9 (1988), pp. 669–686, <https://doi.org/10.1137/0909044>, <http://pubs.siam.org/doi/abs/10.1137/0909044> (accessed 2012-11-08).
- [9] D. COLTON AND R. KRESS, *Inverse Acoustic and Electromagnetic Scattering Theory*, Springer, 2nd ed., Jan. 1998.
- [10] L. D. DALCIN, R. R. PAZ, P. A. KLER, AND A. COSIMO, *Parallel distributed computing using Python*, Advances in Water Resources, 34 (2011), pp. 1124–1139, <https://doi.org/10.1016/j.advwatres.2011.04.013>. New Computational Methods and Software Tools.
- [11] V. DRUSKIN, S. GÜTTEL, AND L. KNIZHNERMAN, *Near-optimal perfectly matched layers for indefinite helmholtz problems*, SIAM Review, 58 (2016), pp. 90–116, <https://doi.org/10.1137/140966927>, <https://doi.org/10.1137/140966927>, <https://arxiv.org/abs/https://doi.org/10.1137/140966927>.
- [12] O. G. ERNST AND M. J. GANDER, *Why it is difficult to solve Helmholtz problems with classical iterative methods*, in Numerical analysis of multiscale problems, Springer, 2012, pp. 325–363.
- [13] M. J. GANDER, I. G. GRAHAM, AND E. A. SPENCE, *Applying GMRES to the Helmholtz equation with shifted Laplacian preconditioning: what is the largest shift for which wavenumber-independent convergence is guaranteed?*, Numerische Mathematik, (2015), pp. 1–48.
- [14] Z. GIMBUTAS AND L. GREENGARD, *A fast and stable method for rotating spherical harmonic expansions*, Journal of Computational Physics, 228 (2009), pp. 5621–5627, <https://doi.org/10.1016/j.jcp.2009.05.014>, <http://www.sciencedirect.com/science/article/pii/S0021999109002691> (accessed 2012-07-07).
- [15] D. GIVOLI, *Non-reflecting boundary conditions*, Journal of computational physics, 94 (1991),

pp. 1–29.

- [16] D. GIVOLI, *A spatially exact non-reflecting boundary condition for time dependent problems*, Computer methods in applied mechanics and engineering, 95 (1992), pp. 97–113.
- [17] D. GIVOLI AND D. COHEN, *Nonreflecting boundary conditions based on Kirchhoff-type formulae*, Journal of Computational Physics, 117 (1995), pp. 102–113.
- [18] N. GMATI AND B. PHILIPPE, *Comments on the GMRES convergence for preconditioned systems*, in International Conference on Large-Scale Scientific Computing, Springer, 2007, pp. 40–51.
- [19] C. I. GOLDSTEIN, *A finite element method for solving Helmholtz type equations in waveguides and other unbounded domains*, Mathematics of Computation, 39 (1982), pp. 309–324.
- [20] C. I. GOLDSTEIN, *The finite element method with non-uniform mesh sizes applied to the exterior Helmholtz problem*, Numerische Mathematik, 38 (1982), pp. 61–82.
- [21] P. GRISVARD, *Elliptic problems in nonsmooth domains*, SIAM, 2011.
- [22] M. J. GROTE AND J. B. KELLER, *Exact nonreflecting boundary conditions for the time dependent wave equation*, SIAM Journal on Applied Mathematics, 55 (1995), pp. 280–297.
- [23] T. HAGSTROM AND S. KIM, *Complete radiation boundary conditions for the helmholtz equation i: Waveguides*, Numer. Math., 141 (2019), pp. 917–966, <https://doi.org/10.1007/s00211-018-1012-0>, <https://doi.org/10.1007/s00211-018-1012-0>.
- [24] D. HOCH, *Nonreflecting boundary conditions obtained from equivalent sources for time-dependent scattering problems*, PhD thesis, California Institute of Technology, 2008.
- [25] H. HOU-DE AND W. XIAO-NAN, *Approximation of infinite boundary condition and its application to finite element methods*, Journal of Computational Mathematics, (1985), pp. 179–192.
- [26] V. E. HOWLE AND R. C. KIRBY, *Block preconditioners for finite element discretization of incompressible flow with thermal convection*, Numerical Linear Algebra with Applications, 19 (2012), pp. 427–440.
- [27] C. JOHNSON AND J. C. NEDELEC, *On the coupling of boundary integral and finite element methods*, Mathematics of Computation, (1980), pp. 1063–1079.
- [28] J. B. KELLER AND D. GIVOLI, *Exact non-reflecting boundary conditions*, Journal of Computational Physics, 82 (1989), pp. 172 – 192, [https://doi.org/https://doi.org/10.1016/0021-9991\(89\)90041-7](https://doi.org/https://doi.org/10.1016/0021-9991(89)90041-7), <http://www.sciencedirect.com/science/article/pii/0021999189900417>.
- [29] A. KLÖCKNER, A. BARNETT, L. GREENGARD, AND M. O’NEIL, *Quadrature by expansion: A new method for the evaluation of layer potentials*, Journal of Computational Physics, 252 (2013), pp. 332 – 349, <https://doi.org/10.1016/j.jcp.2013.06.027>.
- [30] A. KLÖCKNER ET AL., *meshmode Source Code Repository*, 2020, <https://github.com/inducer/meshmode>.
- [31] A. KLÖCKNER ET AL., *pyfmm Source Code Repository*, 2020, <https://github.com/inducer/pyfmm>.
- [32] A. KLÖCKNER ET AL., *pytential Source Code Repository*, 2020, <https://github.com/inducer/pytential>.
- [33] D. KOYAMA, *Error estimates of the DtN finite element method for the exterior Helmholtz problem*, Journal of Computational and Applied Mathematics, 200 (2007), pp. 21–31.
- [34] R. KRESS, *Linear Integral Equations*, Springer, Jan. 1999.
- [35] D. P. NICHOLLS AND N. NIGAM, *Exact non-reflecting boundary conditions on general domains*, Journal of Computational Physics, 194 (2004), pp. 278 – 303, <https://doi.org/https://doi.org/10.1016/j.jcp.2003.09.006>, <http://www.sciencedirect.com/science/article/pii/S0021999103004832>.
- [36] D. P. NICHOLLS AND N. NIGAM, *Error analysis of an enhanced DtN-FE method for exterior scattering problems*, Numerische Mathematik, 105 (2006), p. 267.
- [37] L. N. OLSON AND J. B. SCHRODER, *Smoothed aggregation for Helmholtz problems*, Numerical Linear Algebra with Applications, 17 (2010), pp. 361–386.
- [38] L. N. OLSON AND J. B. SCHRODER, *PyAMG: Algebraic multigrid solvers in Python v4.0*, 2018, <https://github.com/pyamg/pyamg>. Release 4.0.
- [39] M. RACHH, A. KLÖCKNER, AND M. O’NEIL, *Fast algorithms for Quadrature by Expansion I: Globally valid expansions*, Journal of Computational Physics, 345 (2017), pp. 706–731, <https://doi.org/10.1016/j.jcp.2017.04.062>, <http://www.sciencedirect.com/science/article/pii/S0021999117303418> (accessed 2020-09-10).
- [40] F. RATHGEBER, D. A. HAM, L. MITCHELL, M. LANGE, F. LUPORINI, A. T. T. MCRAE, G.-T. BERCEA, G. R. MARKALL, AND P. H. J. KELLY, *Firedrake: automating the finite element method by composing abstractions*, arXiv:1501.01809.
- [41] S. RIASANOW AND O. STEINBACH, *The fast solution of boundary integral equations*, Springer Science & Business Media, 2007.
- [42] M. E. ROGNES, R. C. KIRBY, AND A. LOGG, *Efficient assembly of $H(\text{div})$ and $H(\text{curl})$ con-*

forming finite elements, SIAM Journal on Scientific Computing, 31 (2009), pp. 4130–4151, <https://doi.org/10.1137/08073901X>.

[43] J. W. RUGE AND K. STÜBEN, *Algebraic multigrid*, in Multigrid methods, SIAM, 1987, pp. 73–130.

[44] A. SAFIN, S. MINKOFF, AND J. ZWECK, *A preconditioned finite element solution of the coupled pressure-temperature equations used to model trace gas sensors*, SIAM Journal on Scientific Computing, 40 (2018), pp. B1470–B1493.

[45] F.-J. SAYAS, *The validity of Johnson–Nédélec’s BEM–FEM coupling on polygonal interfaces*, SIAM Journal on Numerical Analysis, 47 (2009), pp. 3451–3463.

[46] O. STEINBACH, *Numerical approximation methods for elliptic boundary value problems: finite and boundary elements*, Springer Science & Business Media, 2007.

[47] L. TING AND M. J. MIKSIS, *Exact boundary conditions for scattering problems*, The Journal of the Acoustical Society of America, 80 (1986), pp. 1825–1827.

[48] B. VIOREANU AND V. ROKHLIN, *Spectra of Multiplication Operators as a Numerical Tool*, SIAM J. Sci. Comput., 36 (2014), pp. A267–A288, <https://doi.org/10.1137/110860082>, <https://pubs.siam.org/proxy2.library.illinois.edu/doi/abs/10.1137/110860082> (accessed 2018-05-09).

[49] M. WALA AND A. KLÖCKNER, *A fast algorithm with error bounds for Quadrature by Expansion*, Journal of Computational Physics, 374 (2018), pp. 135–162, <https://doi.org/10.1016/j.jcp.2018.05.006>, <http://www.sciencedirect.com/science/article/pii/S0021999118302985> (accessed 2020-09-10).

[50] M. WALA AND A. KLÖCKNER, *A fast algorithm for Quadrature by Expansion in three dimensions*, Journal of Computational Physics, 388 (2019), pp. 655–689, <https://doi.org/10.1016/j.jcp.2019.03.024>, <http://www.sciencedirect.com/science/article/pii/S0021999119302074> (accessed 2020-09-10).

[51] M. WALA AND A. KLÖCKNER, *On the Approximation of Local Expansions of Laplace Potentials by the Fast Multipole Method*, arXiv:2008.00653 [cs, math], (2020), <http://arxiv.org/abs/2008.00653> (accessed 2020-09-10). arXiv: 2008.00653.

[52] H. XIAO AND Z. GIMBUTAS, *A numerical algorithm for the construction of efficient quadrature rules in two and higher dimensions*, Computers & Mathematics with Applications, 59 (2010), pp. 663–676, <https://doi.org/10.1016/j.camwa.2009.10.027>.

h	κ	Iteration Count
0.5	0.1	88
0.25	0.1	293
0.125	0.1	675
0.5	1.0	89
0.25	1.0	283
0.125	1.0	519

TABLE 1

Number of iterations to convergence for the nonlocal method with no preconditioner, degree 1 polynomials, a KSP relative tolerance of 10^{-12} , and GMRES restart set to 200. Iteration counts were higher with a GMRES restart of 100, and much higher with the default GMRES restart of 30.