

A spectrally accurate direct solution technique for frequency-domain scattering problems with variable media

A. Gillman¹, A. H. Barnett¹, and P. G. Martinsson²

¹ Department of Mathematics, Dartmouth College ² Department of Applied Mathematics, University of Colorado at Boulder

Abstract: This paper presents a direct solution algorithm for the scattering of time-harmonic waves from a bounded region of the plane in which the wavenumber varies smoothly in space. The algorithm constructs the interior Dirichlet-to-Neumann (DtN) map for the bounded region via bottom-up recursive merges of (discretizations of) certain boundary operators on a quadtree of boxes. These operators take the form of impedance-to-impedance (ItI) maps. Since ItI maps are unitary, this formulation is inherently numerically stable, and is immune to problems of artificial internal resonances. The ItI maps on the smallest (leaf) boxes are built by spectral collocation on tensor-product grids of Chebyshev nodes. At the top level the DtN map is recovered from the ItI map and coupled to a boundary integral formulation of the free space exterior problem, to give a provably second kind equation. We show that the scheme can solve challenging problems 70 wavelengths on a side to 9-digit accuracy with 4 million unknowns, in under 5 minutes on a desktop workstation. Each additional solve corresponding to a different incident wave (right-hand side) then requires only 0.04 seconds.

1. INTRODUCTION

1.1. Problem formulation. Consider time-harmonic waves propagating in a medium where the wave speed varies smoothly, but is constant outside of a bounded domain $\Omega \subset \mathbb{R}^2$. This manuscript presents a technique for numerically solving the scattering problem in such a medium. Specifically, we seek to compute the scattered wave u^s that results when a given incident wave u^i (which satisfies the free space Helmholtz equation) impinges upon the region with variable wave speed, as in Figure 1. Mathematically, the scattered field u^s satisfies the variable coefficient Helmholtz equation

$$(1) \quad \Delta u^s(\mathbf{x}) + \kappa^2(1 - b(\mathbf{x}))u^s(\mathbf{x}) = \kappa^2 b(\mathbf{x})u^i(\mathbf{x}), \quad \mathbf{x} \in \mathbb{R}^2,$$

and the outgoing Sommerfeld radiation condition

$$(2) \quad \frac{\partial u^s}{\partial r} - i\kappa u^s = o(r^{-1/2}), \quad r := |\mathbf{x}| \rightarrow \infty,$$

uniformly in angle. The real number κ in (1) and (2) is the free space wavenumber (or frequency), and the so called “scattering potential” $b = b(\mathbf{x})$ is a given smooth function that specifies how the wave speed (phase velocity) $v(\mathbf{x})$ at the point $\mathbf{x} \in \mathbb{R}^2$ deviates from the free space wave speed v_{free} ,

$$(3) \quad b(\mathbf{x}) = 1 - \left(\frac{v_{\text{free}}}{v(\mathbf{x})} \right)^2.$$

One may interpret $\sqrt{1 - b}$ as a spatially-varying refractive index. Observe that b is identically zero outside Ω . Together, equations (1) and (2) completely specify the problem. For $1 - b$ everywhere real and positive, the problem is known to have a unique solution for each positive κ [9, Thm. 8.7].

This transmission problem (1)-(2), and its generalizations, has applications in acoustics, electromagnetics, optics, and quantum mechanics; in particular we might mention underwater acoustics [3], ultrasound and microwave tomography [13, 34], wave propagation in metamaterials and photonic crystals, and seismology [35]. The solution technique in this paper is high-order accurate, yet robust and computationally highly efficient. It is based on a direct (as opposed to iterative) solver, and thus is particularly effective when the response of a given potential b to multiple incident waves u^i

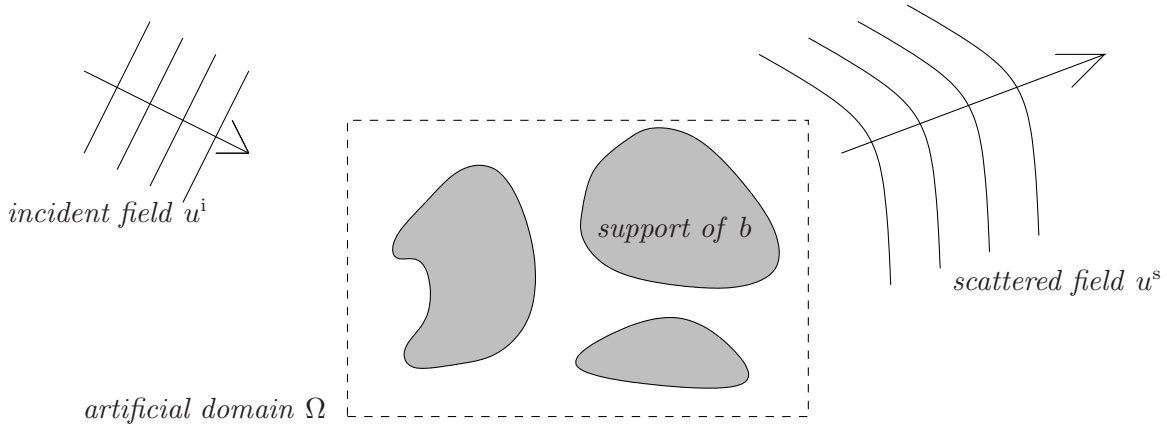


FIGURE 1. Geometry of scattering problem. Waves propagate in a medium with constant wave speed v_{free} everywhere except in the shaded areas, where the wave speed is given by $v(\mathbf{x}) = v_{\text{free}}/\sqrt{1-b(\mathbf{x})}$ where $b = b(\mathbf{x})$ is a given, smooth, compactly supported “scattering potential.” An incident field u^i hits the scattering potential and induces the scattered field u^s . The dashed line marks the artificial domain Ω which encloses the support of the scattering potential.

is desired, as arises in optical device characterization, or computing radar scattering cross-sections. The complexity of the method is $O(N^{3/2})$ where N is the number of discretization points in Ω . Additional solves with the same scattering potential b and wavenumber κ require only $O(N)$ operations. (Further reductions in asymptotic complexity can sometimes be attained; see section 4.2.) For simplicity of presentation, the solution technique is presented in \mathbb{R}^2 ; however, the method can be directly extended to \mathbb{R}^3 .

Remark 1.1. Equation (1) is derived by requiring that the total field $u = u^s + u^i$ should satisfy the variable coefficient Helmholtz equation

$$(4) \quad \Delta u(\mathbf{x}) + \kappa^2 \left(\frac{v_{\text{free}}}{v(\mathbf{x})} \right)^2 u(\mathbf{x}) = 0, \quad \mathbf{x} \in \mathbb{R}^2.$$

Now simply combine (4) with the condition that the incident field u^i satisfy the free space equation $(\Delta + \kappa^2)u^i = 0$ inside Ω , and the definition of the scattering potential (3), to obtain (1).

1.2. Outline of proposed method. We will solve (1)-(2) by splitting the problem into two parts, namely a variable-coefficient problem on the bounded domain Ω , and a constant coefficient problem on the exterior domain $\Omega^c := \mathbb{R}^2 \setminus \overline{\Omega}$. For each of the two domains, we will construct a solution operator in the form of a Dirichlet-to-Neumann (DtN) map on the boundary $\partial\Omega$, and then we will “glue” the two solution operators together on $\partial\Omega$ to form a solution operator for the full problem. The end result will be a discretized boundary integral operator that takes as input the restriction of the incoming field u^i (and its normal derivative) to $\partial\Omega$, and constructs the restriction of the scattered field u^s on $\partial\Omega$ (and its normal derivative). Once these quantities are known, the total field can rapidly be computed at any point $\mathbf{x} \in \mathbb{R}^2$.

1.2.1. Solution technique for the variable-coefficient problem on Ω . For the interior domain Ω , we will construct a solution operator for the homogeneous variable-coefficient equation that the total field $u = u^s + u^i$ must satisfy, namely,

$$(5) \quad \Delta u(\mathbf{x}) + \kappa^2(1-b(\mathbf{x}))u(\mathbf{x}) = 0, \quad \mathbf{x} \in \Omega,$$

$$(6) \quad u(\mathbf{x}) = h(\mathbf{x}) \quad \mathbf{x} \in \partial\Omega.$$

Note that, for now, we specify the Dirichlet data h ; when we consider the full problem h will become an unknown that will also be solved for. It is known that, for all but a countable set of wavenumbers, the interior Dirichlet BVP (5)-(6) has a unique solution u [25, Thm. 4.10]. The values $\{\kappa_j\}_{j=1}^{\infty}$ at which the solution is not unique are (the square roots of) the interior Dirichlet eigenvalues of the penetrable domain Ω ; we will call them *resonant wavenumbers*.

Definition 1.1 (interior Dirichlet-to-Neumann map). Suppose that $\kappa > 0$ is not a resonant wavenumber of Ω . Then the interior Dirichlet-to-Neumann (DtN) operator¹ $T_{\text{int}} : H^1(\partial\Omega) \rightarrow L^2(\partial\Omega)$ is defined by

$$(7) \quad T_{\text{int}}h = u_n,$$

where u the solution to (5)–(6) corresponding to Dirichlet data h , for any $h \in H^1(\partial\Omega)$.

Remark 1.2. The operator T_{int} is a pseudo-differential operator of order +1 [14]; that is, in the limit of high-frequency boundary data it behaves like a differentiation operator on $\partial\Omega$. The boundedness as a map $T_{\text{int}} : H^1(\partial\Omega) \rightarrow L^2(\partial\Omega)$ holds for Ω any bounded Lipschitz domain since the PDE is strongly elliptic [25, Thm. 4.25].

We will construct a discrete approximation to T_{int} via a variation of recent composite spectral methods in [15, 24]. These methods partition Ω into a collection of small “leaf” boxes and construct approximate DtN operators for each box via a brute force calculation on a local spectral grid. The DtN operator for Ω is then constructed via a hierarchical merge process. Unfortunately, at any given κ , each of the many leaves and merging subdomains may hit a resonance as described above, causing its DtN to fail to exist. As κ approaches any such resonance the norm of the DtN grows without bound. Thus a technique based on the DtN alone is not robust.

Remark 1.3. We remind the reader that any such “box” resonance is purely artificial and is caused by the introduction of the solution regions. It is important to distinguish these from resonances that the physical scattering problem (1)-(2) itself might possess (e.g. due to nearly trapped rays), whose effect of course cannot be avoided in *any* accurate numerical method.

One contribution of the present work is to present a robust improvement to the methods of [15, 24], built upon hierarchical merges of impedance-to-impedance (ItI) rather than DtN operators; see section 2. The idea of using impedance coupling builds upon the work of Kirsch–Monk [21]. ItI operators are inherently stable, with condition number $O(1)$, and thus exclude the possibility of inverting arbitrarily ill-conditioned matrices as in the original DtN formulation. For instance, for the *lens* experiment in section 5, the DtN method has condition numbers as large as 2×10^5 , while for the new ItI method the condition number never grows larger than 20. We do still need to handle the DtN of the whole domain Ω ; however, if Ω has a resonance, we may simply change its size slightly to avoid the resonance (see remark 2.6).

The discretization methods in [15, 24], and in this paper are related to earlier work on spectral collocation methods on composite (“multi-domain”) grids, such as, e.g., [22, 38], and in particular Pfeiffer et al [28]. For a detailed review of the similarities and differences, see [24].

1.2.2. *Solution technique for the constant coefficient problem on Ω^c .* On the exterior domain, we build a solution operator for the constant coefficient problem

$$(8) \quad \Delta u^s(\mathbf{x}) + \kappa^2 u^s(\mathbf{x}) = 0, \quad \mathbf{x} \in \Omega^c,$$

$$(9) \quad u^s(\mathbf{x}) = h(\mathbf{x}), \quad \mathbf{x} \in \partial\Omega,$$

$$(10) \quad \frac{\partial u^s}{\partial r} - i\kappa u^s = o(r^{-1/2}), \quad r := |\mathbf{x}| \rightarrow \infty,$$

¹this is also known as the Steklov–Poincaré operator [25].

obtained by restricting (1)-(2) to Ω^c . (Again, the Dirichlet data h later will become an unknown that is solved for.) It is known that (8)-(10) has a unique solution for every wavenumber κ [9, Ch. 3]. This means that the following DtN for the exterior domain is always well-defined.

Definition 1.2 (exterior Dirichlet-to-Neumann map). Suppose that $\kappa > 0$. The exterior DtN operator $T_{\text{ext}} : H^1(\partial\Omega) \rightarrow L^2(\partial\Omega)$ is defined by

$$(11) \quad T_{\text{ext}}h = u_n^s$$

for u^s the unique solution to the exterior Dirichlet BVP (8)-(10).

Numerically, we construct an approximation to T_{ext} by reformulating (8)-(10) as a boundary integral equation (BIE), as described in section 3.1, and then discretizing it using a Nyström method based on a high order Gaussian composite quadrature [18].

1.2.3. *Combining the two solution operators.* Once the DtN operators T_{int} and T_{ext} have been determined (as described in sections 1.2.1 and 1.2.2), and the restriction of the incident field to $\partial\Omega$ is given, it is possible to determine the scattered field on $\partial\Omega$ as follows. First observe that the total field $u = u^s + u^i$ must satisfy

$$(12) \quad T_{\text{int}}(u^i|_{\partial\Omega} + u^s|_{\partial\Omega}) = u_n^i + u_n^s.$$

We also know that the scattered field u^s must satisfy

$$(13) \quad T_{\text{ext}}u^s|_{\partial\Omega} = u_n^s.$$

Combining (12) and (13) to eliminate u_n^s , we obtain the equation (analogous to [21, Eq. (2.12)]),

$$(14) \quad (T_{\text{int}} - T_{\text{ext}})u^s|_{\partial\Omega} = u_n^i - T_{\text{int}}u^i|_{\partial\Omega}.$$

As discussed in Remark 1.2, both T_{int} and T_{ext} have order +1. Lamentably, this behavior *adds rather than cancels* in (14), so that $(T_{\text{int}} - T_{\text{ext}})$ also has order +1, and is therefore unbounded. This makes any numerical discretization of (14) ill-conditioned, with condition number growing linearly in the number of boundary nodes. To remedy this, we present in section 3 a new method for combining T_{int} and T_{ext} to give a provably second kind integral equation, which thus gives a well-conditioned linear system.

Once the scattered field is known on the boundary, the field at any exterior point may be found via Green's representation formula; see section 3.3. The interior transmitted wave u may be reconstructed anywhere in Ω by applying solution operators which were built as part of the composite spectral method.

1.3. **Prior work.** Perhaps the most common technique for solving the scattering problem stated in Section 1.1 is to discretize the variable coefficient PDE (1) via a finite element or finite difference method, while approximating the radiation condition in one of many ways, including perfectly matched layers (PML) [19], absorbing boundary conditions (ABC) [11], separation of variables or their perturbations [26], local impedance conditions [7], or a Nyström method [21] (as in the present work). However, the accuracy of finite element and finite difference schemes for the Helmholtz equation is limited by so-called ‘‘pollution’’ (dispersion) error [2, 4], demanding an increasing number of degrees of freedom per wavelength in order to maintain fixed accuracy as wavenumber κ grows. In addition, the resulting linear system will be large and sparse and is often ill-conditioned in such a way that standard pre-conditioning techniques fail, although hybrid direct-iterative solvers such as [12] have proven effective in certain environments. While there do exist fast direct solvers for such linear systems (for low wavenumbers κ) [31, 30, 36, 23], the accuracy of the solution is limited by the discretization. The solver's performance worsens when increasing the order of the discretization—and yet high order is precisely what would be required to overcome the above-mentioned pollution error.

Scattering problems on infinite domains are also commonly handled by rewriting them as volume integral equations (e.g. the Lippmann–Schwinger equation) defined on a domain (such as Ω) that contains the support of the scattering potential [1, 8]. This approach is appealing in that the Sommerfeld condition (2) is enforced analytically, and in that high-order discretizations can be implemented without loss of stability [10]. Principal drawbacks are that the resulting linear systems have dense coefficient matrices, and tend to be challenging to solve using iterative solvers [10].

1.4. Outline. Section 2 describes in detail the stable hierarchical procedure for constructing an approximation to the DtN map T_{int} for the interior problem (5)–(6). Section 3 describes how boundary integral equation techniques are used to approximate the DtN map T_{ext} for the exterior problem (8)–(10), how to couple the DtN maps T_{int} and T_{ext} to solve the full problem (1)–(2), and the proof (Theorem 3.1) that the formulation is second kind. Section 4 details the computational cost of the method and explains the reduced cost for multiple incident waves. Finally, section 5 illustrates the performance of the method in several challenging scattering potential configurations.

2. CONSTRUCTING AND MERGING IMPEDANCE-TO-IMPEDANCE MAPS

This section describes a technique for building a discrete approximation to the Dirichlet-to-Neumann (DtN) operator for the interior variable coefficient BVP (5)–(6) on a square domain Ω . It relies on the hierarchical construction of impedance-to-impedance (ItI) maps; these are defined in section 2.1. Section 2.2 defines a hierarchical tree on the domain Ω . Section 2.3 explains how the ItI maps are built on the (small) leaf boxes in the tree. Section 2.4 describes the merge procedure whereby the global ItI map is built, and then how the global DtN map is recovered from the global ItI map.

2.1. The impedance-to-impedance map. We start by defining the ItI map on a general Lipschitz domain, and giving some of its properties. (In this section only, Ω will refer to such a general domain.)

Proposition 2.1. Let $\Omega \subset \mathbb{R}^2$ be a bounded Lipschitz domain, and $b(\mathbf{x})$ be real. Let $\eta \in \mathbb{C}$, and $\text{Re } \eta \neq 0$. Then the interior Robin BVP,

$$(15) \quad [\Delta + \kappa^2(1 - b(\mathbf{x}))]u(\mathbf{x}) = 0 \quad \mathbf{x} \in \Omega ,$$

$$(16) \quad u_n + i\eta u|_{\partial\Omega} = f \quad \text{on } \partial\Omega ,$$

has a unique solution u for all real $\kappa > 0$.

Proof. We first prove uniqueness. Consider u a solution to the homogeneous problem $f \equiv 0$. Then using Green’s 1st identity and (15), (16),

$$-i\eta \int_{\partial\Omega} |u|^2 = \int_{\partial\Omega} \bar{u}u_n = \int_{\Omega} |\nabla u|^2 - \kappa^2(1 - b)|u|^2.$$

Taking the imaginary part shows that u , and hence u_n , vanishes on $\partial\Omega$, hence $u \equiv 0$ in Ω by unique continuation of the Cauchy data. Existence of $u \in H^1(\Omega)$ now follows for data $f \in H^{-1/2}(\partial\Omega)$ from the Fredholm alternative, as explained in this context by McLean [25, Thm. 4.11]. \square

Definition 2.1 (interior impedance-to-impedance map). Fix $\eta \in \mathbb{C}$, and $\text{Re } \eta \neq 0$. Let

$$(17) \quad f := u_n + i\eta u|_{\partial\Omega}$$

$$(18) \quad g := u_n - i\eta u|_{\partial\Omega}$$

be Robin traces of u . We will refer to f and g as the “incoming” and “outgoing” (respectively) impedance data. For any $\kappa > 0$, the *interior ItI operator* $R : L^2(\partial\Omega) \rightarrow L^2(\partial\Omega)$ is defined by

$$(19) \quad Rf = g$$

for f and g the Robin traces of u the solution of (15)–(16), for all $f \in L^2(\partial\Omega)$.




FIGURE 2. The square domain Ω shown split hierarchically into boxes. The case where there are 4×4 leaf boxes is shown, ie there are $M = 2$ levels. These are then gathered into a binary tree of successively larger boxes as described in Section 2. One possible enumeration of the boxes in the tree is shown, but note that the only restriction is that if box τ is the parent of box σ , then $\tau < \sigma$.

We shall choose the impedance parameter η (on dimensional grounds) to be $\eta = \kappa$. Numerically, in what follows, we observe very little sensitivity to the exact value or sign of η .

For the following, we will need the result that the DtN map T_{int} is self-adjoint for real κ and $b(\mathbf{x})$. This holds since, for any functions u and v satisfying $(\Delta + \kappa^2(1-b))u = 0$ in Ω and $(\Delta + \kappa(1-b))v = 0$ in Ω ,

$$0 = \int_{\Omega} \bar{v}(\Delta + \omega^2(1-b))u - u(\Delta + \overline{\kappa^2(1-b)})\bar{v} = \int_{\partial\Omega} \bar{v}u_n - u\bar{v}_n = (v|_{\partial\Omega}, T_{\text{int}}u|_{\partial\Omega}) - (T_{\text{int}}v|_{\partial\Omega}, u|_{\partial\Omega})$$

by Green's second identity. This allows us to prove the following property of the ItI map that will be the key to the numerical stability of the method.

Proposition 2.2. Let Ω be a bounded Lipschitz domain, let $b(\mathbf{x})$ be real, and let $\eta \in \mathbb{C}$ and $\text{Re } \eta \neq 0$. Then the ItI map R for Ω exists for all real frequencies κ , and is unitary whenever η is also real.

Proof. Existence of R for all real κ follows from Proposition 2.1. To prove R is unitary, we insert the definitions of f and g into (19) and use the definition of the DtN to rewrite $u_n = T_{\text{int}}u|_{\partial\Omega}$, giving

$$(20) \quad R(T_{\text{int}} + i\eta)u|_{\partial\Omega} = (T_{\text{int}} - i\eta)u|_{\partial\Omega} ,$$

which holds for any data $u|_{\partial\Omega} \in H^1(\partial\Omega)$. Thus we have as operators,

$$(21) \quad R = (T_{\text{int}} - i\eta)(T_{\text{int}} + i\eta)^{-1} .$$

Since T_{int} is self-adjoint and η is real, this formula shows that R is unitary. \square

As a unitary operator, R has unit operator L^2 -norm, pseudo-differential order 0, and eigenvalues lying on the unit circle. From (21) and the pseudo-differential order of T_{int} one may see that the eigenvalues of R accumulate only at $+1$.

2.2. Partitioning of Ω into hierarchical tree of boxes. Recall that Ω is the square domain containing the support of b . We partition Ω into a collection of 4^M equally-sized square boxes called *leaf boxes*, where M sets the number of levels; see Figure 2. We place q Gauss–Legendre interpolation nodes on each edge of each leaf, which will serve to discretize all interactions of this leaf with its neighbors; see Figure 3(a). The size of the leaf boxes, and the parameter q , should be chosen so that any potential transmitted wave u , as well as its first derivatives, can be accurately interpolated on each box edge from their values at these nodes.

Next we construct a binary tree over the collection of leaf boxes. This is achieved by forming the union of adjacent pairs boxes (forming rectangular boxes), then forming the pairwise union of the rectangular boxes. The result is a collection of squares with twice the side length of a leaf box. We continue the process until the only box is Ω itself, as in Figure 2. The boxes should be ordered so

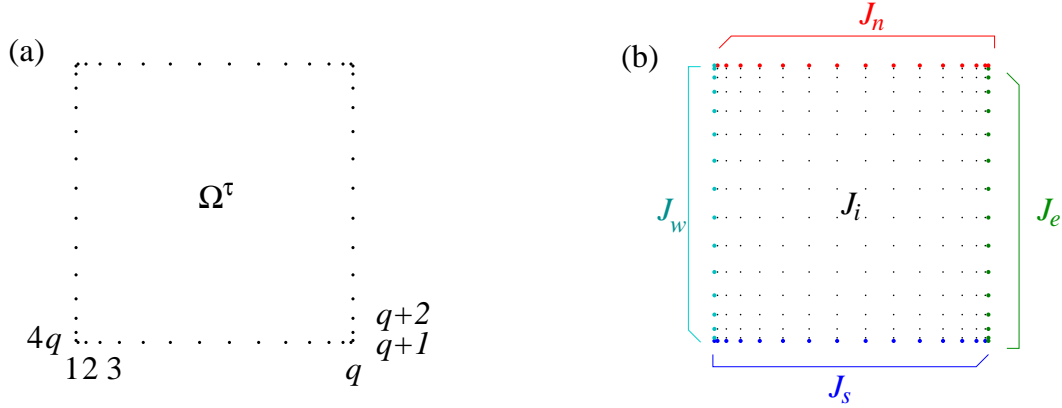


FIGURE 3. (a) The set of $4q$ Gauss–Legendre points used to represent the ItI on the boundary of a leaf box Ω^τ . The case $q = 14$ is shown. (b) Chebyshev discretization used for the PDE spectral solution on the same leaf box, for the case $p = 16$. Interior nodes J_i are shown by small black dots. The four sets of $p - 1$ Chebyshev boundary nodes are shown (each a different color); each set includes the start corner (in counter-clockwise ordering) but not the end corner. The set of all Chebyshev boundary nodes is $J_b = [J_s, J_e, J_n, J_w]$.

that if τ is a parent of a box σ , then $\tau < \sigma$. We also assume that the root of the tree (i.e. the full box Ω) has index $\tau = 1$. We let Ω^τ denote the domain associated with box τ .

Remark 2.1. The method we present easily generalizes to rectangular boxes, and to more complicated domains Ω in the same manner as [24].

2.3. Spectral approximation of the ItI map on a leaf box. Let Ω^τ denote a single leaf box, and let $\mathbf{f} = \mathbf{f}^\tau$ and $\mathbf{g} = \mathbf{g}^\tau$ be a pair of vectors of associated incoming and outgoing impedance data, sampled at the $4q$ Gauss–Legendre boundary nodes, with entries ordered in a counter-clockwise fashion starting from the leftmost node of the bottom edge of the box, as in Figure 3(a). In this section we describe how to construct a matrix approximation to the ItI operator on this leaf box. Namely, we build a $4q \times 4q$ matrix \mathbf{R} such that

$$\mathbf{g} \approx \mathbf{R}\mathbf{f}$$

holds to high-order accuracy, for all incoming data vectors $\mathbf{f} \in \mathbb{R}^{4q}$ corresponding to smooth transmitted wave solutions u .

We discretize the PDE (15) on the square leaf box Ω^τ using a spectral method on a $p \times p$ tensor product Chebyshev grid filling the box, as in Figure 3(b), comprised of the nodes at locations

$$\left(\frac{a+b}{2} + hx_i, \frac{c+d}{2} + hx_j \right), \quad i, j = 1, \dots, p$$

where $x_j := \cos \frac{\pi(j-1)}{p-1}$, $j = 1, \dots, p$ are the Chebyshev points on $[-1, 1]$. We label the Chebyshev node locations $\mathbf{x}_j \in \mathbb{R}^2$, for $j = 1, \dots, p^2$. For notational purposes, we order these nodes in the following fashion: the indices $J_b = \{1, 2, \dots, 4(p-1)\}$ correspond to the Chebyshev nodes lying on the boundary of Ω^τ , ordered counter-clockwise starting from the node located at the south-west corner (a, c) . The remaining $(p-2)^2$ interior nodes have indices $J_i = \{4(p-1) + 1, \dots, p^2\}$ and may be ordered arbitrarily (a Cartesian ordering is convenient).

Let $\mathbf{D}^{(1)}, \mathbf{D}^{(2)} \in \mathbb{R}^{p^2 \times p^2}$ be the standard spectral differentiation matrices constructed on the full set of Chebyshev nodes, which approximate the $\partial/\partial x_1$ (horizontal) and $\partial/\partial x_2$ (vertical) derivative

operators, respectively. As explained in [32, Ch. 7], these are constructed from Kronecker products of the $p \times p$ identity matrix and $h^{-1}\mathbf{D}$, where $\mathbf{D} \in \mathbb{R}^{p \times p}$ is the usual one-dimensional differentiation matrix on the nodes $\{x_i\}_{i=1}^p$. Namely \mathbf{D} has entries $\mathbf{D}_{ij} = \frac{w_j}{w_i(x_i - x_j)}$ where $\{w_j\}_{j=1}^p = [1/2, -1, 1, -1, \dots, (-1)^p, (-1)^{p-1}/2]$ is the vector of barycentric weights for the Chebyshev nodes (see [32, Ch. 6] and [29, Eqn.(8)].) Let the matrix $\mathbf{A} \in \mathbb{R}^{p^2 \times p^2}$ be the spectral discretization of the operator $\Delta + \kappa^2(1 - b(\mathbf{x}))$ on the product Chebyshev grid, namely

$$\mathbf{A} = (\mathbf{D}^{(1)})^2 + (\mathbf{D}^{(2)})^2 + \text{diag} \{ \kappa^2(1 - b(\mathbf{x}_j)) \}_{j=1}^{p^2},$$

where ‘‘diag S’’ indicates the diagonal matrix whose entries are the elements of the ordered set S .

Remark 2.2. The matrices $\mathbf{D}^{(1)}$, $\mathbf{D}^{(2)}$, and \mathbf{A} must have rows and columns ordered as explained above (i.e. boundary then interior) for the Chebyshev nodes; this requires permuting rows and columns of the matrices constructed by Kronecker products. For example, the structure of \mathbf{A} is

$$\begin{bmatrix} \mathbf{A}_{bb} & \mathbf{A}_{bi} \\ \mathbf{A}_{ib} & \mathbf{A}_{ii} \end{bmatrix},$$

where the notation \mathbf{A}_{bi} indicates the submatrix block $\mathbf{A}(J_b, J_i)$, etc.

We now break the $4(p-1)$ boundary Chebyshev nodes into four sets $J_b = [J_s, J_e, J_n, J_w]$, denoting the south, east, north, and west edges, as in Figure 3(b). Note that J_s includes the south-western corner $J_s(1)$ but not the south-eastern corner (which in turn is the first element of J_e), etc.

We are now ready to derive the linear system required for constructing the approximate ItI operator. We first build a matrix $\mathbf{N} \in \mathbb{R}^{4(p-1) \times p^2}$ which maps values of u at all Chebyshev nodes to the outgoing normal derivatives at the boundary Chebyshev nodes, as follows,

$$(22) \quad \mathbf{N} = \begin{bmatrix} -\mathbf{D}^{(2)}(J_s, :) \\ \mathbf{D}^{(1)}(J_e, :) \\ \mathbf{D}^{(2)}(J_n, :) \\ -\mathbf{D}^{(1)}(J_w, :) \end{bmatrix},$$

where (as is standard in MATLAB) the notation $\mathbf{A}(S, :)$ denotes the matrix formed from the subset of rows of a matrix \mathbf{A} given by the index set S . Then, recalling (16), the matrix $\mathbf{F} \in \mathbb{R}^{4(p-1) \times p^2}$ which maps the values of u at all Chebyshev nodes to incoming impedance data on the boundary Chebyshev nodes is

$$(23) \quad \mathbf{F} = \mathbf{N} + i\eta \mathbf{I}_{p^2}(J_b, :),$$

where \mathbf{I}_{p^2} denotes the identity matrix of size p^2 . Using $\mathbf{u} \in \mathbb{R}^{p^2}$ for the vector of u values at all Chebyshev nodes, the linear system for the unknown \mathbf{u} that imposes the spectral discretization of the PDE at all interior nodes, and incoming impedance data $\mathbf{f}_c \in \mathbb{R}^{4(p-1)}$ at the boundary Chebyshev nodes, is

$$(24) \quad \mathbf{B}\mathbf{u} = \begin{bmatrix} \mathbf{f}_c \\ \mathbf{0} \end{bmatrix}$$

where $\mathbf{0}$ is an appropriate column vector of zeros, and the square size- p^2 system matrix is

$$\mathbf{B} := \begin{bmatrix} \mathbf{F} \\ \mathbf{A}(J_i, :) \end{bmatrix}.$$

Remark 2.3. At each of the four corner nodes, only one boundary condition is imposed, namely the one associated with the edge lying in the counter-clockwise direction. This results in a square linear system, which we observe is around twice as fast to solve as a similar-sized rectangular one.

To construct the $p^2 \times 4(p-1)$ “solution matrix” \mathbf{X} for the linear system, we solve (24) for each unit vector in $\mathbb{R}^{4(p-1)}$, namely

$$\mathbf{B}\mathbf{X} = \begin{bmatrix} \mathbf{I}_{4p-4} \\ \mathbf{0}_{(p-2)^2 \times (4p-4)} \end{bmatrix}.$$

In practice, \mathbf{X} is found using the backwards-stable solver available via MATLAB’s `mldivide` command. If desired, the tabulated solution \mathbf{u} can now be found at all the Chebyshev nodes by applying \mathbf{X} to the right hand side of (24).

Recall that the goal is to construct matrices that act on boundary data on Gauss (as opposed to Chebyshev) nodes. With this in mind, let \mathbf{P} be the matrix which performs Lagrange polynomial interpolation [33, Ch. 12] from q Gauss to p Chebyshev points on a single edge, and let \mathbf{Q} be the matrix from Chebyshev to Gauss points. Let $\mathbf{P}_0 \in \mathbb{R}^{(p-1) \times q}$ be \mathbf{P} with the last row omitted. For example, \mathbf{P}_0 maps from Gauss points on the south edge to the Chebyshev points J_s .

Then the solution matrix which takes incoming impedance data on Gauss nodes to the values \mathbf{u} at all Chebyshev nodes is

$$(25) \quad \mathbf{Y} = \mathbf{X} \begin{bmatrix} \mathbf{P}_0 & & & \\ & \mathbf{P}_0 & & \\ & & \mathbf{P}_0 & \\ & & & \mathbf{P}_0 \end{bmatrix}.$$

Finally, we must extract outgoing impedance data on Gauss nodes from the vector \mathbf{u} , to construct an approximation \mathbf{R} to the full ItI map on the Gauss nodes. This is done by extracting (as in (22)) the relevant rows of the spectral differentiation matrices, then interpolating back to Gauss points. Let $J'_s := [J_s, J_e(1)]$ be the indices of all p Chebyshev nodes on the south edge, and correspondingly for the other three edges. Then the index set $J'_b := [J'_s, J'_e, J'_n, J'_w]$ counts each corner *twice*.² Then let $\mathbf{G} \in \mathbb{R}^{4p \times p^2}$ be the matrix mapping values of \mathbf{u} to the outgoing impedance data with respect to each edge, given by

$$\mathbf{G} = \begin{bmatrix} -\mathbf{D}^{(2)}(J'_s, :) \\ \mathbf{D}^{(1)}(J'_e, :) \\ \mathbf{D}^{(2)}(J'_n, :) \\ -\mathbf{D}^{(1)}(J'_w, :) \end{bmatrix} - i\eta \mathbf{I}_{p^2}(J'_b, :).$$

Then, in terms of (25),

$$\mathbf{R} = \begin{bmatrix} \mathbf{Q} & & & \\ & \mathbf{Q} & & \\ & & \mathbf{Q} & \\ & & & \mathbf{Q} \end{bmatrix} \mathbf{G}\mathbf{Y}$$

is the desired spectral approximation to the ItI map on the leaf box.

The computation time is dominated by the solution step for \mathbf{X} , which takes effort $O(p^6)$. We observe empirically that one must choose $p > q + 1$ in order that \mathbf{R} not acquire a spurious numerical null space. We typically choose $q = 14$ and $p = 16$.

2.4. Merging ItI maps. Once the approximate ItI maps are constructed on the boundary Gauss nodes on the leaf boxes, the ItI map defined on Ω is constructed by merging two boxes at a time, moving up the binary tree as described in section 2.2. In this section, with τ labeling a box in the tree with children α and β , we first demonstrate the purely local construction of the ItI operator \mathbf{R}^τ from \mathbf{R}^α and \mathbf{R}^β .

²Including both endpoints will allow more accurate interpolation back to Gauss nodes; functions on each edge are available at all Chebyshev nodes for that edge.

We begin by introducing some notation. Let Ω^τ denote a box with children Ω^α and Ω^β where $\Omega^\tau = \Omega^\alpha \cup \Omega^\beta$. For concreteness, consider the case where Ω^α and Ω^β share a vertical edge. As shown in Figure 4, we partition the Gauss points on $\partial\Omega^\alpha$ and $\partial\Omega^\beta$ into three sets:

- J_1 : Boundary nodes of Ω^α that are not boundary nodes of Ω^β .
- J_2 : Boundary nodes of Ω^β that are not boundary nodes of Ω^α .
- J_3 : Boundary nodes of both Ω^α and Ω^β that are *not* boundary nodes of the union box Ω^τ .

We now define interior and exterior outgoing data via

$$\mathbf{g}^i = \mathbf{g}_3^\alpha \quad \text{and} \quad \mathbf{g}^e = \begin{bmatrix} \mathbf{g}_1^\alpha \\ \mathbf{g}_2^\beta \end{bmatrix}.$$

The incoming vectors \mathbf{f}^i and \mathbf{f}^e are defined similarly. Our goal is to obtain an equation mapping \mathbf{f}^e to \mathbf{g}^e . Since we already have ItI operators for α and β , we have the following two equations

$$(26) \quad \begin{bmatrix} \mathbf{R}_{11}^\alpha & \mathbf{R}_{13}^\alpha \\ \mathbf{R}_{31}^\alpha & \mathbf{R}_{33}^\alpha \end{bmatrix} \begin{bmatrix} \mathbf{f}_1^\alpha \\ \mathbf{f}_3^\alpha \end{bmatrix} = \begin{bmatrix} \mathbf{g}_1^\alpha \\ \mathbf{g}_3^\alpha \end{bmatrix}; \quad \begin{bmatrix} \mathbf{R}_{22}^\beta & \mathbf{R}_{23}^\beta \\ \mathbf{R}_{32}^\beta & \mathbf{R}_{33}^\beta \end{bmatrix} \begin{bmatrix} \mathbf{f}_2^\beta \\ \mathbf{f}_3^\beta \end{bmatrix} = \begin{bmatrix} \mathbf{g}_1^\beta \\ \mathbf{g}_3^\beta \end{bmatrix}.$$

Since the normals of the two leaf boxes are opposed on the interior ‘‘3’’ edge, from (17)-(18) we have $\mathbf{g}_3^\alpha = -\mathbf{f}_3^\beta$ and $\mathbf{f}_3^\alpha = -\mathbf{g}_3^\beta$. These allow us to rewrite the bottom row equations using only α on the interior edge, namely

$$(27) \quad \mathbf{R}_{31}^\alpha \mathbf{f}_1^\alpha + \mathbf{R}_{33}^\alpha \mathbf{f}_3^\alpha = \mathbf{g}_3^\alpha$$

and

$$(28) \quad \mathbf{R}_{32}^\beta \mathbf{f}_2^\beta - \mathbf{R}_{33}^\beta \mathbf{g}_3^\alpha = -\mathbf{f}_3^\alpha.$$

Let $\mathbf{W} := (\mathbf{I} - \mathbf{R}_{33}^\beta \mathbf{R}_{33}^\alpha)^{-1}$. Plugging (27) into (28) results in the following equation

$$\mathbf{R}_{32}^\beta \mathbf{f}_2^\beta - \mathbf{R}_{33}^\beta (\mathbf{R}_{31}^\alpha \mathbf{f}_1^\alpha + \mathbf{R}_{33}^\alpha \mathbf{f}_3^\alpha) = -\mathbf{f}_3^\alpha.$$

By collecting like terms and solving for \mathbf{f}_3^α , we find

$$(29) \quad \begin{aligned} \mathbf{f}_3^\alpha &= \mathbf{W} \left(\mathbf{R}_{33}^\beta \mathbf{R}_{31}^\alpha \mathbf{f}_1^\alpha - \mathbf{R}_{32}^\beta \mathbf{f}_2^\beta \right) \\ &= \begin{bmatrix} \mathbf{W} \mathbf{R}_{33}^\beta \mathbf{R}_{31}^\alpha & -\mathbf{W} \mathbf{R}_{32}^\beta \end{bmatrix} \begin{bmatrix} \mathbf{f}_1^\alpha \\ \mathbf{f}_2^\beta \end{bmatrix}. \end{aligned}$$

Note that the matrix $\mathbf{S}^\alpha := \begin{bmatrix} \mathbf{W} \mathbf{R}_{33}^\beta \mathbf{R}_{31}^\alpha & -\mathbf{W} \mathbf{R}_{32}^\beta \end{bmatrix}$ maps the incoming impedance data on τ to the incoming (with respect to α) impedance data on the interior edge. By combining the relationship between the impedance boundary data on neighbor boxes and (27), the matrix $\mathbf{S}^\beta = -\begin{bmatrix} \mathbf{R}_{33}^\alpha + \mathbf{W} \mathbf{R}_{33}^\beta \mathbf{R}_{31}^\alpha & -\mathbf{W} \mathbf{R}_{32}^\beta \end{bmatrix}$ computes the impedance data \mathbf{f}_3^β .

The outgoing impedance data \mathbf{g}_3^α is found by plugging \mathbf{f}_3^α into equation (27). Now the top row equations of (26) can be rewritten without reference to the interior edge. The top row equation of (26) from Ω^α is now

$$\left(\mathbf{R}_{11}^\alpha + \mathbf{R}_{13}^\alpha \mathbf{W} \mathbf{R}_{33}^\beta \mathbf{R}_{31}^\alpha \right) \mathbf{f}_1^\alpha - \mathbf{R}_{13}^\alpha \mathbf{Q} \mathbf{R}_{32}^\beta \mathbf{f}_2^\beta = \mathbf{g}_1^\alpha$$

and the top row equation of (26) from Ω^β is

$$\left(\mathbf{R}_{22}^\beta + \mathbf{R}_{23}^\beta \mathbf{R}_{33}^\alpha \mathbf{W} \mathbf{R}_{32}^\beta \right) \mathbf{f}_2^\beta - \mathbf{R}_{23}^\beta \left(\mathbf{R}_{31}^\alpha + \mathbf{R}_{33}^\alpha \mathbf{W} \mathbf{R}_{33}^\beta \mathbf{R}_{31}^\alpha \right) \mathbf{f}_1^\alpha = \mathbf{g}_2^\beta.$$

Writing these equations as a system, we find \mathbf{f}_1^α and \mathbf{f}_2^β satisfy

$$(30) \quad \begin{bmatrix} \mathbf{R}_{11}^\alpha + \mathbf{R}_{13}^\alpha \mathbf{W} \mathbf{R}_{33}^\beta \mathbf{R}_{31}^\alpha & -\mathbf{R}_{13}^\alpha \mathbf{W} \mathbf{R}_{32}^\beta \\ -\mathbf{R}_{23}^\beta \left(\mathbf{R}_{31}^\alpha + \mathbf{R}_{33}^\alpha \mathbf{W} \mathbf{R}_{33}^\beta \mathbf{R}_{31}^\alpha \right) & \mathbf{R}_{22}^\beta + \mathbf{R}_{23}^\beta \mathbf{R}_{33}^\alpha \mathbf{Q} \mathbf{R}_{32}^\beta \end{bmatrix} \begin{bmatrix} \mathbf{f}_1^\alpha \\ \mathbf{f}_2^\beta \end{bmatrix} = \begin{bmatrix} \mathbf{g}_1^\alpha \\ \mathbf{g}_2^\beta \end{bmatrix}.$$




FIGURE 4. Notation for the merge operation described in Section 2.4. The rectangular domain Ω^τ is the union of two squares Ω^α and Ω^β . The sets J_1 and J_2 form the exterior Gauss nodes (black), while J_3 consists of the interior Gauss nodes (white). (An unrealistically small node number $q = 7$ is used for visual clarity.)

Thus the block matrix on the left hand side of (30) is \mathbf{R}^τ , the ItI operator for τ .

Remark 2.4. In practice, the matrix products $\mathbf{WR}_{33}^\beta \mathbf{R}_{31}^\alpha$ and \mathbf{WR}_{32}^β should be computed once per merge.

Remark 2.5. We note that the formula (30) is quite different from that for merging DtN maps appearing in prior work [15, 24]. The root cause is the way the equivalence of incoming and outgoing data on the interior edge differs from the case of Dirichlet and Neumann data.

Algorithm 1 outlines the implementation of the hierarchical construction of the impedance operators, by repeated application of the above merge operation. Algorithm 2 illustrates the downwards sweep to construct from incoming impedance data f the solution at all Chebyshev discretization points in Ω . Note that the latter requires the solution matrices \mathbf{S} at each level, and \mathbf{Y} for each of the leaf boxes, that were precomputed by Algorithm 1.

The resulting approximation to the top-level ItI operator $\mathbf{R} = \mathbf{R}^1$ is a square matrix which acts on incoming impedance data living on the total of $4q2^M$ composite Gauss nodes on $\partial\Omega$. An approximation \mathbf{T}_{int} to the interior DtN map on these same nodes now comes from inverting equation (20) for T_{int} , to give

$$(31) \quad \mathbf{T}_{\text{int}} = -i\eta (\mathbf{R} - \mathbf{I})^{-1} (\mathbf{R} + \mathbf{I}) ,$$

where \mathbf{I} is the identity matrix of size $4q2^M$. The need for conversion from the ItI to the DtN for the domain Ω will become clear in the next section.

Remark 2.6. Due to the pseudo-differential order $+1$ of T_{int} , we expect the norm of \mathbf{T}_{int} to grow linearly in the number of boundary nodes. However, it is also possible that κ falls close enough to a resonant wavenumber of Ω that the norm of \mathbf{T}_{int} is actually much larger, resulting in a loss of accuracy due to the inversion of the nearly-singular matrix $\mathbf{R} - \mathbf{I}$ in (31). In our extensive numerical experiments, this latter problem has never happened. However, it is important to include a condition number check when formula (31) is evaluated. Should there be a problem, it would be a simple matter to modify slightly the domain to avoid a resonance. For instance, one can add a column of leaf boxes to the side of Ω , and then inexpensively update the computed ItI operator for Ω to the enlarged domain.

ALGORITHM 1 (build solution matrices)

This algorithm builds the global Dirichlet-to-Neumann matrix \mathbf{T}_{int} for (5)-(6). For τ a leaf box, the algorithm builds the solution matrix \mathbf{Y}^τ that maps impedance data at Gauss nodes to the solution at the interior Chebyshev nodes. For non-leaf boxes τ , it builds the matrices \mathbf{S}^τ required for constructing \mathbf{f} impedance data on interior Gauss nodes. It is assumed that if box τ is a parent of box σ , then $\tau < \sigma$.

-
- (1) **for** $\tau = N_{\text{boxes}}, N_{\text{boxes}} - 1, N_{\text{boxes}} - 2, \dots, 1$
 - (2) **if** (τ is a leaf)
 - (3) Construct \mathbf{R}^τ and \mathbf{Y}^τ via the process described in Section 2.3.
 - (4) **else**
 - (5) Let α and β be the children of τ .
 - (6) Split J_b^α and J_b^β into vectors J_1, J_2 , and J_3 as shown in Figure 4.
 - (7) $\mathbf{W} = \left(\mathbf{I} - \mathbf{R}_{33}^\beta \mathbf{R}_{33}^\alpha \right)^{-1}$
 - (8)
$$\mathbf{R}^\tau = \begin{bmatrix} \mathbf{R}_{11}^\alpha + \mathbf{R}_{13}^\alpha \mathbf{W} \mathbf{R}_{33}^\beta \mathbf{R}_{31}^\alpha & -\mathbf{R}_{13}^\alpha \mathbf{W} \mathbf{R}_{32}^\beta \\ -\mathbf{R}_{23}^\beta (\mathbf{R}_{31}^\alpha + \mathbf{R}_{33}^\alpha \mathbf{W} \mathbf{R}_{33}^\beta \mathbf{R}_{31}^\alpha) & \mathbf{R}_{22}^\beta + \mathbf{R}_{23}^\beta \mathbf{R}_{33}^\alpha \mathbf{W} \mathbf{R}_{32}^\beta \end{bmatrix}$$
 - (9) $\mathbf{S}^\alpha = \begin{bmatrix} \mathbf{W} \mathbf{R}_{33}^\beta \mathbf{R}_{31}^\alpha & -\mathbf{W} \mathbf{R}_{32}^\beta \end{bmatrix}$.
 - (10) $\mathbf{S}^\beta = - \begin{bmatrix} \mathbf{R}_{33}^\alpha + \mathbf{W} \mathbf{R}_{33}^\beta \mathbf{R}_{33}^\alpha & -\mathbf{W} \mathbf{R}_{32}^\beta \end{bmatrix}$.
 - (10) Delete \mathbf{R}^α and \mathbf{R}^β .
 - (11) **end if**
 - (12) **end for**
 - (13) $\mathbf{T}_{\text{int}} = -i\eta(\mathbf{R}^1 - \mathbf{I})^{-1}(\mathbf{R}^1 + \mathbf{I})$

ALGORITHM 2 (solve BVP (15)-(16) once solution matrices have been built)

This program constructs an approximation \mathbf{u} to the solution u of (15)-(16). It assumes that all matrices $\mathbf{S}^\tau, \mathbf{Y}^\tau$ have already been constructed. It is assumed that if box τ is a parent of box σ , then $\tau < \sigma$. The indices of nodes in box τ we call J^τ .

-
- (1) **for** $\tau = 1, 2, 3, \dots, N_{\text{boxes}}$
 - (2) **if** (τ is a leaf)
 - (3) $\mathbf{u}(J^\tau) = \mathbf{Y}^\tau \mathbf{f}^\tau$.
 - (4) **else**
 - (5) Let α and β be the children of τ .
 - (6) $\mathbf{f}_3^\alpha = \mathbf{S}^\alpha \mathbf{f}^\tau, \mathbf{f}_3^\beta = \mathbf{S}^\beta \mathbf{f}^\tau$.
 - (7) **end if**
 - (8) **end for**

3. WELL-CONDITIONED BOUNDARY INTEGRAL FORMULATION FOR SCATTERING

In this section, we present an improved boundary integral equation alternative to the scattering formulation (14) from the introduction.

3.1. Formula for the exterior DtN operator T_{ext} . We first construct the exterior DtN operator T_{ext} via potential theory. The starting point is Green's exterior representation formula [9, Thm. 2.5],

which states that any radiative solution u^s to the Helmholtz equation in Ω^c may be written,

$$(32) \quad u^s(\mathbf{x}) = (\mathcal{D}u^s|_{\partial\Omega})(\mathbf{x}) - (\mathcal{S}u_n^s)(\mathbf{x}), \quad \text{for } \mathbf{x} \in \Omega^c,$$

where $(\mathcal{D}\phi)(\mathbf{x}) := \int_{\partial\Omega} \frac{\partial}{\partial n_{\mathbf{y}}} \left(\frac{i}{4} H_0^{(1)}(\kappa|\mathbf{x}-\mathbf{y}|) \right) \phi(\mathbf{y}) ds_{\mathbf{y}}$ and $(\mathcal{S}\phi)(\mathbf{x}) := \int_{\partial\Omega} \frac{i}{4} H_0^{(1)}(\kappa|\mathbf{x}-\mathbf{y}|) \phi(\mathbf{y}) ds_{\mathbf{y}}$ are respectively the frequency- κ Helmholtz double- and single-layer potentials with density ϕ , with $H_0^{(1)}$ the outgoing Hankel function of order zero. $n_{\mathbf{y}}$ denotes the outward unit normal vector at $\mathbf{y} \in \partial\Omega$. Letting \mathbf{x} approach $\partial\Omega$ in (32), one finds via the jump relations,

$$(33) \quad \frac{1}{2}u^s(\mathbf{x}) = (Du^s|_{\partial\Omega})(\mathbf{x}) - (Su_n^s)(\mathbf{x}), \quad \mathbf{x} \in \partial\Omega,$$

where D and S are the double- and single-layer boundary integral operators on $\partial\Omega$. See [9, Ch. 3.1] for an introduction to these representations and operators. Rearranging (33) gives $u_n^s = S^{-1} \left(D - \frac{1}{2}I \right) u^s|_{\partial\Omega}$, thus the exterior DtN operator is given in terms of the operators of potential theory by

$$(34) \quad T_{\text{ext}} = S^{-1} \left(D - \frac{1}{2}I \right) .$$

3.2. The new integral formulation. We apply from the left the single layer integral operator S to both sides of (14), and use (34), to obtain

$$(35) \quad \left(\frac{1}{2}I - D + ST_{\text{int}} \right) u^s|_{\partial\Omega} = S \left(u_n^i - T_{\text{int}}u^i|_{\partial\Omega} \right) ,$$

a linear equation for $u^s|_{\partial\Omega}$, the restriction of the scattered wave to the domain boundary $\partial\Omega$. Let

$$(36) \quad A := \frac{1}{2}I - D + ST_{\text{int}}$$

be the boundary integral operator appearing in the above formulation. In the trivial case $b \equiv 0$ (no scattering potential) it is easy to check that $A = I$, by using $T_{\text{int}} = S^{-1}(D + \frac{1}{2}I)$ which can be derived in this case similarly to (34). However, we now prove that introducing a general scattering potential perturbs A only compactly, that is, our left-regularization of the original ill-conditioned (14) has produced a well-conditioned equation.

Theorem 3.1. *Let $\Omega \in \mathbb{R}^2$ be a bounded Lipschitz domain containing the support of a bounded scattering potential function b . Let $\kappa > 0$ not be a resonant wavenumber of Ω . Then the operator (36) takes the form*

$$A = I + K$$

where $K : L^2(\partial\Omega) \rightarrow L^2(\partial\Omega)$ is compact; thus the formulation (35) is of Fredholm second kind.

Proof. Let u satisfy (5) in Ω , then by Green's interior representation formula (third identity) [9, Eq. (2.4)],

$$(37) \quad u(\mathbf{x}) = (\mathcal{S}u_n)(\mathbf{x}) - (\mathcal{D}u|_{\partial\Omega})(\mathbf{x}) - \kappa^2(\mathcal{V}bu)(\mathbf{x}) \quad \text{for } \mathbf{x} \in \Omega ,$$

where $(\mathcal{V}\phi)(\mathbf{x}) := \int_{\Omega} \frac{i}{4} H_0^{(1)}(\kappa|\mathbf{x}-\mathbf{y}|) \phi(\mathbf{y}) d\mathbf{y}$ denotes the Helmholtz volume potential [9, Sec. 8.2] acting on a function ϕ with support in Ω . By defining P to be the solution operator for the interior Dirichlet problem (5)-(6), i.e. $u(\mathbf{y}) = (Pu|_{\partial\Omega})(\mathbf{y})$ for $\mathbf{y} \in \Omega$, and B the operator that multiplies a function pointwise by $b(\mathbf{y})$, we can write the last term in (37) as $-\kappa^2\mathcal{V}BPu|_{\partial\Omega}$. Taking \mathbf{x} to $\partial\Omega$ from inside in (37), the jump relations give

$$\frac{1}{2}u(\mathbf{x}) = (Su_n)(\mathbf{x}) - (Du|_{\partial\Omega})(\mathbf{x}) - \kappa^2(\mathcal{V}BPu|_{\partial\Omega})(\mathbf{x}), \quad \mathbf{x} \in \partial\Omega ,$$

where V is \mathcal{V} restricted to evaluation on $\partial\Omega$. We substitute $u_n = T_{\text{int}}u|_{\partial\Omega}$, and since the above must hold for all data $u|_{\partial\Omega}$, we have as operators

$$ST_{\text{int}} = \frac{1}{2}I + D + \kappa^2\mathcal{V}BP$$

which, substituting into (36) results in cancellation of the D terms, giving

$$A = I + \kappa^2\mathcal{V}BP .$$




FIGURE 5. Eigenvalues of the discretized operators plotted in the complex plane. (a) shows the operator $T_{\text{int}} - T_{\text{ext}}$ from (14), whilst (b) shows the new regularized operator A from (35). In both cases $\Omega = (-0.5, 0.5)^2$, $\kappa = 20$ and $b(\mathbf{x}) = -1.5e^{-160|\mathbf{x}|^2}$.

Now $P : L^2(\partial\Omega) \rightarrow L^2(\Omega)$ is bounded [25, Thm. 4.25], and B is bounded. \mathcal{V} is two orders of smoothing: it is bounded from $L^2(\Omega)$ to $H^2(R)$ for R any bounded domain, and thus also bounded to $H^1(R)$. For $\partial\Omega$ Lipschitz the Sobolev trace theorem has certain restrictions on the order [25, Thm. 3.38], but, for example, the trace operator is bounded from $H^1(\Omega)$ to $H^{1/2}(\partial\Omega)$. Thus $V : L^2(\Omega) \rightarrow H^{1/2}(\partial\Omega)$ is bounded. Since $H^{1/2}(\partial\Omega)$ compactly imbeds into $L^2(\partial\Omega)$ on a Lipschitz boundary [25, Thm. 3.27 and p.99], the operator VBP is compact and the proof complete. \square

Note that D in (36) is *not compact* when $\partial\Omega$ has corners [9, Sec. 3.5], yet the theorem holds with corners since D is cancelled in the proof.

Figure 5 compares the spectrum of the unregularized (14) and regularized (35) operators, in a simple computational example. The improvement in the eigenvalue distribution is dramatic: the spectrum of (14) has small eigenvalues but extends to large eigenvalues of order 10^5 , while the spectrum of (35) is tightly clustered around $+1$ with no eigenvalues of magnitude larger than 2.

3.3. Reconstructing the scattered field on the exterior. Once equation (35) is solved for the scattered wave on $\partial\Omega$, we are able to find the scattered wave at any point in Ω^c via the representation (32). All that is needed is the normal derivative of u^s on $\partial\Omega$, which we easily get from equation (12) to be

$$(38) \quad u_n^s = T_{\text{int}}(u^i|_{\partial\Omega} + u^s|_{\partial\Omega}) - u_n^i.$$

For evaluation of (32), the native Nyström quadrature on $\partial\Omega$ is sufficient for 10-digit accuracy for all points further away from Ω than the size of one leaf box; however, as with any boundary integral method, for highly accurate evaluation very close to $\partial\Omega$ a modified quadrature would be needed.

3.4. Numerical discretization of the boundary integral equation. We discretize the BIE (35) on $\partial\Omega$ via a Nyström method with composite (panel-based) quadrature with n nodes in total. These panels on $\partial\Omega$ are coincident with the edges of the leaf boxes from the interior discretization, apart from the eight panels touching corners, where six levels of dyadic panel refinement are used on each to achieve around 10-digit accuracy.³ On each of these panels, a 10-point Gaussian rule is used.

³We note that some refinement is necessary even though the solution is smooth near the (fictitious) corners. However, the extra cost of refinement, as opposed to, say, local corner rounding, is negligible.

For building $n \times n$ matrix approximations to the operators S and D in (35), the plain Nyström method is used for matrix elements corresponding to non-neighboring panels, while generalized Gaussian quadrature for matrix elements corresponding to the self- or neighbor-interaction of each panel [18]. The matrix \mathbf{T}_{int} computed by (31) must also be interpolated from the $4q2^M$ Gauss nodes on $\partial\Omega$ to the n new nodes; since the panels are mostly coincident, this is a local operation analogous to the use of \mathbf{P} and \mathbf{Q} matrices in section 2.3.

4. COMPUTATIONAL COMPLEXITY

The computational cost of the solution technique is determined by the cost of constructing the approximate DtN operator \mathbf{T}_{int} and the cost of solving the boundary integral equation (35). Let N denote the total number of discretization points in Ω required for constructing \mathbf{R} . As there are p^2 Chebyshev points for each leaf box, the total number of discretization points is roughly $4^M p^2$ (to be precise, since points are shared on leaf box edges, it is $N = 4^M(p-1)^2 + 2^{M+1}(p-1) + 1$). Recall that n is the number of points on $\partial\Omega$ required to solve the integral equation. Note that $n \sim \sqrt{N}$.

4.1. Using dense linear algebra. Using dense linear algebra, the cost of constructing \mathbf{R} via the technique in section 2 is dominated by the cost at the top level where a matrix of size $\sqrt{N} \times \sqrt{N}$ is inverted. Thus the computational cost is $O(N^{3/2})$. The cost of approximating the DtN operator \mathbf{T}_{int} is also $O(N^{3/2})$. However, the computational cost of applying \mathbf{T} is $O(N)$. If the solution in the interior of Ω is desired, the computation cost of the solve is $O(N)$ as well.

The cost of inverting the linear system resulting from the (eg. Nyström) discretization of (35) is $O(n^3)$. It is possible to accelerate the solve by using iterative methods such as GMRES, which, given its second kind nature, would converge in $O(1)$ iterations.

When there are multiple incident waves at the same wavenumber κ , the solution technique should be separated into two steps: precomputation and solve. The precomputation step consist of constructing the approximate ItI, and DtN operators \mathbf{R} and \mathbf{T}_{int} , respectively. Also included in the precomputation should be the discretization and inversion of the BIE (35). The solve step then consist of applying the inverse of the system in (35). The precomputation need only be done once per wavenumber with a computational cost $O(N^{3/2})$. The cost of each solve (one for each incident wave) is simply the cost of applying an $n \times n$ dense matrix $\sim O(N)$.

4.2. Using fast algorithms. The matrices \mathbf{R}^T in Algorithm 1 that approximate ItI operators, as well as the matrices \mathbf{T}_{int} and \mathbf{T}_{ext} approximating DtN operators, all have internal structure that could be exploited to accelerate the matrix algebra. Specifically, the off-diagonal blocks of these matrices tend to have low numerical ranks, which means that they can be represented efficiently in so called “data-sparse” formats such as, e.g., \mathcal{H} or \mathcal{H}^2 -matrices [17, 6, 5], or, even better, the Hierarchically Block Separable (HBS) format [16, 20] (which is closely related to the “HSS” format [37]). If the wavenumber κ is kept fixed as N increases, it turns out to be possible to accelerate all computations in the build stage to optimal $O(N)$ asymptotic complexity, and the solve stage to optimal $O(N^{1/2})$ complexity, see [15]. However, the scaling constants suppressed by the big-O notation depend on κ in such a way that the use of accelerated matrix algebra is worthwhile primarily for problems of only moderate size (say a few dozen wavelengths across). Moreover, for high-order methods such as ours, it is common to keep the number of discretization nodes per wavelength fixed as N increases (so that $\kappa \sim N^{1/2}$), and in this environment, the scaling of the “accelerated” methods revert to $O(N^{3/2})$ and $O(N)$ for the build and the solve stages, respectively.




FIGURE 6. Plots (a) and (b) illustrate $b(\mathbf{x})$ for experiments *Bumps 1* and *Bumps 2* in section 5.1. Plots (c) and (d) illustrate the real part of the total field for each experiment respectively.

5. NUMERICAL EXPERIMENTS

This section reports on the performance of the new solution technique for several choices of potential $b(\mathbf{x})$ where the (numerical) support of b is contained in $\Omega = (-0.5, 0.5)^2$. The incident wave is a plane wave $u^i(\mathbf{x}) = e^{i\kappa\mathbf{w}\cdot\mathbf{x}}$ with incident unit direction vector $\mathbf{w} \in \mathbb{R}^2$.

Firstly, in section 5.1 the method is applied to problems where $b(\mathbf{x})$ is a single Gaussian “bump.” In this case the radial symmetry allows for an independent semi-analytic solution, which we use to verify the accuracy of the method. Then section 5.2 reports on the performance of the method when applied to more complicated problems. Finally, section 5.3 illustrates the computational cost in practice.

For all the experiments, for the composite spectral method described in section 2 we use a $p \times p$ Chebyshev tensor product grid of points per leaf with $p = 16$, and the number of Gaussian nodes per side of a leaf is $q = 14$.

We implemented the methods based on dense matrix algebra with $O(N^{3/2})$ asymptotic complexity described in section 4.1. (We do not use the $O(N)$ accelerated techniques of section 4.2 since we are primarily interested in scatterers that are large in comparison to the wavelength.)

All experiments were executed on a desktop workstation with two quad-core Intel Xeon E5-2643 processors and 128 GB of RAM. All computations were done in MATLAB (version 2012b), apart from the evaluation of Hankel functions in the Nyström and scattered wave calculations, which use Fortran. We expect that, by careful implementation of the whole scheme in a compiled language, execution times could be improved substantially.

5.1. Accuracy of the method. In this section, we consider problems where the scattering potential $b(\mathbf{x})$ is given by a Gaussian bump. Since b has radial symmetry, we may compute an accurate reference scattering solution by solving a series of ODEs, as explained in Appendix A. With $\kappa = 40$ (so that the square Ω is around six free-space wavelengths on a side), and $\mathbf{w} = (1, 0)$, we consider two problems given as follows,

$$\begin{aligned} \text{Bump 1: } b(\mathbf{x}) &= -1.5e^{-160|\mathbf{x}|^2}, \\ \text{Bump 2: } b(\mathbf{x}) &= 1.5e^{-160|\mathbf{x}|^2}. \end{aligned}$$

For *Bump 1* the bump region has an increased refractive index, varying from 1 to around 1.58, which can be interpreted as an attractive potential. For *Bump 2* the potential is repulsive, causing the waves to become slightly evanescent near the origin (here the refractive index decreases to zero then becomes purely imaginary, but note that this does not correspond to absorption.) Figure 6 illustrates the geometry and the resulting real part of the total field for each experiment.

Let \tilde{u} denote the approximate total field constructed via the proposed method, and u denote the reference total field computed as in Appendix A. Table 1 reports

- N : the number of discretization points used by the composite spectral method in Ω ,
- n : the number of discretization points used for discretizing the BIE,
- $\text{Re } \tilde{u}(0.5, 0)$: real part of the approximate solution at $(0.5, 0)$ (on $\partial\Omega$),
- $e_1 = |u(0.5, 0) - \tilde{u}(0.5, 0)|$,
- $\text{Re } \tilde{u}(1, 0.5)$: real part of the approximate solution at $(1, 0.5)$ (outside of Ω),
- $e_2 = |u(1, 0.5) - \tilde{u}(1, 0.5)|$.

In the table, the number of levels M grows from 2 to 5, roughly quadrupling N each time. High-order convergence is apparent, reaching an accuracy of 9-10 digits (accuracy does not increase much beyond 10 digits). At the highest N , there are about 50 gridpoints per wavelength at the shortest wavelength occurring at the center of *Bump 1*.

	N	n	$\text{Re } \tilde{u}(0.5, 0)$	e_1	$\text{Re } \tilde{u}(1, 0.5)$	e_2
<i>Bump 1</i>	3721	640	-0.98792561833285	7.75e-05	-1.12207402682737	5.09e-05
	14641	800	-0.987981264965721	1.78e-07	-1.12205758254400	8.18e-08
	58081	1120	-0.987981217277174	2.56e-09	-1.12205766387011	1.15e-10
	231361	1760	-0.987981215350216	9.31e-10	-1.12205766378840	7.90e-11
<i>Bump 2</i>	3721	640	-0.0470314782486572	4.82e-05	-1.01063677552351	3.23e-05
	14641	800	-0.0470619180279044	4.25e-08	-1.01065022958956	7.32e-08
	58081	1120	-0.0470619010992819	1.32e-09	-1.01065028579517	1.29e-10
	231361	1760	-0.0470619007119554	5.07e-10	-1.01065028569638	4.36e-11

TABLE 1. Approximate solutions and pointwise errors for the experiments in section 5.1

5.2. Performance for challenging scattering potentials. This section illustrates the performance of the numerical method for problems with smoothly varying wave speed inside of Ω .




FIGURE 7. Plots of the different scattering potentials b inside Ω in section 5.2.

We consider three different test cases of scattering potential b . They are

Lens: A vertically-graded lens (Figure 7(a)), at wavenumber $\kappa = 300$. Specifically, $b(\mathbf{x}) = 4(x_2 - 0.2)[1 - \text{erf}(25(|\mathbf{x}| - 0.3))]$, where $\mathbf{x} = (x_1, x_2)$. The maximum refractive index is around 2.1

Random bumps: The sum of 200 wide Gaussian bumps randomly placed in Ω , rolled off to zero (see Figure 7(b)) giving a smooth random potential. Wavenumber $\kappa = 160$. The maximum refractive index is around 4.3.

Photonic crystal: 20×20 square array of small Gaussian bumps (with peak refractive index 6.7) with a “waveguide” channel removed (Figure 7(c)). The wavenumber $\kappa = 77.1$ is chosen carefully to lie in the first complete bandgap of the crystal.

For the first two cases, Ω is around 70 wavelengths on a side, measured using the typical wavelength occurring in the medium (for the *lens* case, it is 100 wavelengths on a side at the minimum wavelength). This is quite a high frequency for a variable-medium problem at the accuracies we achieve. In these two cases the waves mostly *propagate*; in contrast, in the third case the waves

mostly *resonate* within each small bump, in such a way that large-scale propagation through the crystal is impossible (hence evanescent), except in the channel.⁴

For each choice of varying wave speed, the incident wave is in the direction $\mathbf{w} = (1, 0)$ (we remind the reader that the method works for arbitrary incident direction.) For the *Photonic crystal*, we also consider the incident wave direction $\mathbf{w} = (-\sqrt{2}/2, \sqrt{2}/2)$. There are no reference solutions available for these problems, hence we study convergence.

In addition to the number of discretization points N and n , Table 2 reports

- Re $\tilde{u}(0.25, 0)$: real part of the approximate solution at $(0.25, 0)$ (inside of Ω),
- e_1 : $= |\tilde{u}_N(0.25, 0) - \tilde{u}_{4N}(0.25, 0)|$, an estimate of the pointwise error,
- Re $\tilde{u}(1, 0.5)$: real part of the approximate solution at $(1, 0.5)$ (outside of Ω),
- e_2 : $= |\tilde{u}_N(1, 0.5) - \tilde{u}_{4N}(1, 0.5)|$, an estimate of the pointwise error.

	N	n	Re $\tilde{u}(0.5, 0)$	e_1	Re $\tilde{u}(1, 0.5)$	e_2
Lens	58081	1120	-0.373405022283892	2.02e-01	-0.547735180732198	5.09e-01
	231361	1760	-0.221345395796661	3.53e-03	0.161212542340161	2.86e-03
	923521	3040	-0.218651605400620	1.61e-07	0.158422280450864	1.87e-07
	3690241	5600	-0.218651458554288	6.85e-10	0.158422464920298	6.99e-10
	14753281	10720	-0.218651458391577	-	0.158422464625727	-
Bumps	58081	1120	1.29105948477323	1.91	-0.612141744074168	0.44
	231361	1760	0.359271869087464	1.99e-02	-0.931198083868205	1.87e-02
	923521	3040	0.374697595070227	2.76e-06	-0.945752835546445	7.95e-07
	3690241	5600	0.374698518812982	3.06e-09	-0.945753626496863	1.58e-09
	14753281	10720	0.374698518930658	-	-0.945753627849080	-
Crystal	58081	1120	-0.406418011063883	2.03	-0.129067996215635	3.42e-01
	231361	1760	0.0424527158875615	1.53e-03	0.195870563479998	1.82e-4
	923521	3040	0.0437392735790711	7.30e-07	0.195981633749759	2.81e-07
	3690241	5600	0.0437393320806644	4.84e-10	0.195981570696692	7.56e-10
	14753281	10720	0.0437393324622741	-	0.195981570519668	-
Crystal $\mathbf{w} = \left(-\frac{\sqrt{2}}{2}, \frac{\sqrt{2}}{2}\right)$	58081	1120	-0.0420633119821246	1.28	-1.20915538109562	1.37e-01
	231361	1760	0.0367128964251903	3.47e-03	-1.09529393341122	1.10e-3
	923521	3040	0.0376839447575234	1.58e-06	-1.09445429347097	7.25e-07
	3690241	5600	0.0376833752064704	6.79e-10	-1.09445431363369	4.90e-9
	14753281	10720	0.0376833752704930	-	-1.09445430874556	-

TABLE 2. Convergence results for the experiments in section 5.2.

Table 2 shows that typically 9-digit accuracy is reached when $N \approx 3.7 \times 10^6$ ($M = 7$), which corresponds to 1921 Chebyshev nodes in each direction, or around 20 nodes per wavelength at the shortest wavelengths in each medium.

5.3. Scaling of the method. Recall that in the case of multiple incident waves, the solution technique should be broken into two steps: precomputation and solve. Since a direct solver is used, the timing results are independent of the particular scattering potential.

For each choice of N and n , Table 3 reports

⁴The choice of bump height and width needs to be made carefully to ensure that a usable bandgap exists; this was done by creating a separate spectral solver for the band structure of the periodic problem.

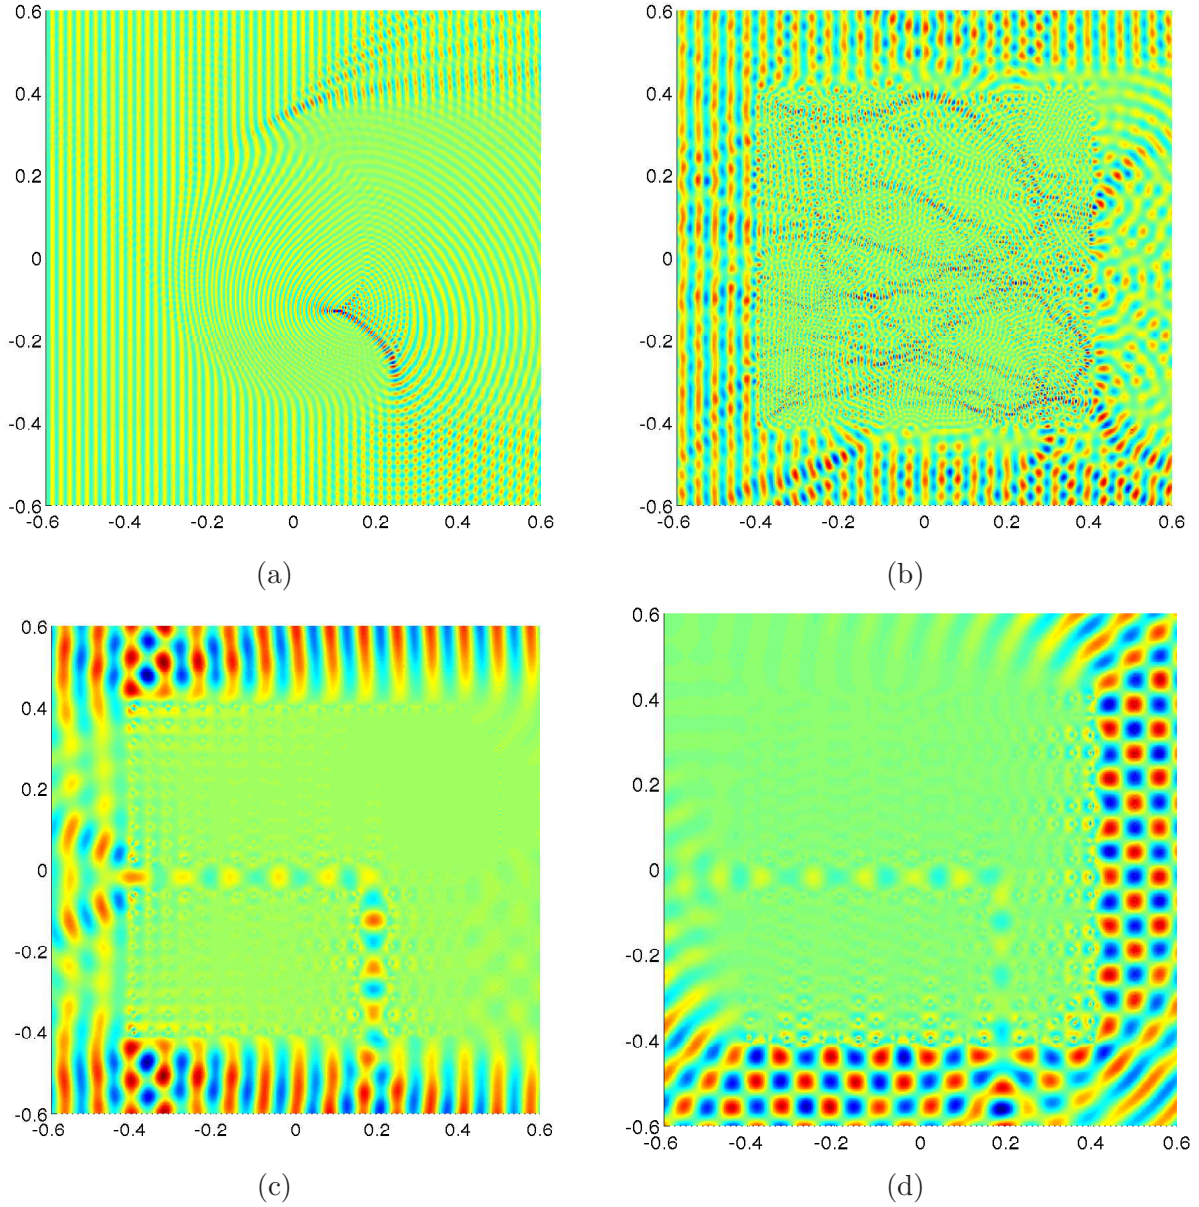


FIGURE 8. Plots of the real part of the total field for the four experiments in section 5.2 whose scattering potentials are shown in Figure 7: (a) lens, (b) random bumps, (c) and (d) photonic crystal with different incident wave directions.

T_{build} : Time in seconds to building the approximate ItI and DtN,
 T_{solve} : Time in seconds to discretize and invert the BIE (35),
 T_{apply} : Time in seconds to apply the inverse A^{-1} of the discretized integral equation
 R_{build} : Memory in MB required to store the ItI and solution operators in the hierarchical scheme,
 R_{solve} : Memory in MB required to store the discretized inverse A^{-1} .

Figure 9 plots the timings against the problem size N . (The total precomputation time is the sum of T_{build} and T_{solve} .) The results show that even at the largest N tested, the precomputation time has not reached its asymptotic $O(N^{3/2})$; the large dense linear algebra has not yet started to

dominate T_{build} (this may be due to MATLAB overheads). However, the cost for solving (35) and applying the inverse scale closer to expectations.

The memory usage scales as the expected $O(N \log N)$. We are not able to test beyond 15 million unknowns ($M = 8$) since by that point the memory usage approaches 100 GB. However, note that if all that is needed is the *far-field* solution for arbitrary incident waves at one wavenumber, the \mathbf{S}^τ and \mathbf{Y}^τ solution matrices need not be stored, reducing memory significantly, and the final solution matrix only requires 2 GB. We note that, extrapolating from the convergence study, this N should be sufficient for 9-digit accuracy for problems up to 200 wavelengths on a side.

N	n	T_{build}	T_{solve}	T_{apply}	R_{build}	R_{solve}
3721	640	0.506	1.78	5.39e-04	9.71	6.25
14641	800	0.709	2.01	8.28e-04	48.07	9.77
58081	1120	2.90	3.01	1.73e-03	229.05	19.14
231361	1760	12.09	5.40	3.32e-03	1063.23	47.37
923521	3040	51.67	13.23	1.05e-02	4841.01	141.02
3690241	5600	231.18	40.79	4.03e-02	21716.21	478.52
14753281	10720	1081.09	185.54	1.13e-01	96273.17	1753.52

TABLE 3. T_{build} and R_{build} report the time in seconds and memory in MB, respectively, required for building the interior ItI operator and constructing the discretized integral equation (35). T_{solve} reports the time in seconds required to invert the discretized system, while R_{solve} reports the memory in MB to store the inverse. T_{apply} reports the time in seconds required to apply the inverse to the incident wave dependent data. This table is independent of the choice of potential or wavenumber.





FIGURE 9. Time in seconds for each step in the proposed method. (For the comparison power law graphs, the prefactors are $C_1 = 2 \times 10^{-7}$ and $C_2 = 2 \times 10^{-9}$.)

6. CONCLUDING REMARKS

This paper presents a robust, high accuracy direct method for solving scattering problems involving smoothly varying media. Numerical results show that the method converges to high order, as

expected, for choices of refractive index that are representative of challenging problems that occur in applications. Namely, for problems dominated either by propagation (lenses) or by resonances (a bandgap photonic crystal), with of order 100 wavelengths on a side, the method converges to around 9-digit accuracy with 3.7 million unknowns. The method is ideal for problems where the far field scattering is desired for multiple incident waves, since each additional incident wave requires merely applying a dense matrix to its boundary data. For example, a problem involving 14 million unknowns requires 21 minutes of precomputation (to build the necessary operators), but each additional solve takes approximately 0.1 seconds. As discussed in section 4.2, for low frequency problems, these timings, and asymptotic behavior, could be improved by replacing the dense linear algebra by faster algorithms exploiting compressed representations. Another open question is the existence of a convenient second-kind formulation which involves the ItI map (and not the DtN map) of the domain Ω .

ACKNOWLEDGMENTS

We are grateful for a helpful discussion with Michael Weinstein. The work of AHB is supported by NSF grant DMS-1216656. ***** OTHER SUPPORT for AG and PGM.

APPENDIX A. REFERENCE SOLUTION FOR PLANE WAVE SCATTERING FROM RADIAL POTENTIALS

In this appendix we describe how we generate reference solutions with around 13 digits of accuracy for the scattering problem from smooth radially-symmetric potentials such as

$$(39) \quad b(r) = \pm 1.5e^{-160r^2},$$

which are needed in section 5.1. Here (r, θ) are polar coordinates; in what follows $(x_1, x_2) \in \mathbb{R}^2$ indicate Cartesian coordinates. We choose a solution domain radius $R > 0$ such that b is numerically negligible outside the ball $r < R$. A plane wave incident in the positive x_1 -direction is decomposed into a polar Fourier (“angular momentum”) basis via the Jacobi–Anger expansion [27, 10.12.5],

$$e^{i\kappa x_1} = J_0(\kappa r) + 2 \sum_{l=1}^{\infty} i^l J_l(\kappa r) \cos l\theta .$$

We write $J_l(z) = (H_l^{(1)}(z) + H_l^{(2)}(z))/2$, and then notice that the effect of the potential b on this field is to modify only the *outgoing* scattering coefficients. Thus, restricting to a maximum order L , the full field becomes

$$(40) \quad u(r, \theta) \approx \frac{1}{2}[H_0^{(1)}(\kappa r) + a_0 H_0^{(2)}(\kappa r)] + \sum_{l=1}^L i^l [H_l^{(1)}(\kappa r) + a_l H_l^{(2)}(\kappa r)] \cos l\theta, \quad r > R .$$

The coefficients $\{a_l\}$ are known as *scattering phases*; by flux conservation they lie on the unit circle if $b(r)$ is a real-valued function. Convergence with respect to L is exponential, once L exceeds κR . For the case of (39) we choose $R = 0.5$ and $L = 30$.

The phases are found in the following way. For each $l = 0, \dots, L$ we solve the homogeneous radial ODE,

$$u_l'' + \frac{1}{r}u_l' + \left[-\frac{l^2}{r^2} + (1 - b(r))\kappa^2 \right] u_l = 0, \quad 0 < r < R$$

with initial conditions that correspond to a regular solution of the form $u_l(r) \sim cr^l$ as $r \rightarrow 0^+$ (we implement the initial condition by restricting the domain to $[r_0, R]$ for some small number $r_0 > 0$ chosen such that the solution growing with increasing r dominates sufficiently over the decaying one). For the numerical solution we use MATLAB’s `ode45` with machine precision requested for absolute and relative tolerances. (We note that the standard transformation $u(r) = r^l U(r)$ which mollifies the

behavior at $r = 0$ resulted in no improvement in accuracy.) After extracting each interior solution's Robin constant $\beta_l := u_l'(R)/u_l(R)$, and matching value and derivative to (40) at $r = R$, we get after simplification,

$$a_l = -\frac{\alpha_l^*}{\alpha_l}, \quad \text{where } \alpha_l = \kappa H_l^{(1)'}(\kappa R) - H_l^{(1)}(\kappa R)\beta_l,$$

which completes the recipe for the phases. The computation time required is a few seconds, due to the large number of steps required by `ode45`. A simple accuracy test is independence of the phases with respect to variation in R . Values of $u(r, \theta)$ for $r \geq R$ may then be found via evaluating the sum in (40), and for $r < R$ by summation of the interior solutions $\{u_l(r)\}$.

REFERENCES

- [1] W. Ang. *A beginner's course in boundary element methods*. Universal Publishers, USA, 2007.
- [2] I. M. Babuska and S. A. Sauter. Is the pollution effect of the FEM avoidable for the Helmholtz equation considering high wave numbers? *SIAM J. Numer. Anal.*, 34(6):2392–2423, 1997.
- [3] A. Bayliss, C. I. Goldstein, and E. Turkel. The numerical solution of the Helmholtz equation for wave propagation problems in underwater acoustics. *Comput. Math. Appl.*, 11(7–8):655–665, 1985.
- [4] A. Bayliss, C. I. Goldstein, and E. Turkel. On accuracy conditions for the numerical computation of waves. *J. Comput. Phys.*, 59(3):396–404, 1985.
- [5] M. Bebendorf. *Hierarchical matrices*, volume 63 of *Lecture Notes in Computational Science and Engineering*. Springer-Verlag, Berlin, 2008.
- [6] S. Börm. *Efficient numerical methods for non-local operators*, volume 14 of *EMS Tracts in Mathematics*. European Mathematical Society (EMS), Zürich, 2010.
- [7] S. Britt, S. Tsynkov, and E. Turkel. Numerical simulation of time-harmonic waves in inhomogeneous media using compact high order schemes. *Commun. Comput. Phys.*, 9(3):520–541, 2011.
- [8] Y. Chen. A fast, direct algorithm for the Lippmann-Schwinger integral equation in two dimensions. *Adv. Comput. Math.*, 16(2-3):175–190, 2002.
- [9] D. Colton and R. Kress. *Inverse acoustic and electromagnetic scattering theory*, volume 93 of *Applied Mathematical Sciences*. Springer-Verlag, Berlin, second edition, 1998.
- [10] R. Duan and V. Rokhlin. High-order quadratures for the solution of scattering problems in two dimensions. *J. Comput. Physics*, 228(6):2152–2174, 2009.
- [11] B. Engquist and A. Majda. Absorbing boundary conditions for the numerical simulation of waves. *Math. Comp.*, 31:629–651, 1977.
- [12] B. Engquist and L. Ying. Sweeping preconditioner for the Helmholtz equation: Moving perfectly matched layers. *Multiscale Mod. Sim.*, 9(2):686–710, 2011.
- [13] Q. Fang, P. M. Meaney, and K. D. Paulsen. Viable three-dimensional medical microwave tomography: Theory and numerical experiments. *IEEE Trans. Antennas Propag.*, 58(2):449–458, 2010.
- [14] L. Friedlander. Some inequalities between Dirichlet and Neumann eigenvalues. *Arch. Rational Mech. Anal.*, 116(2):153–160, 1991.
- [15] A. Gillman and P. Martinsson. A direct solver with $O(N)$ complexity for variable coefficient elliptic PDEs discretized via a high-order composite spectral collocation method, 2013. In review.
- [16] A. Gillman, P. Young, and P. Martinsson. A direct solver $O(N)$ complexity for integral equations on one-dimensional domains. *Front. Math. China*, 7:217–247, 2012. 10.1007/s11464-012-0188-3.
- [17] W. Hackbusch. A sparse matrix arithmetic based on H-matrices; Part I: Introduction to H-matrices. *Computing*, 62:89–108, 1999.
- [18] 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, 2013. in press, *Adv. Comput. Math.*
- [19] E. Heikkola, T. Rossi, and J. Toivanen. Fast direct solution of the Helmholtz equation with a perfectly matched layer/an absorbing boundary condition. *Int. J. Numer. Meth. Engng*, 57:2007–2025, 2003.
- [20] K. Ho and L. Greengard. A fast direct solver for structured linear systems by recursive skeletonization. *SIAM J. Sci. Comput.*, 34(5):A2507–A2532, 2012.
- [21] A. Kirsch and P. Monk. An analysis of the coupling of finite-element and Nyström methods in acoustic scattering. *IMA J. Numer. Anal.*, 14:523–544, 1994.
- [22] D. Kopriva. A staggered-grid multidomain spectral method for the compressible Navier-Stokes equations. *J. Comput. Phys*, 143(1):125–158, 1998.
- [23] P. Martinsson. A fast direct solver for a class of elliptic partial differential equations. *J. Sci. Comput.*, 38(3):316–330, 2009.

- [24] P. Martinsson. A direct solver for variable coefficient elliptic PDEs discretized via a composite spectral collocation method. *J. Comput. Phys.*, 242:460–479, 2013.
- [25] W. McLean. *Strongly elliptic systems and boundary integral equations*. Cambridge University Press, Cambridge, 2000.
- [26] D. P. Nicholls and N. Nigam. Exact non-reflecting boundary conditions on general domains. *J. Computat. Phys.*, 194(1):278–303, 2004.
- [27] F. Olver, D. Lozier, R. Boisvert, and C. Clark, editors. *NIST Handbook of Mathematical Functions*. Cambridge University Press, 2010. <http://dlmf.nist.gov>.
- [28] H. Pfeiffer, L. Kidder, M. Scheel, and S. Teukolsky. A multidomain spectral method for solving elliptic equations. *Comput. Phys. Commun.*, 152(3):253–273, 2003.
- [29] H. E. Salzer. Lagrangian interpolation at the Chebyshev points $x_{n,\nu} \equiv \cos(\nu\pi/n)$, $\nu = 0(1)n$; some unnoted advantages. *Comput. J.*, 15:156–159, 1972.
- [30] P. Schmitz and L. Ying. A fast direct solver for elliptic problems on Cartesian meshes in 3D, 2011. preprint.
- [31] P. Schmitz and L. Ying. A fast direct solver for elliptic problems on general meshes in 2D. *J. Comput. Phys.*, 231(4):1314 – 1338, 2012.
- [32] L. Trefethen. *Spectral Methods in MATLAB*. SIAM, Philadelphia, 2000.
- [33] E. E. Tyrtyshnikov. *A brief introduction to numerical analysis*. Birkhäuser Boston, 1997.
- [34] E. Wadbro and M. Berggren. High contrast microwave tomography using topology optimization techniques. *J. Comput. Appl. Math.*, 234:1773–1780, 2010.
- [35] S. Wang, M. V. de Hoop, and J. Xia. On 3D modeling of seismic wave propagation via a structured parallel multifrontal direct Helmholtz solver. *Geophysical Prospecting*, 59(5):857–873, 2011.
- [36] J. Xia, S. Chandrasekaran, M. Gu, and X. Li. Superfast multifrontal method for large structured linear systems of equations. *SIAM J. Matrix Anal. Appl.*, 31(3):1382–1411, 2009.
- [37] J. Xia, S. Chandrasekaran, M. Gu, and X. S. Li. Fast algorithms for hierarchically semiseparable matrices. *Numer. Linear Algebra Appl.*, 17(6):953–976, 2010.
- [38] B. Yang and J. Hesthaven. Multidomain pseudospectral computation of Maxwell’s equations in 3-D general curvilinear coordinates. *Appl. Numer. Math.*, 33(1–4):281 – 289, 2000.