

MIT Open Access Articles

Compressed Absorbing Boundary Conditions via Matrix Probing

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

Citation: Belanger-Rioux, R., and L. Demanet. "Compressed Absorbing Boundary Conditions via Matrix Probing." *SIAM J. Numer. Anal.* 53, no. 5 (January 2015): 2441–2471. © 2015 Society for Industrial and Applied Mathematics

As Published: <http://dx.doi.org/10.1137/14095563x>

Publisher: Society for Industrial and Applied Mathematics

Persistent URL: <http://hdl.handle.net/1721.1/100894>

Version: Final published version: final published article, as it appeared in a journal, conference proceedings, or other formally published context

Terms of Use: Article is made available in accordance with the publisher's policy and may be subject to US copyright law. Please refer to the publisher's site for terms of use.



COMPRESSED ABSORBING BOUNDARY CONDITIONS VIA MATRIX PROBING*

R. BÉLANGER-RIOUX[†] AND L. DEMANET[‡]

Abstract. Absorbing layers are sometimes required to be impractically thick in order to offer an accurate approximation of an absorbing boundary condition for the Helmholtz equation in a heterogeneous medium. It is always possible to reduce an absorbing layer to an operator at the boundary by layer stripping elimination of the exterior unknowns, but the linear algebra involved is costly. We propose bypassing the elimination procedure and directly fitting the surface-to-surface operator in compressed form from a few exterior Helmholtz solves with random Dirichlet data. The result is a concise description of the absorbing boundary condition, with a complexity that grows slowly (often, logarithmically) in the frequency parameter.

Key words. absorbing boundary condition, nonreflecting boundary condition, radiating boundary condition, open boundary condition, Helmholtz equation, matrix probing, heterogeneous media

AMS subject classifications. 65N99, 65T40, 65F35

DOI. 10.1137/14095563X

1. Introduction. This paper investigates arbitrarily accurate realizations of absorbing (a.k.a. open, radiating) boundary conditions (ABCs), including absorbing layers, for the two-dimensional (2D) acoustic high-frequency Helmholtz equation in certain kinds of heterogeneous media. Instead of considering a specific modification of the partial differential equation, such as a perfectly matched layer, studied here is the broader question of compressibility of the nonlocal kernel that appears in the exact boundary integral form $\mathcal{D}u = \partial_\nu u$ of the ABC, where $\partial_\nu u$ is the outward normal derivative. The operator \mathcal{D} is called the *Dirichlet-to-Neumann* (DtN) map. This boundary integral viewpoint invites one to rethink ABCs as a two-step numerical scheme, where

1. a precomputation sets up an expansion of the kernel of the boundary integral equation; then
2. a fast algorithm is used for each application of this integral kernel at the open boundaries in a Helmholtz solver.

This two-step approach may pay off in scenarios when the precomputation is amortized over a large number of solves of the original equation with different data.

This paper addresses the precomputation step: a basis for the efficient expansion of the integral kernel of the ABC in some simple 2D settings is described, and a randomized probing procedure to quickly find the coefficients in the expansion is discussed. This framework is, interestingly, halfway between a purely analytical or physical method and a purely numerical one. It uses both the theoretical grounding of analytic knowledge and the intuition from understanding the physics of the problem in order to obtain a useful solution.

*Received by the editors February 7, 2014; accepted for publication (in revised form) August 26, 2015; published electronically October 29, 2015. This research was supported by AFOSR, ONR, NSF, Total SA, and the Alfred P. Sloan Foundation.

<http://www.siam.org/journals/sinum/53-5/95563.html>

[†]Department of Mathematics, Harvard University, Cambridge, MA 02138 (rbr@math.harvard.edu). The work of this author was supported by the National Sciences and Engineering Research Council of Canada.

[‡]Department of Mathematics, Massachusetts Institute of Technology, Cambridge, MA 02139 (laurent@math.mit.edu).

The numerical realization of ABC typically involves absorbing layers that become very thick (on the order of $O(N)$ grid points) for difficult $c(\mathbf{x})$ or for high accuracy (see section 2.2). In this paper, we propose instead to realize the ABC by directly compressing the integral kernel of \mathcal{D} , so that the computational cost of its setup and application would become competitive when the Helmholtz equation is to be solved multiple times. Hence this paper is not concerned with the design of a new ABC, but rather with the reformulation of existing ABCs that otherwise require a lot of computational work per solve.

In many situations of practical interest we show that it is possible to “learn” the integral form of \mathcal{D} , as a precomputation, from a small number of solves of the original problem with the expensive ABC. In fact the computational cost of the present method is approximately that of solving the Helmholtz equation a few times (see section 3.2), with the advantage of speeding up further solves (see section 2.3). By “a few times” we mean a quantity essentially independent of the number of discretization points N along one dimension—in practice as small as 1 or as large as 50. As mentioned, this new approach becomes competitive when a large number of solves is required. This strategy is called *matrix probing* [12, 17]. In these situations, we also show that the accurate expansion of \mathcal{D} requires a number of parameters p “essentially independent” of N (see section 4.4). This is in sharp contrast to the scaling of $O(N)$ grid points for the layer width mentioned above. However, the basis for the expansion needs to be chosen carefully. A rationale for this choice is provided in section 3.3, and a proof of convergence in a special case is presented in section 5. This leads to the successful design of a basis for a variety of heterogeneous media.

Once a matrix realization \tilde{D} of the DtN map \mathcal{D} is obtained from matrix probing, it can be used in a Helmholtz solver. However, a solver would use matrix-vector multiplications to apply the dense matrix \tilde{D} . Hence the second step of our numerical scheme: \tilde{D} needs to be compressed further, in a form that provides a fast matrix-vector product. This second step can only come after the first, since it is the first step that gives us access to the entries in D and allows us to use compression algorithms of interest. Work on this second step appears here [6]. Note that the feasibility of probing and the availability of a fast algorithm for matrix-vector multiplication are two different goals that require different expansion schemes.

Section 2 introduces the DtN map, and how to obtain it both analytically and numerically from an ABC. However, as will be explained, ABCs are computationally expensive when solving the Helmholtz equation in a heterogeneous medium, and this is why matrix probing is introduced in section 3 to compress such an ABC. Section 4 presents numerical results documenting the complexity and accuracy of matrix probing, including results on using an ABC, compressed using matrix probing, in a Helmholtz solver. Section 5 presents a proof that matrix probing converges rapidly to what is called the half-space DtN map in uniform medium, which justifies particular choices made in the setup of the matrix probing algorithm of section 3.

2. Background and setup. Consider the Helmholtz equation in \mathbb{R}^2 ,

$$(2.1) \quad \Delta u(\mathbf{x}) + \frac{\omega^2}{c^2(\mathbf{x})}u(\mathbf{x}) = f(\mathbf{x}), \quad \mathbf{x} = (x_1, x_2),$$

with compactly supported f . Throughout this article the unique solution determined by the Sommerfeld radiation condition (SRC) at infinity will be considered: when

$c(\mathbf{x})$ extends to a constant c outside of a bounded set, the SRC is [40]

$$(2.2) \quad \lim_{\rho \rightarrow \infty} \rho^{1/2} \left(\frac{\partial u}{\partial \rho} - iku \right) = 0, \quad k = \frac{\omega}{c},$$

where ρ is the radial coordinate.

2.1. The Dirichlet-to-Neumann map. The SRC can be reformulated to the boundary $\partial\Omega$, so that the resulting solution inside Ω matches that of the free-space problem (2.1), (2.2). There are many ways to do this numerically, as shall be seen in section 2.2. However, the following analytical reformulation is important because it introduces a fundamental concept, the *Dirichlet-to-Neumann (DtN) map*.

Let $G(\mathbf{x}, \mathbf{y})$ be the fundamental solution for this problem, i.e., the solution when $f(\mathbf{x}) = \delta(\mathbf{x} - \mathbf{y})$. Define the single- and double-layer potentials, respectively, on some closed contour Γ by the following, for ψ, ϕ on Γ (see details in [40, 13]):

$$(2.3) \quad \mathcal{S}\psi(\mathbf{x}) = \int_{\Gamma} G(\mathbf{x}, \mathbf{y}) \psi(\mathbf{y}) dS_{\mathbf{y}}, \quad \mathcal{T}\phi(\mathbf{x}) = \int_{\Gamma} \frac{\partial G}{\partial \nu_{\mathbf{y}}}(\mathbf{x}, \mathbf{y}) \phi(\mathbf{y}) dS_{\mathbf{y}},$$

where ν is the outward-pointing normal to the curve Γ , and \mathbf{x} is not on Γ . Now let u^+ satisfy the Helmholtz equation (2.1) in the exterior domain $\mathbb{R}^2 \setminus \overline{\Omega}$, along with the SRC (2.2). Then Green's third identity is satisfied in the exterior domain: using $\Gamma = \partial\Omega$, it follows that

$$(2.4) \quad \mathcal{T}u^+ - \mathcal{S} \frac{\partial u^+}{\partial \nu} = u^+, \quad \mathbf{x} \in \mathbb{R}^2 \setminus \overline{\Omega}.$$

Finally, using the jump condition of the double layer \mathcal{T} , Green's identity on the boundary $\partial\Omega$ is obtained:

$$\left(\mathcal{T} - \frac{1}{2}\mathcal{I} \right) u^+ - \mathcal{S} \frac{\partial u^+}{\partial \nu} = 0, \quad \mathbf{x} \in \partial\Omega,$$

where \mathcal{I} is the identity operator.

When the single-layer potential \mathcal{S} is invertible,¹ we can let $\mathcal{D} = \mathcal{S}^{-1}(\mathcal{T} - \frac{1}{2}\mathcal{I})$ and equivalently write (dropping the $+$ in the notation)

$$(2.5) \quad \frac{\partial u}{\partial \nu} = \mathcal{D}u, \quad \mathbf{x} \in \partial\Omega.$$

The operator \mathcal{D} is called the *exterior Dirichlet-to-Neumann map* because it maps the Dirichlet data u to the Neumann data $\partial u / \partial \nu$. The DtN map is independent of the right-hand side f of (2.1) as long as f is supported in Ω . The notion that (2.5) can serve as an exact ABC was made clear in a uniform medium, e.g., in [21] and in [36]. Equation (2.5) continues to hold even when $c(\mathbf{x})$ is heterogeneous in the vicinity of $\partial\Omega$, provided the correct (often unknown) Green's function is used. The medium is indeed heterogeneous near $\partial\Omega$ in many situations of practical interest, such as in geophysics.

Note that the interior DtN map has a much higher degree of complexity than the exterior one, because it needs to encode all the waves that travel along the broken

¹This is the case when there is no interior resonance at frequency ω . If there is such a resonance, it could be circumvented by the use of combined field integral equations as in [13]. The existence and regularity of \mathcal{D} ultimately do not depend on the invertibility of \mathcal{S} .

geodesic rays that go from one part of the domain Ω to another (i.e., rays bouncing inside the domain). In contrast, the exterior DtN map rarely needs to take into account multiple scattering if the solution is outgoing. Only the exterior DtN map is considered in this work and is referred to as the DtN map for simplicity.

The DtN map \mathcal{D} is symmetric.² The proof of the symmetry of \mathcal{D} was shown in a slightly different setting in [37] and is adapted to the present situation. Indeed, consider again $\mathcal{D} = \mathcal{S}^{-1}(\mathcal{T} - \frac{1}{2}\mathcal{I})$ from above, and use it in the alternative Steklov–Poincaré identity (mentioned in section 3.3, where \mathcal{H} and \mathcal{T}^* are defined) $\mathcal{D} = \mathcal{H} - (\mathcal{T}^* - \frac{1}{2}\mathcal{I})\mathcal{D}$, to obtain the symmetric expression $\mathcal{D} = \mathcal{H} - (\mathcal{T}^* - \frac{1}{2}\mathcal{I})\mathcal{S}^{-1}(\mathcal{T} - \frac{1}{2}\mathcal{I})$.

Much more is known about DtN maps, such as the many boundedness and coercivity theorems between adequate fractional Sobolev spaces (mostly in free space, with various smoothness assumptions on the boundary). It was not attempted to leverage these properties of \mathcal{D} in the scheme presented here, but more can be found in [41], as well as in [47] and references therein.

In this paper, unless otherwise noted, \mathcal{D} is used to denote the exact, analytical DtN map as described above, while D refers to the discrete realization of the DtN map as will be made precise in section 2.3, and \tilde{D} will be used for the approximation of matrix D in the probing framework.

2.2. Discrete absorbing boundary conditions. There are many ways to realize an absorbing boundary condition (ABC) for the wave or Helmholtz equation. Some ABCs are surface-to-surface, such as in [21, 31, 32, 34, 36, 46]. Others involve surrounding the computational domain Ω by an absorbing layer, or AL [2, 3, 8]. The latter approach is desirable because the parameters of the layer can usually be adjusted to obtain the required accuracy.

The perfectly matched layer (PML) of Bérenger [8] is a convincing solution to this problem in a uniform acoustic medium. Its performance often carries through in a general heterogeneous acoustic medium $c(\mathbf{x})$, though its derivation strictly speaking does not. It is possible to define a layer-based scheme from a transformation of the spatial derivatives which mimics the one done in a homogeneous medium, by replacing the Laplacian operator Δ by some Δ_{layer} inside the PML, but this layer will not be perfectly matched anymore and is called a *pseudo-PML* (pPML) [42]. In this case, reflections from the interface between Ω and the layer might not be small. In fact, the layer might even cause the solution to grow exponentially inside it, instead of forcing it to decay [19, 39]. It has been shown in [42] that a pPML for Maxwell’s equations can still lead to an accurate solution, but the layer needs to be made very thick in order to minimize reflections at the interface. Then the equations have to be solved in a very large computational domain, where most of the work will consist in solving for the pPML. In the setting where (2.1) has to be solved a large number of times, a precomputation to speed up the application of the pPML (or any other accurate but slow ABC or AL) might be of interest.

Other discrete ALs may also need to be quite wide in practice, or may be otherwise computationally costly, and this can be understood with the following physical explanation. Call L the required width (in meters) of a layer, for a certain accuracy. Although this is not a limitation of the framework presented in this paper, discretize the Helmholtz operator in the most elementary way using the standard five-point difference stencil. Put $h = 1/N$ for the grid spacing, where N is the number of points per

²For example, the half-space DtN map in a uniform medium is certainly symmetric, as can be seen from the explicit formulas (5.1) and (5.2) of section 5.

dimension for the interior problem, inside the unit square $\Omega = [0, 1]^2$. Let $w = L/h$ be the AL width in number of grid points. While Ω contains N^2 points, the total number of unknowns is $O((N + 2w)^2)$ in the presence of the layer. In a uniform medium, the AL width L needed is a fraction of the wavelength, i.e., $L \sim \lambda = \frac{2\pi}{\omega} \sim \frac{1}{N}$, so that a constant number of points is needed independently of N : $w = L/h = LN \sim 1$. However, in nonuniform media, the heterogeneity of $c(\mathbf{x})$ can limit the accuracy of the layer. Consider an otherwise uniform medium with an embedded scatterer outside of Ω ; then the AL will have to be large enough to enclose this scatterer, no matter the size of N . For more general, heterogeneous media such as the ones considered in this paper, we have often observed empirically (results not shown) that convergence as a function of L or w is delayed compared to a uniform medium. That means that $L \sim L_0$ so that $w \sim NL_0$ or $w = O(N)$, as is assumed in this paper from now on.

In the case of a second-order discretization, the rate at which one must increase N in order to preserve a constant accuracy in the solution, as ω grows, is about $N \sim \omega^{1.5}$. This unfortunate phenomenon, called the pollution effect, is well known: it begs to *increase* the resolution, or number of points per wavelength, of the scheme as ω grows [4, 5]. As was just explained, the width of the AL may be as wide as $O(N)$ points, which will scale as $O(\omega^{1.5})$ grid points.

Before presenting the approach of this paper for compressing an ABC, a straightforward way of obtaining the DtN map from an ABC is presented. It consists of eliminating the unknowns in the AL in order to obtain a reduced system on the interior nodes. This solution technique is enlightening but also computationally impractical, as will be seen.

2.3. Layer stripping for the Dirichlet-to-Neumann map. The system for the discrete Helmholtz equation is

$$(2.6) \quad \begin{pmatrix} A & P \\ P^T & C \end{pmatrix} \begin{pmatrix} u_{\text{out}} \\ u_{\Omega} \end{pmatrix} = \begin{pmatrix} 0 \\ f_{\Omega} \end{pmatrix},$$

with $A = \Delta_{\text{layer}} + k^2 I$ and $C = \Delta + k^2 I$, with Δ overloaded to denote discretization of the Laplacian operator, and Δ_{layer} the discretization of the Laplacian operator inside the AL. Eliminating the exterior unknowns u_{out} will give a new system which only depends on the interior unknowns u_{Ω} . The most obvious way of eliminating those unknowns is to form the Schur complement $S = C - P^T A^{-1} P$ of A by any kind of Gaussian elimination. For instance, in the standard raster scan ordering of the unknowns, the computational cost of this method³ is $O(w^4)$ —owing to the fact that A is a sparse banded matrix of size $O(w^2)$ and band at least $N + 2w$. Alternatively, elimination of the unknowns can be performed by layer stripping, starting with the outermost layer of unknowns from u_{out} , until the layer of points that is just outside of $\partial\Omega$ is eliminated. The computational cost will be $O(w^4)$ in this case as well. To see this, let u_j represent the vector of the solution evaluated at the set of points on the j th layer outside of Ω , so that in particular u_w is the solution at the points on the outermost layer. The system (2.6) is thus rewritten as

$$\begin{pmatrix} A_w & P_w \\ P_w^T & C_w \end{pmatrix} \begin{pmatrix} u_w \\ \vdots \end{pmatrix} = \begin{pmatrix} 0 \\ \vdots \end{pmatrix},$$

³The cost of the Schur complement procedure is dominated by that of Gaussian elimination to apply A^{-1} to P . Gaussian elimination on a sparse banded matrix of size s and band b is $O(sb^2)$, as can easily be inferred from Algorithm 20.1 of [50].

where A_w is the discretization of $\Delta_{\text{layer}} + k^2$ at points in layer w , and C_w is the discretization of either $\Delta_{\text{layer}} + k^2$ or $\Delta + k^2$ at all other points. Note that, because of the five-point stencil, P_w has nonzeros exactly on the columns corresponding to u_{w-1} . Hence the matrix $P_w^T A_w^{-1} P_w$ in the first Schur complement $S_w = C_w - P_w^T A_w^{-1} P_w$ is nonzero exactly at the entries corresponding to u_{w-1} . Then, in the next Schur complement, the matrix A_{w-1} (the block of S_w corresponding to the points u_{w-1}) to be inverted will be full. For the same reason, every matrix A_j to be inverted thereafter, for every subsequent layer to be eliminated, will be a full matrix. Hence at every step the cost of forming the corresponding Schur complement is at least on the order of the cubic power of the number of points in that layer. Hence the total cost of eliminating the exterior unknowns by layer stripping is approximately $\sum_{j=1}^w (4(N+2j))^3 = O(w^4)$.

Similar arguments can be used for the Helmholtz equation in 3 dimensions. In this case, the computational complexity of the Schur complement or layer stripping methods would be, respectively, $O(w^3(w^2)^2) = O(w^7)$ or $\sim \sum_{j=1}^w (6(N+2j)^2)^3 = O(w^7)$. Therefore, direct elimination of the exterior unknowns is quite costly. Better methods than Gaussian elimination could be modified to reduce the system and obtain D —for example, those in [20, 27, 29]. Some new insight is needed to construct the DtN map more efficiently.

Note that eliminating the exterior unknowns, whether in one pass or by layer stripping, creates a reduced system. This new system looks just like the original Helmholtz system on the interior unknowns u_Ω , except for the top left block, corresponding to u_0 the unknowns on $\partial\Omega$, which has been modified by the elimination procedure. Hence with the help of some dense matrix D the reduced, $N^2 \times N^2$ system is written as

$$(2.7) \quad \mathcal{L}u = \begin{pmatrix} (hD - I)/h^2 & I/h^2 & 0 & \cdots \\ I/h^2 & & & \\ 0 & & [\Delta + k^2 I] & \\ \vdots & & & \end{pmatrix} \begin{pmatrix} u_0 \\ u_{-1} \\ u_{-2} \\ \vdots \end{pmatrix} = \begin{pmatrix} 0 \\ f_{-1} \\ f_{-2} \\ \vdots \end{pmatrix}.$$

This creates an absorbing condition which may be used on the boundary of Ω , independent of the right-hand side f . Indeed, let u_{-1} be the solution on the first layer of points inside Ω , so that the first block row of the system gives $(I - hD)u_0 = u_{-1}$, or

$$(2.8) \quad \frac{u_0 - u_{-1}}{h} = Du_0,$$

a numerical realization of the DtN map in (2.5), using the ABC of choice. Indeed, elimination can be used to reformulate any computationally intensive ABC or AL into a realization of (2.5), by reducing any extra equations coming from the ABC or AL to relations involving only unknowns on the boundary and on the first layer (or first few layers) inside the boundary, to obtain a numerical DtN map D . A drawback is that forming this matrix D by elimination can be prohibitive, as explained above, and a dense matrix D is obtained. Instead, this paper suggests adapting the framework of matrix probing in order to obtain D in reasonable complexity, and in compressed form. Matrix probing is also at least as memory efficient as the methods in [20, 27, 29], which have a memory usage of about $N^2 \log N$, whereas the approximate D obtained through matrix probing has memory usage on the order of N^2 .

As for the complexity of solving the Helmholtz equation, reducing the ABC confers the advantage of making the number of nonzeros in the matrix \mathcal{L} of (2.7) independent of the width of the AL or complexity of the ABC. Indeed, \mathcal{L} has about $20N^2$ nonzero entries, instead of the $5N^2$ one would expect from a five-point stencil discretization of the Helmholtz equation, because D is a full $4N \times 4N$ matrix. Obtaining a fast matrix-vector product for our approximation of D is the second step of the algorithm proposed in this paper. This could reduce the application cost of L from $20N^2$ to something closer to $5N^2$. This is work the authors have completed, and that will appear elsewhere. The present paper, again, focuses on the first step of the proposed algorithm: obtaining D efficiently by adapting the framework of matrix probing. The next section presents matrix probing and other algorithms used in this article.

3. Matrix probing for the DtN map. The idea of matrix probing is that a matrix D with adequate structure can sometimes be recovered from the knowledge of a fixed, small number of matrix-vector products $Dg^{(j)}$, $j = 1, 2, \dots$, where $g^{(j)}$ are typically random vectors. In the case where D is the DtN map, each $g^{(j)}$ consists of Dirichlet data on $\partial\Omega$, and each application $Dg^{(j)}$ requires solving an exterior Helmholtz problem, as explained in section 3.1. Matrix probing is then introduced in section 3.2. For matrix probing to be an efficient expansion scheme, a careful choice of the expansion basis is necessary, as shown in section 3.3. Some remarks on traveltimes follow in section 3.4, and in section 3.5 a short note is given on using a probed DtN map inside a Helmholtz solver.

3.1. Setup for the exterior problem. The exterior problem is the heterogeneous medium Helmholtz equation at angular frequency ω , outside $\Omega = [0, 1]^2$, with Dirichlet boundary condition $u_0 = g$ on $\partial\Omega$. In the present paper, this problem is solved numerically with the five-point stencil of finite differences and a pPML, although other methods could be used. The pPML starts at a fixed, small distance away from Ω , so there is a small strip around Ω where the equations are unchanged. Recall that the width of the pPML is in general as large as $O(\omega^{1.5})$ grid points (see the end of section 2.2). Number the four sides of $\partial\Omega$ counterclockwise starting from $(0, 0)$: side 1 is the bottom edge $(x, 0)$, $0 \leq x \leq 1$, side 2 is the right edge, etc. The exterior DtN map \mathcal{D} is defined from $\partial\Omega$ to $\partial\Omega$. Thus its numerical realization D has a 4×4 block structure and is $4N \times 4N$. As an integral kernel, \mathcal{D} has singularities at the junctions between these blocks (due to the fact that $\partial\Omega$ has corners), and this feature is taken into account by probing D block by block. Such a generic block is called M or the (j_M, ℓ_M) block of D , referring to its indices in the 4×4 block structure.

The method by which the system for the exterior problem (and the original problem in section 3.5) is solved is immaterial in the scope of this paper, though for reference the experiments in this paper use UMFPACK's sparse direct solver [16]. For treating large problems, a better solver should be used, such as those in [23, 24, 25, 28, 29, 48].

For a given Dirichlet boundary condition g , solve the system and obtain a solution u in the exterior computational domain. In particular, consider u_1 , the solution on the layer of points just outside of $\partial\Omega$. Let

$$(3.1) \quad \frac{u_1 - g}{h} = Dg.$$

This relation is not exactly the same as in (2.8), but needs not be interpreted as a first-order approximation of the continuous DtN map in (2.8). The matrix D in (3.1) is an algebraic object of interest that will be "probed" from repeated applications

to different vectors g . This D is then used (see section 3.5) to close the Helmholtz system on Ω by eliminating the ghost points on the layer just outside $\partial\Omega$, where u_1 is defined.

For probing the (j_M, ℓ_M) block M of D , one needs matrix-vector products of D with vectors g of the form $[z, 0, 0, 0]^T$, $[0, z, 0, 0]^T$, etc., to indicate that the Dirichlet boundary condition is z on the side indexed by ℓ_M , and zero on the other sides. The application Dg is then restricted to side j_M .

3.2. Matrix probing. The dimensionality of D needs to be limited for recovery from a few $Dg^{(j)}$ to be possible, but matrix probing is *not* an all-purpose low-rank approximation technique. Instead, it is the property that D has an efficient representation in some adequate preset basis that makes recovery from probing possible. As opposed to the randomized SVD method, which requires the number of matrix-vector applications to be greater than the rank [33], matrix probing can recover interesting $N \times N$ operators from a single matrix-vector application, even when the number of matrix parameters approaches N [12, 17]. This technique has also been investigated in [44, 43], where matrices are expanded in dictionaries of time-frequency shifts.

We now describe a model for M , any $N \times N$ block of D , that will sufficiently lower its dimensionality to make probing possible. Assume M can be written as

$$(3.2) \quad M \approx \sum_{j=1}^p c_j B_j,$$

where the B_j 's are fixed, known basis matrices, which need to be chosen carefully in order to give an accurate approximation of M . For illustration, let B_j be a discretization of the integral kernel

$$(3.3) \quad \frac{e^{ik|x-y|}}{(h + |x-y|)^{j/2}},$$

where $h = 1/N$ is the discretization parameter. Choices of basis matrices and their rationales will be discussed in section 3.3.

For now, note that the advantage of the specific choice of basis matrices (3.3), and its generalizations explained in section 3.3, is that it results in accurate expansions with a number of basis matrices p “essentially independent” of N , namely, with a p that grows either logarithmically in N , or at most like a very sublinear fractional power law (such as $N^{0.12}$; see section 4.4). This scaling is much better than that of the layer width, $w = O(N)$ grid points, mentioned earlier in section 2.2. The form of B_j suggested in (3.3) is also confirmed by the fact that it provides a great expansion basis for the uniform-medium half-space DtN map in \mathbb{R}^2 , as shown in section 5.

To recover the matrix probing expansion coefficients c_j , let z be a Gaussian independent and identically distributed random vector of N entries, with mean 0 and variance 1 (other choices are possible; see [12]). Applying the expansion (3.2) to z gives

$$(3.4) \quad v = Mz \approx \sum_{j=1}^p c_j B_j z = \Psi_z \mathbf{c}.$$

Multiplying this equation on the left by the pseudoinverse of the $N \times p$ matrix Ψ_z will give an approximation to \mathbf{c} , the coefficient vector for the expansion (3.2) of M . Note that both Ψ_z and the resulting coefficient vector \mathbf{c} depend on the vector z .

In order to improve the conditioning in taking the pseudoinverse of Ψ_z and reduce the error in the coefficient vector \mathbf{c} , one may use $q > 1$ random realizations of M , that is, $v^{(1)} = Mz^{(1)}, \dots, v^{(q)} = Mz^{(q)}$. Then \mathbf{v} will be a long column vector of Nq entries containing the concatenation of the $v^{(j)}$'s; Ψ will similarly⁴ be the concatenation of the $\Psi_z^{(j)}$'s and will have size $Nq \times p$. Now the solution \mathbf{c} of $\mathbf{v} = \Psi\mathbf{c}$ can be obtained. There is a limit to the range of p for which this system is well-posed: [12] covers the precise conditions on p , N , and the following two parameters, called “weak condition numbers,” for which recoverability of \mathbf{c} is accurate with high probability.

DEFINITION 3.1 (weak condition number λ).

$$\lambda = \max_j \frac{\|B_j\|_2 \sqrt{N}}{\|B_j\|_F}.$$

DEFINITION 3.2 (weak condition number κ).

$$\kappa = \text{cond}(\mathbf{B}), \quad \mathbf{B}_{j\ell} = \text{Tr}(B_j^T B_\ell).$$

Note that a small λ translates into a high-rank condition on the basis matrices [12]. A small κ translates into a necessary condition for the set of basis matrices to be a *Riesz basis* (see equation (7) of [15], where constants A, B can be taken, respectively, to be the minimum and maximum singular values of \mathbf{B}). Having small weak condition numbers will guarantee a small failure probability of matrix probing and a bound on the condition number of Ψ_z or Ψ , while allowing for a higher p and thus potentially higher accuracy. These results are contained in the following theorem.

THEOREM 3.3 (Chiu and Demanet [12]). *Let z be a Gaussian independent and identically distributed random vector of length qN , and let Ψ be as defined above. Choose a value of p which is not too large, namely, which satisfies*

$$(3.5) \quad qN \geq C p (\kappa \lambda \log N)^2$$

for some number $C > 1$ independent of all other quantities. Then $\text{cond}(\Psi) \leq 2\kappa + 1$ with probability at least $1 - 2C_2 p N^{1-\alpha}$, where $\alpha = \frac{C}{2e\sqrt{2}}$.

The constant C_2 in Theorem 3.3 might be large, but $N^{-\alpha}$ decays very fast with larger N and C . Theorem 3.3 in turn implies the accuracy of matrix probing. Indeed, as is shown in Proposition 1.5 of [12], a small error in truncating to p the expansion of M in the probing basis leads to a small error between the probing expansion $\sum^p c_j B_j$ and M , which leads to small errors in the coefficients \mathbf{c} . Also, Theorem 3.3 says that using $q > 1$ allows one to use a larger p , leading to a smaller truncation error of M , hence greater accuracy of $\sum^p c_j B_j$. In a nutshell, recovery of \mathbf{c} works under mild assumptions on the B_j 's, when p is a small fraction of Nq up to log factors.⁵

As noted previously, the work necessary for probing M is on the order of q solves of the original problem. Indeed, computing $Mz^{(1)}, \dots, Mz^{(q)}$ means solving q times the exterior problem with the AL. This is roughly equivalent to solving the original

⁴Note the use of Ψ_z and v when a single random realization of M is employed, and the use of Ψ and \mathbf{v} when $q > 1$ random realizations are employed.

⁵It is difficult to use Theorem 3.3 and [12, Proposition 1.5] to find appropriate values of p, q, N . In practice, first pick an N which makes sense with the desired accuracy of the Helmholtz equation solution u . Then choose an affordable q and obtain the products $Mz^{(\ell)}$. Finally, pick a small p , obtain the coefficients \mathbf{c} , and test for the error in M . Increase p if needed, until the error is satisfactory or the approximation gets worse (p is too large). If p is too large, increase q and start again with a lower p . See section 4 for details on finding or approximating the error in M .

Helmholtz problem with the AL q times, assuming the AL width w is at least as large as N . Then computing the qp products of the p basis matrices with the q random vectors totals at most qpN^2 work, or less if the basis matrices have a fast matrix-vector product. Finally, computing the pseudoinverse of Ψ has cost Nqp^2 [50]. Hence, as long as $p, q \ll N$, the dominant cost of matrix probing⁶ comes from solving q times the exterior problem with a random Dirichlet boundary condition. In our experiments, $q = O(1)$ and p can be as large as a few hundred for high accuracy.

Note that the information from the q solves can be reused for any other block which is in the same block column as M . However, if it is needed to probe blocks of D which are not all in the same block column, then another q solves need to be performed, with a Dirichlet boundary condition on the appropriate side of $\partial\Omega$. This of course increases the total number of solves needed.⁷

3.3. Choice of basis matrices. The essential information of the DtN map needs to be summarized in broad strokes in the basis matrices, with the details of the numerical fit left to the probing procedure. In the case of D , most of its physics is contained in its *oscillations* and *diagonal singularity*, as predicted by geometrical optics. A heuristic argument to obtain the form of D starts from the Green's formula (2.4), differentiating it one more time in the normal direction. After accounting for the correct jump condition, an alternative Steklov–Poincaré identity is obtained, where \mathcal{T}^* denotes the adjoint of the operator \mathcal{T} defined in (2.3), namely,

$$\mathcal{D} = \left(\mathcal{T}^* + \frac{1}{2}\mathcal{I} \right)^{-1} \mathcal{H},$$

where \mathcal{H} is the hypersingular integral operator with kernel $\partial_{\nu_{\mathbf{x}}}\partial_{\nu_{\mathbf{y}}}G$, where $G(\mathbf{x}, \mathbf{y})$ is the Green's function and $\nu_{\mathbf{x}}, \nu_{\mathbf{y}}$ are the normals to $\partial\Omega$ in \mathbf{x} and \mathbf{y} , respectively. The presence of $(\mathcal{T}^* + \frac{1}{2}\mathcal{I})^{-1}$ is inconsequential to the form of \mathcal{D} , as it involves solving a well-posed second-kind integral equation. As a result, the oscillations and singularity of \mathcal{D} are qualitatively similar to those of \mathcal{H} . (The exact construction of \mathcal{D} from \mathcal{G} is of course already known in a few special cases, such as the uniform-medium half-space problem considered in section 5.)

In turn, geometrical optics reveals the form of G . In a context where there is no multipathing, that is, where there is a single traveltime $\tau(\mathbf{x}, \mathbf{y})$ between any two points $\mathbf{x}, \mathbf{y} \in \Omega$, one may write a high- ω asymptotic series for G as

$$(3.6) \quad G(\mathbf{x}, \mathbf{y}) \sim e^{i\omega\tau(\mathbf{x}, \mathbf{y})} \sum_{j \geq 0} a_j(\mathbf{x}, \mathbf{y}) \omega^{-j},$$

$\tau(\mathbf{x}, \mathbf{y})$ is the traveltime between points \mathbf{x} and \mathbf{y} , found by solving the Eikonal equation

$$(3.7) \quad \|\nabla_{\mathbf{x}}\tau(\mathbf{x}, \mathbf{y})\| = \frac{1}{c(\mathbf{x})},$$

⁶In this paper we also perform a QR factorization on the basis matrices (see section 3.3), and this has cost N^2p^2 . This operation has a cost similar to or smaller than the cost of an exterior solve using current Helmholtz solvers. It might also be possible to avoid the QR factorization if a set of basis matrices closer to orthonormal is used.

⁷Another option would be to probe all of D at once, using basis matrices that have the same size as D , but are 0 except on the support of each distinct block in turn. In this case, κ remains 1 if the basis matrices are orthonormalized, but N becomes $4N$ in (3.5), and λ doubles. Hence a higher value of q might be needed for the same p , or the same accuracy. This approach, which might become more advantageous for a more complicated polygonal domain, is not investigated in the present paper.

and the amplitudes a_j satisfy transport equations. The derivation of geometrical optics is classical; see [51, 22]. In the case of multipathing (multiple possible traveltimes between any two points), the representation (3.6) of G becomes instead

$$G(\mathbf{x}, \mathbf{y}) \sim \sum_j e^{i\omega\tau_j(\mathbf{x}, \mathbf{y})} \sum_{\ell \geq 0} a_{j\ell}(\mathbf{x}, \mathbf{y}) \omega^{-\ell},$$

where the τ_j 's each obey (3.7) away from caustic curves. The amplitudes $a_{j\ell}(\mathbf{x}, \mathbf{y})$ are singular at caustic curves in addition to the diagonal $\mathbf{x} = \mathbf{y}$ and contain the information of the Maslov indices. Traveltimes are symmetric: $\tau_j(\mathbf{x}, \mathbf{y}) = \tau_j(\mathbf{y}, \mathbf{x})$.

The singularity of the Green's function in a uniform medium is $O(\log|\mathbf{x} - \mathbf{y}|)$ in 2D and $O(|\mathbf{x} - \mathbf{y}|^{-1})$ [11, 1]. Thus the amplitude factor in (3.6) for a uniform medium, at $\mathbf{x} = \mathbf{y}$, is $O(\log|\mathbf{x} - \mathbf{y}|)$ in 2D and $O(|\mathbf{x} - \mathbf{y}|^{-1})$ in 3D, and similarly in smooth heterogeneous media.⁸ After differentiating \mathcal{G} twice to obtain the kernel \mathcal{H} , the homogeneity on the diagonal becomes $O(|\mathbf{x} - \mathbf{y}|^{-2})$ in 2D and $O(|\mathbf{x} - \mathbf{y}|^{-3})$ in 3D. For the decay at infinity, the scalings are different and can be obtained from Fourier analysis of square root singularities; the kernel of \mathcal{H} decays like $O(|\mathbf{x} - \mathbf{y}|^{-3/2})$ in 2D, and $O(|\mathbf{x} - \mathbf{y}|^{-5/2})$ in 3D. In between, the amplitude is smooth as long as the traveltime is single-valued.

For all these reasons, we define the basis matrices B_j as follows. Assume τ is single-valued. In 1D, denote the tangential component of \mathbf{x} by x , and similarly that of \mathbf{y} by y , in coordinates local to each edge with $0 \leq x, y \leq 1$. Each block M of D relates to a couple of edges of the square domain. Let $\mathbf{j} = (j_1, j_2)$ with j_1, j_2 nonnegative integers. First consider the general forms

$$(3.8) \quad \beta_{\mathbf{j}}(x, y) = e^{i\omega\tau(x, y)} (h + |x - y|)^{-\frac{j_1}{\alpha}} (h + \gamma(x, y))^{-\frac{j_2}{\alpha}}$$

and

$$(3.9) \quad \beta_{\mathbf{j}}(x, y) = e^{i\omega\tau(x, y)} (h + |x - y|)^{-\frac{j_1}{\alpha}} P_{j_2}(\gamma(x, y)),$$

where P_n is the Legendre polynomial of degree n . These forms have three factors: oscillations $e^{i\omega\tau(x, y)}$, inverse powers of $h + |x - y|$ to account for the expected decay of D across diagonals, and a function of γ , where $\gamma(x, y)$ depends on the block and medium of interest, to account for the behavior of D along diagonals. The more favorable choices for γ respect the singularities created at the vertices of the square; typically $\gamma(x, y) = 1 - |x + y - 1|$, $\gamma(x, y) = \min(x, y, 1 - x, 1 - y)$, or $\gamma(x, y) = \min(x, y)$. The parameter α can be taken to be 2, a good choice in view of the numerics of section 5.2 and in the light of the asymptotic behaviors on the diagonal and at infinity discussed in the previous paragraph.

If several traveltimes are needed, then different sets of $\beta_{\mathbf{j}}$ are defined for each traveltime. For example, if the four traveltimes $+\tau_1, -\tau_1, +\tau_2, -\tau_2$ are to be used (more details on this in section 3.4), then the following forms could be considered:

$$\left\{ \left\{ \frac{e^{i\omega\tau_1}}{(h + |x - y|)^{-\frac{j_1}{\alpha}} (h + \gamma(x, y))^{-\frac{j_2}{\alpha}}} \right\}_{\mathbf{j}}, \left\{ \frac{e^{-i\omega\tau_1}}{(h + |x - y|)^{-\frac{j_1}{\alpha}} (h + \gamma(x, y))^{-\frac{j_2}{\alpha}}} \right\}_{\mathbf{j}}, \right. \\ \left. \left\{ \frac{e^{i\omega\tau_2}}{(h + |x - y|)^{-\frac{j_1}{\alpha}} (h + \gamma(x, y))^{-\frac{j_2}{\alpha}}} \right\}_{\mathbf{j}}, \left\{ \frac{e^{-i\omega\tau_2}}{(h + |x - y|)^{-\frac{j_1}{\alpha}} (h + \gamma(x, y))^{-\frac{j_2}{\alpha}}} \right\}_{\mathbf{j}} \right\}.$$

⁸The careful analysis of the remainder of a geometrical optics expansion [51] or progressive wave expansion [14] in smooth heterogeneous media would be a way to make this statement precise.

The basis matrices B_ℓ are then obtained by first discretizing the above forms and then applying a QR factorization, where orthogonality between two matrices A and B is defined in the sense of the Frobenius inner product $\langle A, B \rangle = \text{Tr}(AB^T)$. This automatically sets the weak condition number κ of probing to 1. In many of our test cases it appears that the “triangular” condition $j_1 + 2j_2 < \text{const}$ works well. The number of couples (j_1, j_2) satisfying this relation will be p/T , where p is the number of basis matrices in the matrix probing algorithm and T is the number of distinct traveltimes. The ordering of the basis matrices B_ℓ respects the increase of $j_1 + 2j_2$.

3.4. Traveltimes. Determining the traveltime(s) $\tau(\mathbf{x}, \mathbf{y})$ is the more “supervised” part of this method, but is needed to keep the number p of parameters small in the probing expansion. A few different scenarios can arise.

1. In the case when $\nabla c(\mathbf{x})$ is locally perpendicular to a straight segment of the boundary, this segment is itself a ray and the waves are called interfacial, or “creeping.” The direct traveltime between any two points \mathbf{x} and \mathbf{y} on this segment is then given by the line integral of $1/c(\mathbf{x})$. An infinite sequence of additional interfacial waves results from successive reflections at the endpoints of the segment. Consider points \mathbf{x}, \mathbf{y} on the same side of $\partial\Omega$ —for illustration, let $\mathbf{x} = (x, 0)$ and $\mathbf{y} = (y, 0)$ on the bottom side of $\Omega = [0, 1]^2$, with $x \leq y$ (this is sufficient since traveltimes are symmetric). Then traveltimes are predicted as follows:

$$\begin{aligned}\tau_1(\mathbf{x}, \mathbf{y}) &= \int_x^y \frac{1}{c(t, 0)} dt, \\ \tau_2(\mathbf{x}, \mathbf{y}) &= \tau_1(\mathbf{x}, \mathbf{y}) + 2 \min \left(\int_0^x \frac{1}{c(t, 0)} dt, \int_y^1 \frac{1}{c(t, 0)} dt \right), \\ \tau_3(\mathbf{x}, \mathbf{y}) &= \tau_1(\mathbf{x}, \mathbf{y}) + 2 \max \left(\int_0^x \frac{1}{c(t, 0)} dt, \int_y^1 \frac{1}{c(t, 0)} dt \right) \\ &= 2 \int_0^1 \frac{1}{c(t, 0)} dt - \tau_2(\mathbf{x}, \mathbf{y}), \\ \tau_4(\mathbf{x}, \mathbf{y}) &= 2 \int_0^1 \frac{1}{c(t, 0)} dt - \tau_1(\mathbf{x}, \mathbf{y}), \quad \text{etc.}\end{aligned}$$

All of the first four traveltimes can be expressed as a sum of $\pm\tau_1$, $\pm\tau_2$, and the constant phase $2 \int_0^1 (c(t, 0))^{-1} dt$, which does not depend on \mathbf{x} or \mathbf{y} . In fact, any subsequent traveltime corresponding to traveling solely along the bottom boundary of $\partial\Omega$ will be again a combination of those quantities. This means that using $\pm\tau_1$ and $\pm\tau_2$ in the basis matrices captures all the traveltimes relative to a single side, which can help to achieve higher accuracy when probing the diagonal blocks of D .

This analysis can be adapted to deal with creeping waves that start on one side of the square and terminate on another side, which is important for the nondiagonal blocks of D . For example, let us consider a point $\mathbf{x} = (x, 0)$ on the bottom side of $\Omega = [0, 1]^2$ and a point $\mathbf{y} = (1, y)$ on the right side of Ω . The first few traveltimes between \mathbf{x} and \mathbf{y} , assuming that those traveltimes correspond to creeping waves that only travel along the bottom or right side of $\partial\Omega$, are then as follows:

$$\begin{aligned}\tau_1(\mathbf{x}, \mathbf{y}) &= \int_x^1 \frac{1}{c(t, 0)} dt + \int_0^y \frac{1}{c(1, t)} dt, \\ \tau_2(\mathbf{x}, \mathbf{y}) &= \tau_1(\mathbf{x}, \mathbf{y}) + 2 \min \left(\int_0^x \frac{1}{c(t, 0)} dt, \int_y^1 \frac{1}{c(1, t)} dt \right),\end{aligned}$$

$$\begin{aligned}
\tau_3(\mathbf{x}, \mathbf{y}) &= \tau_1(\mathbf{x}, \mathbf{y}) + 2 \max \left(\int_0^x \frac{1}{c(t, 0)} dt, \int_y^1 \frac{1}{c(1, t)} dt \right) \\
&= 2 \left(\int_0^1 \frac{1}{c(t, 0)} dt + \int_0^1 \frac{1}{c(1, t)} dt \right) - \tau_2(\mathbf{x}, \mathbf{y}), \\
\tau_4(\mathbf{x}, \mathbf{y}) &= 2 \left(\int_0^1 \frac{1}{c(t, 0)} dt + \int_0^1 \frac{1}{c(1, t)} dt \right) - \tau_1(\mathbf{x}, \mathbf{y}), \quad \text{etc.}
\end{aligned}$$

Once again, all of the first four traveltimes are expressed as a sum of $\pm\tau_1$, $\pm\tau_2$, and the constant phase $2(\int_0^1(c(t, 0))^{-1} dt + \int_0^1(c(1, t))^{-1} dt)$. In fact, any subsequent traveltime corresponding to traveling solely along the bottom or right boundary of $\partial\Omega$ will be a combination of those quantities. This means that using $\pm\tau_1$ and $\pm\tau_2$ in the basis matrices captures all the traveltimes between the two adjacent sides, which can help achieve higher accuracy for probing nondiagonal blocks of D .

2. When $c(\mathbf{x})$ increases outward in a smooth fashion, there may be body waves going off into the exterior and coming back to $\partial\Omega$. The traveltime for these waves needs to be solved either by a Lagrangian method (solving the ordinary differential equation for the rays) or by an Eulerian method (solving the Eikonal equation shown earlier). In this paper we use the fast marching method of Sethian [45] to deal with these waves in the case that we label “slow disk” in section 4.1.3.

3. In the case when $c(\mathbf{x})$ has singularities in the exterior domain, each additional reflection creates a traveltime that should (ideally) be predicted. Such is the case of the “diagonal fault” example introduced in the next section, where a straight jump discontinuity of $c(\mathbf{x})$ intersects $\partial\Omega$ at a nonnormal angle. If points \mathbf{x} and \mathbf{y} on a side of the boundary intersected by the jump discontinuity in $c(\mathbf{x})$ are such that there is a path leaving the boundary at \mathbf{x} , reflecting off of the discontinuity and coming back to the boundary at \mathbf{y} , then the corresponding traveltime can be constructed. This is now a body traveltime, as opposed to a creeping traveltime, and is called $\tau'_2(\mathbf{x}, \mathbf{y})$. Note again that $\tau'_2(\mathbf{x}, \mathbf{y})$ only exists for some pairs of points on the boundary. In the case of the “diagonal fault,” we found empirically that replacing the quantity τ_2 that was described previously by τ'_2 where τ'_2 exists (calling this new traveltime $\tilde{\tau}_2$) increased accuracy by an order of magnitude, as mentioned in the numerical results of section 4.1.5. The intuition behind the replacement of τ_2 with $\tilde{\tau}_2$ is that the reflections of creeping waves off of corners of Ω , whose traveltimes are in τ_2 , should be weaker than the body waves reflected from the fault, whose traveltimes are in τ'_2 . Using the 6 traveltimes $\pm\tau_1, \pm\tau_2, \pm\tau'_2$ where they are defined did not improve upon using simply $\pm\tau_1, \pm\tilde{\tau}_2$.

In every case described above, except when using the fast marching method [45], traveltimes were computed to an accuracy of both six and ten digits, and we did not notice a difference in the probing results. We thus used an accuracy of ten digits for these traveltimes. Because of the complexity of the fast marching method, we were only able to compute traveltime using this method to an accuracy of three digits.

3.5. Solving the Helmholtz equation with a compressed ABC. Once approximations M of each block M are obtained through matrix probing, we construct block by block the approximation \tilde{D} of D and use it in a solver for the Helmholtz equation on the domain $\Omega = [0, 1]^2$. To close the system at the boundary, we used the method of ghost points alluded to earlier and eliminated the ghost points on the layer just outside Ω using the relationship (3.1).

4. Numerical experiments. Our benchmark media $c(\mathbf{x})$ are as follows (color figures available in online version):

1. a uniform wave speed of 1, $c(\mathbf{x}) \equiv 1$ (Figure 1),
2. a “Gaussian waveguide” (Figure 2),
3. a “Gaussian slow disk” (Figure 3) large enough to encompass Ω (this will cause some waves going out of Ω to come back in),
4. a “vertical fault” (Figure 4),
5. a “diagonal fault” (Figure 5),
6. and a discontinuous “periodic medium” (Figure 6). The periodic medium consists of square holes of velocity 1 in a background of velocity $1/\sqrt{12}$.

The media are continued in the obvious way (i.e., they are *not* put to a constant) outside the domain shown in the figures. The outline of the $[0, 1]^2$ box is in black.




FIG. 1. Color plot of $c(\mathbf{x})$ for the uniform medium.




FIG. 3. Color plot of $c(\mathbf{x})$ for the Gaussian slow disk.




FIG. 5. Color plot of $c(\mathbf{x})$ for the diagonal fault.




FIG. 2. Color plot of $c(\mathbf{x})$ for the Gaussian waveguide.




FIG. 4. Color plot of $c(\mathbf{x})$ for the vertical fault.




FIG. 6. Color plot of $c(\mathbf{x})$ for the periodic medium.

We use a Helmholtz equation solver to estimate the relative error in the Helmholtz equation solution caused by the finite difference discretization (the *FD error*⁹) and the error caused by using the specified pPML width.¹⁰ Those errors are presented in Table 1, with some parameters used in this section, including the position of the point source f . Whenever possible, we use an AL with error smaller than the precision we seek with matrix probing, so with a width w greater than that showed in Table 1. This makes probing easier, i.e., p and q can be smaller.

Note that some blocks in D are the same up to transpositions or flips (inverting the order of columns or rows) if the medium $c(\mathbf{x})$ has symmetries.

DEFINITION 4.1 (multiplicity of a block of D). *Let M be the (j_M, ℓ_M) block of D .*

⁹To find this FD error, use a large pPML and compare the solution u for different values of N . The FD error is the relative ℓ_2 error in u inside Ω .

¹⁰To obtain the error caused by the absorbing layer, fix N and compare the solution u for different layer widths w to calculate the relative ℓ_2 error in u inside Ω .

TABLE 1

For each medium considered, we show parameters N and $\omega/2\pi$, along with the discretization error caused by the finite difference (FD error) formulation. We also show the width w of the pPML needed, in number of points, to obtain an error caused by the pPML of less than $1e-1$, and the total number of basis matrices (P) needed to probe the entire DtN map with an accuracy of about $1e-1$ as found in section 4.1. Finally, we show the position of the point source (Source).

Medium	N	$\omega/2\pi$	FD error	w	P	Source
$c(\mathbf{x}) \equiv 1$	1024	51.2	2.5e-01	4	8	(0.5, 0.25)
Waveguide	1024	51.2	2.0e-01	4	56	(0.5, 0.5)
Slow disk	1024	51.2	1.8e-01	4	43	(0.5, 0.25)
Vertical fault, left source	1024	51.2	1.1e-01	4	48	(0.25, 0.5)
Vertical fault, right source	1024	51.2	2.2e-01	4	48	(0.75, 0.5)
Diagonal fault	1024	51.2	2.6e-01	256	101	(0.5, 0.5)
Periodic medium	320	6	1.0e-01	1280	792	(0.5, 0.5)

The multiplicity $n(M)$, sometimes written as $n((j_M, \ell_M))$, of block M is the number of copies of M appearing in D , up to transpositions or flips.

Only the distinct blocks of D need to be probed. Whenever possible, out of all the blocks that are copies of each other up to transpositions or flips, it is best to choose representative blocks to be probed so as to minimize the number of distinct block-columns of D that contain those representatives, because this in turn will minimize the total number of solves Q (Definition 4.7) that will be needed. Once a block M is chosen, we may calculate its *true probing coefficients* if we have access to M directly (if not, see the remarks after Definition 4.4), given basis matrices $\{B_j\}$.

DEFINITION 4.2 (true probing coefficients of block M). *Let M be a block of D , corresponding to the restriction of D to two sides of $\partial\Omega$. Assume orthonormal probing basis matrices $\{B_j\}$. The true coefficients c_j^t in the probing expansion of M are the inner products $c_j^t = \langle B_j, M \rangle$.*

We may now introduce the p -term approximation error for the block M . Because the blocks on the diagonal of D have a singularity, and we expect the magnitude of entries in D to decay away from the diagonal, the Frobenius norm of blocks on the diagonal can be a few orders of magnitude greater than that of other blocks, and so it is more important to approximate those blocks well. This is why the error is considered relative to D , not to the block M , in the p -term approximation error below. Also, multiplying the error by the square root of the multiplicity of the block gives a better idea of how big the total error on D will be.

DEFINITION 4.3 (p -term approximation error of block M). *Let M be a block of D , corresponding to the restriction of D to two sides of $\partial\Omega$. For orthonormal probing basis matrices $\{B_j\}$, let c_j^t be the true coefficients in the probing expansion of M . Let $M_p = \sum_{j=1}^p c_j^t B_j$ be the probing p -term approximation to M . The p -term approximation error for M is*

$$(4.1) \quad \sqrt{n(M)} \frac{\|M - M_p\|_F}{\|D\|_F},$$

using the matrix Frobenius norm.

For brevity, we shall refer to (4.1) simply as the approximation error when it is clear from the context what M , p , $\{B_j\}$, and D are.

Then, using matrix probing, we will recover a coefficient vector \mathbf{c} close to \mathbf{c}^t , which gives an approximation $\tilde{M} = \sum_{j=1}^p c_j B_j$ to M . We now define the *probing*

error (which depends on q and the random vectors used) for the block M .

DEFINITION 4.4 (probing error of block M). Let $\mathbf{c} = (c_1, c_2, \dots, c_p)$ be the probing coefficients for M obtained with q random realizations $z^{(1)}$ through $z^{(q)}$. Let $\widetilde{M} = \sum_{j=1}^p c_j B_j$ be the probing approximation to M . The probing error of M is

$$(4.2) \quad \sqrt{n(M)} \frac{\|M - \widetilde{M}\|_F}{\|D\|_F}.$$

In order to get a point of reference for the accuracy benchmarks, for small problems only, the actual matrix M is computed explicitly by solving the exterior problem $4N$ times using the standard basis as Dirichlet boundary conditions, and from this we can calculate (4.2) exactly. For larger problems, the only access to M might be through a black box that outputs the product of M with some input vector by solving the exterior problem. We can then estimate (4.2) by comparing the products of M and \widetilde{M} with a few random vectors different from those used in matrix probing.

Again, for brevity, we refer to (4.2) as the probing error when other parameters are clear from the context. Once all distinct blocks of D have been probed, we can consider the *total probing error*.

DEFINITION 4.5 (total probing error). The total probing error is defined as the total error made on D by concatenating all probed blocks \widetilde{M} to produce an approximate \widetilde{D} :

$$(4.3) \quad \frac{\|D - \widetilde{D}\|_F}{\|D\|_F}.$$

Distinct blocks might of course require completely different basis matrices, and also a different number of basis matrices, to achieve a certain accuracy. Hence we define the *total number of basis matrices* P .

DEFINITION 4.6 (total number of basis matrices P). Given \widetilde{D} obtained from concatenating probed blocks \widetilde{M} , the total number of basis matrices P is the sum, over all distinct blocks \widetilde{M} of \widetilde{D} , of the number of basis matrices p used for probing M .

Probing all the distinct blocks might also require more than q solves of the exterior problem, in particular if distinct blocks are in distinct block-columns of D . Hence we define the *total number of solves* Q .

DEFINITION 4.7 (total number of solves Q). Consider the subset of the four block-columns of D that contain the chosen representative blocks of D that will be probed. The total number of solves Q is defined as the sum, over this subset of block-columns, of the maximum number of the exterior problem solves needed for each distinct block of D contained in that block-column.

For example, in the uniform-medium case, because of symmetries, D has three distinct blocks, which can all be taken in the same block-column: $(1, 1)$, $(2, 1)$, and $(3, 1)$. Hence those can all be probed using the same random vectors and thus the same exterior solves, and Q would be simply the maximum of the number q of random vectors needed for each block. This is in contrast to a generic medium where all blocks on or below the diagonal would need to be probed, and so Q would be the sum of the maximum number of solves needed for each of the four block-columns.

Presented in section 4.1 are results on the approximation and probing errors, along with related condition numbers, and then it is verified that using \widetilde{D} as an ABC gives an accurate solution to the Helmholtz equation using the *solution error from probing*.

DEFINITION 4.8 (solution error from probing). *Once we have obtained an approximation \tilde{D} to D from probing the distinct blocks of D , we may use this \tilde{D} as an ABC in a Helmholtz solver to obtain an approximate solution \tilde{u} and compare that to the true solution u using D in the solver. The solution error from probing is the ℓ_2 error on u inside Ω :*

$$(4.4) \quad \frac{\|u - \tilde{u}\|_2}{\|u\|_2} \text{ in } \Omega.$$

4.1. Probing tests. As seen in section 3.2, randomness plays a role in the value of $\text{cond}(\Psi)$ and the probing error. When we show plots for those quantities, we show 10 trials for each value of q used. The error bars display the minimum and maximum of the quantity over the 10 trials, and the line is plotted through the average value.

4.1.1. Uniform medium. For a uniform medium ($c(\mathbf{x}) \equiv 1$, Figure 1) we have three blocks with the following multiplicities: $n((1, 1)) = 4$ (same edge), $n((2, 1)) = 8$ (neighboring edges), and $n((3, 1)) = 4$ (opposite edges). Note that we do not present results for the $(3, 1)$ block: this block has negligible Frobenius norm¹¹ compared to D . Regarding the conditioning of blocks $(1, 1)$ and $(2, 1)$, as expected $\kappa = 1$, λ remains on the order of 50, and $\text{cond}(\Psi)$ increases by a few orders of magnitude as p increases for a fixed q and N (results not shown). This will affect probing in terms of the accuracy (taking the pseudoinverse will introduce larger errors in \mathbf{c}). We notice these two phenomena in Figure 7, where we show the approximation and probing errors in probing the $(1, 1)$ block for various p and q . As expected, as p increases, the probing error (always larger than the approximation error) becomes farther away from the approximation error, but using a larger q allows us to use a larger p to attain smaller probing errors. To obtain these satisfactory results, it was sufficient to use only the first arrival traveltime (as described in section 3.4) in the construction of the basis matrices. However, in order to reach higher accuracies, it was necessary to use all of the first four arrival traveltimes as described in the first scenario of section 3.4. These higher accuracy results are not shown here, but do appear as the last two rows of Table 2 (results of rows 2 to 5 of Table 2 are from Figure 7).

4.1.2. The waveguide. For a waveguide as a velocity field (see Figure 2), we have 5 distinct blocks: $n((1, 1)) = 2$, $n((2, 2)) = 2$, $n((2, 1)) = 8$, $n((3, 1)) = 2$, $n((4, 2)) = 2$. Note that block $(2, 2)$ will be very similar to block $(1, 1)$ from the uniform-medium case, and so here again using the first two traveltimes (thus four different types of oscillations, as explained in section 3.3) allows us to obtain higher accuracies. Also, blocks $(3, 1)$ and $(4, 2)$ have negligible Frobenius norm compared to D . Hence we only show results for the approximation and probing errors of blocks $(1, 1)$ and $(2, 1)$, in Figure 8. Results for using probing in a solver can be found in section 4.2.

4.1.3. The slow disk. Next, we consider the slow disk of Figure 3. Here, we have a choice to make for the traveltime upon which the oscillations depend. We may consider interfacial waves, traveling in straight line segments along $\partial\Omega$, with traveltime τ . There is also the first arrival time of body waves, τ_f , which for some points on $\partial\Omega$ involve taking a path that goes away from $\partial\Omega$, into the exterior where

¹¹We can use probing with $q = 1$ and a single basis matrix (a constant multiplied by the correct oscillations) and have a probing error of less than 10^{-6} for that block. Again, we expect blocks further from the diagonal of D to have smaller norm than blocks on or near the diagonal.




FIG. 7. Approximation error (line) and probing error (with markers) for the blocks of D , $c(\mathbf{x}) \equiv 1$. Circles are for $q = 3$, squares for $q = 5$, stars for $q = 10$.




FIG. 8. Approximation error (line) and probing error (with markers) for the blocks of D , $c(\mathbf{x})$ is the waveguide. Circles are for $q = 3$, squares for $q = 5$, stars for $q = 10$.




FIG. 9. Probing error for the (1,1) block of D , $c(\mathbf{x})$ is the slow disk, comparing the use of the normal travelttime (circles) to the fast marching travelttime (squares).




FIG. 10. Probing error for the (2,1) block of D , $c(\mathbf{x})$ is the slow disk, comparing the use of the normal travelttime (circles) to the fast marching travelttime (squares).




FIG. 11. Approximation error (line) and probing error (with markers) for the blocks of D , $c(\mathbf{x})$ is the vertical fault. Circles are for $q = 3$, squares for $q = 5$, stars for $q = 10$.




FIG. 12. Approximation error (line) and probing error (with markers) for the subblocks of the (1,1) block of D , $c(\mathbf{x})$ is the vertical fault. Circles are for $q = 3$, squares for $q = 5$, stars for $q = 10$.

$c(\mathbf{x})$ is higher, and back towards $\partial\Omega$. We have approximated this τ_f using the fast marching method of Sethian [45]. For this example, it turns out that using either τ or τ_f to obtain oscillations in our basis matrices does not significantly alter the probing accuracy or conditioning, although it does seem that, for higher accuracies at least, using the body waves traveltime makes convergence slightly faster. Figures 9 and 10 demonstrate this for blocks $(1, 1)$ and $(2, 1)$, respectively. We omit plots of the probing and approximation errors and refer the reader to section 4.2 for final probing results and using those in a solver.

4.1.4. The vertical fault. Next, we look at the case of the medium $c(\mathbf{x})$ which has a vertical fault (see Figure 4). We note that this case is harder because some of the blocks will have a 2×2 or 1×2 block structure caused by the discontinuity in the medium. Ideally, each subblock should be probed separately. There are 7 distinct blocks: $n((1, 1)) = 2$, $n((2, 2)) = 1$, $n((4, 4)) = 1$, $n((2, 1)) = 4$, $n((4, 1)) = 4$, $n((3, 1)) = 2$, $n((4, 2)) = 2$. Blocks $(2, 2)$ and $(4, 4)$ are easier to probe than block $(1, 1)$ because they do not exhibit a substructure. Also, since the velocity is smaller on the right side of the fault, the frequency there is higher, which means that blocks involving side 2 are slightly harder to probe than those involving side 4. We first present results for the blocks $(1, 1)$, $(2, 2)$, and $(2, 1)$ of D . In Figure 11 we see the approximation and probing errors for those blocks. Then, in Figure 12, we present results for the errors related to probing the 3 distinct subblocks of the $(1, 1)$ block of D . We can see that probing the $(1, 1)$ block by subblocks helps achieve greater accuracy. We tried splitting other blocks to improve the accuracy of their probing (for example, block $(2, 1)$ has a 1×2 structure because side 1 has a discontinuity in $c(\mathbf{x})$) but the accuracy of the overall DtN map was still limited by the accuracy of probing the $(1, 1)$ block, so we do not show results for other splittings.

4.1.5. The diagonal fault. Now, we look at the case of the medium $c(\mathbf{x})$ which has a diagonal fault, as in Figure 5. As in the case of the vertical fault, some of the blocks will have a 2×2 or 1×2 structure. There are 6 distinct blocks: $n((1, 1)) = 2$, $n((2, 2)) = 2$, $n((2, 1)) = 4$, $n((4, 1)) = 2$, $n((3, 2)) = 2$, $n((3, 1)) = 4$. Again, we split up block $(1, 1)$ in 4 subblocks and probe each of those subblocks separately for greater accuracy, but do not split other blocks. We use two traveltimes for the $(2, 2)$ subblock of block $(1, 1)$. Using as the second arrival time the geometrical traveltime consisting of leaving the boundary and bouncing off the fault, as mentioned in section 3.4, allowed us to increase accuracy by an order of magnitude compared to using only the first arrival traveltime, or compared to using as a second arrival time the usual bounce off the corner (or here, bounce off the fault where it meets $\partial\Omega$). We omit plots of the probing and approximation errors and refer the reader to section 4.2 for final probing results and using those in a solver.

4.1.6. The periodic medium. Finally, we look at the case of the periodic medium of Figure 6. There are 3 distinct blocks: $n((1, 1)) = 4$, $n((2, 1)) = 8$, $n((3, 1)) = 4$. We expect the corresponding DtN map to be harder to probe because its structure will reflect that of the medium; i.e., it will exhibit sharp transitions at points corresponding to sharp transitions in $c(\mathbf{x})$ (similarly as with the two previous discontinuous media). First, note that, in all the previous media we tried, plotting the norm of the antidiagonal entries of diagonal blocks (or subblocks for the faults) shows a rather smooth decay away from the diagonal, as expected from the DtN map. However, that is not the case for the periodic medium: it looks like there is decay away from the diagonal, but variations from that decay can be of relative order 1

(results not shown). This prevents our usual strategy, and so we replace the term decaying term $(h + |x - y|)^{-j_1/\alpha}$ of (3.9) by $P_{j_1}(|x - y|)$.

It is known that solutions to the Helmholtz equation in a periodic medium are Bloch waves with a particular structure [35]. However, using that structure in the basis matrices is not robust. Indeed, using a Bloch wave structure did not succeed very well, probably because our discretization was not accurate enough and so D exhibited that structure only to a very rough degree. Hence we did not use Bloch waves for probing the periodic medium. For these reasons, we tried basis matrices with no oscillations, so without the term $e^{i\omega\tau}$ in (3.9), and obtained the results of section 4.2.

Now that we have probed the DtN map and obtained an approximation \tilde{D} of D , we use this \tilde{D} in a Helmholtz solver as an ABC.

4.2. Using the probed DtN map in a Helmholtz solver. In Figures 13, 14, 15, 16, 17, and 18 we can see the standard solutions to the Helmholtz equation on $[0, 1]^2$ using a large pPML for the various media we consider except the uniform medium, where the solution is well known. We use those as our reference solutions.




FIG. 13. Real part of the solution; c is the slow disk.

FIG. 14. Real part of the solution; c is the waveguide.




FIG. 15. Real part of the solution; c is the vertical fault with source on the left.

FIG. 16. Real part of the solution; c is the vertical fault with source on the right.




FIG. 17. Real part of the solution; c is the diagonal fault.




FIG. 18. Imaginary part of the solution; c is the periodic medium.

TABLE 2
 $c(\mathbf{x}) \equiv 1$.

p for (1, 1)	p for (2, 1)	$q = Q$	$\frac{\ D - \tilde{D}\ _F}{\ D\ _F}$	$\frac{\ u - \tilde{u}\ _F}{\ u\ _F}$
6	1	1	2.0130e-01	3.3191e-01
12	2	1	9.9407e-03	1.9767e-02
20	12	3	6.6869e-04	1.5236e-03
72	20	5	1.0460e-04	5.3040e-04
224	30	10	8.2892e-06	9.6205e-06
360	90	10	7.1586e-07	1.3044e-06

TABLE 3
 $c(\mathbf{x})$ is the waveguide.

p for (1, 1)	p for (2, 1)	q	p for (2, 2)	q	Q	$\frac{\ D - \tilde{D}\ _F}{\ D\ _F}$	$\frac{\ u - \tilde{u}\ _F}{\ u\ _F}$
40	2	1	12	1	2	9.1087e-02	1.2215e-01
40	2	3	20	1	4	1.8685e-02	7.6840e-02
60	20	5	20	3	8	2.0404e-03	1.3322e-02
112	30	10	30	3	13	2.3622e-04	1.3980e-03
264	72	20	168	10	30	1.6156e-05	8.9911e-05
1012	240	20	360	10	30	3.3473e-06	1.7897e-05

We have tested the solver with the probed \tilde{D} as an ABC with success. See Tables 2, 3, 4, 5, 6, and 7 for results corresponding to each medium. We show the number p of basis matrices required for some blocks for that tolerance, the number of solves q of the exterior problem for those blocks, the total number of solves Q , the error in D , and the relative error in Frobenius norm between the solution \tilde{u} using \tilde{D} and the solution u using D . As we can see from the tables, the error $\|u - \tilde{u}\|_F / \|u\|_F$ in the solution u is no more than an order of magnitude greater than the error $\|D - \tilde{D}\|_F / \|D\|_F$ in the DtN map D . Grazing waves, which can arise when the source is close to the boundary of the computational domain, will be discussed in section 4.3.

We note again that, for the uniform medium, using the second arrival traveltime as well as the first for the (1, 1) block allowed us to achieve accuracies of 5 and 6 digits in the DtN map, which was not possible otherwise (see Figure 7). Similarly for the

TABLE 4
 $c(\mathbf{x})$ is the slow disk.

p for (1, 1)	p for (2, 1)	$q = Q$	$\frac{\ D - \tilde{D}\ _F}{\ D\ _F}$	$\frac{\ u - \tilde{u}\ _F}{\ u\ _F}$
40	2	3	1.0730e-01	5.9283e-01
84	2	3	8.0607e-03	4.5735e-02
180	12	3	1.2215e-03	1.3204e-02
264	30	5	1.5073e-04	7.5582e-04
1012	132	20	2.3635e-05	1.5490e-04

TABLE 5
 $c(\mathbf{x})$ is the vertical fault.

Q	$\frac{\ D - \tilde{D}\ _F}{\ D\ _F}$	$\frac{\ u - \tilde{u}\ _F}{\ u\ _F}$, left source	$\frac{\ u - \tilde{u}\ _F}{\ u\ _F}$, right source
5	2.8376e-01	6.6053e-01	5.5522e-01
5	8.2377e-03	3.8294e-02	2.4558e-02
30	1.1793e-03	4.0372e-03	2.9632e-03

TABLE 6
 $c(\mathbf{x})$ is the diagonal fault.

Q	$\frac{\ D - \tilde{D}\ _F}{\ D\ _F}$	$\frac{\ u - \tilde{u}\ _F}{\ u\ _F}$
4	1.6030e-01	4.3117e-01
6	1.7845e-02	7.1500e-02
23	4.2766e-03	1.2429e-02

TABLE 7
 $c(\mathbf{x})$ is the periodic medium.


Q	$\frac{\ D - \tilde{D}\ _F}{\ D\ _F}$	$\frac{\ u - \tilde{u}\ _F}{\ u\ _F}$
50	1.8087e-01	1.7337e-01
50	3.5714e-02	7.1720e-02
50	9.0505e-03	2.0105e-02

waveguide case, block (2, 2). Using a second arrival time for some blocks or subblocks in the cases of the vertical and diagonal faults was also useful. Those results show that probing works best when the medium $c(\mathbf{x})$ is rather smooth. For nonsmooth media such as a fault, it becomes harder to probe the DtN map to a good accuracy, so that the solution to the Helmholtz equation also contains more error.

4.3. Grazing waves. It is well known that ABCs often have difficulties when a source is close to the boundary of the domain, or in general when waves incident to the boundary are almost parallel to it. We wish to verify that the solution \tilde{u} using the result \tilde{D} of probing D does not degrade as the source becomes closer and closer to some side of $\partial\Omega$. For this, we use a right-hand side f to the Helmholtz equation which is a point source, located at the point (x_0, y_0) , where $x_0 = 0.5$ is fixed and $y_0 > 0$ becomes smaller and smaller, until it is a distance $2h$ away from the boundary (the point source's stencil has width h , so a source at a distance h from the boundary does not make sense). We see in Figure 19 that, for $c(\mathbf{x}) \equiv 1$, the solution remains quite good until the source is a distance $2h$ away from the boundary.

In particular, the better the solution is for a source in the middle of the domain, the less it degrades as the source gets closer to the boundary. We obtain similar results for the waveguide, slow disk, and faults (for the vertical fault we locate the source at (x_0, y_0) , where $y_0 = 0.5$ is fixed and x_0 goes to 0 or 1). This shows the probing process itself does not significantly affect how well grazing waves are absorbed.

4.4. Variations of p with N . We now discuss how the number of basis matrices p needed to achieve a desired accuracy depends on N . To do this, we pick 4 consecutive powers of 2 as values for N (128, 256, 512, 1024) and find the appropriate ω such that

FIG. 19. Error in solution u , $c(\mathbf{x}) \equiv 1$, moving point source.

the finite discretization error remains constant at 10^{-1} , so that in fact $N \sim \omega^{1.5}$ as mentioned at the end of section 2.2. We then probe the (1, 1) block of the DtN map, letting ω vary with N and fixing all other parameters, and observe the required p to obtain a fixed probing error (1, 2, or 3 digits). The worst case we have seen in our experiments came from the slow disk, where p seems to follow a very weak power law with N , close to $p \sim 15N^{.12}$ for a probing error of 10^{-1} or $p \sim 15N^{.2}$ for a probing error of 10^{-2} . In all other cases, p is approximately constant with increasing N , or seems to follow a logarithmic law with N .

5. Convergence of basis matrices for probing the half-space DtN map.

This section presents theoretically why p should indeed depend weakly on N or ω (as our numerical results show in section 4.4 for the exterior problem) in the special case of the half-space DtN map in a uniform medium. The goal of this paper is to approximate \mathcal{D} in more general cases, but the result in Theorem 5.1 confirms that a geometrical optics expansion as done in section 3.3 makes sense.

5.1. Statement and proof. Recall from section 3.3 that the form of basis matrices suggested is motivated by a geometrical optics expansion of the Green's function and relating the DtN map to this Green's function. In the half-space case, for the Helmholtz equation $\Delta u(x, y) + k^2 u(x, y) = 0$ in $y > 0$ with the Sommerfeld radiation condition and $c(\mathbf{x}) \equiv 1$ (thus in this section, $\omega = k$), the DtN map is

$$(5.1) \quad \partial_y u(x, y)|_{y=0} = \int_{-\infty}^{\infty} K(|x - x'|) u(x', 0) dx',$$

where

$$(5.2) \quad K(r) = \frac{ik}{2r} H_1^{(1)}(kr),$$

with $H_1^{(1)}$ the Hankel function of the first kind, of order 1. The expression for $K(r)$ is obtained by taking the mixed normal derivative

$$\partial_{\nu_x} \partial_{\nu_y} G(x, y)$$

of the half-space Green's function, which we call G here, as explained, for instance, in [26, p. 92]. (Green's representation formula involves $\partial_{\nu_y} G(x, y)$ to map Dirichlet data to the solution inside the domain; then the DtN map is asking to take an additional

∂_{ν_x} .) For this section, let \mathcal{D} be the operator mapping $u(\cdot, 0)$ to $\partial_y u(x, y)|_{y=0}$ in the half-space, uniform-medium case.

Note that in the case of the uniform medium $c \equiv 1$, there is only one arrival time $\tau(\mathbf{x}, \mathbf{x}')$ between any two points $\mathbf{x} = (x, 0)$ and $\mathbf{x}' = (x', 0)$ on the boundary, and this traveltime is $\tau(\mathbf{x}, \mathbf{x}') = |x - x'|$. In fact, because the medium is uniform, interactions between \mathbf{x} and \mathbf{x}' depend only on the distance $|x - x'|$ between the points, as can be seen from the form (5.2) of the kernel $K(|x - x'|)$. Coming back to the proposed (3.8) and (3.9) of section 3.3, the following choice of basis functions makes sense in the present case:

$$(5.3) \quad \beta_j(x, x') = e^{i\omega|x-x'|} (h + |x - x'|)^{-\frac{j}{\alpha}}.$$

Since $K(r)$ is singular at $r = 0$, we only study the representation of K in the range $r_0 \leq r \leq 1$, with r_0 on the order of $1/k$. Let

$$\tilde{K}(r) = K\chi_{[r_0, 1]}(r).$$

Denote by $\tilde{\mathcal{D}}$ the corresponding operator with integral kernel $\tilde{K}(|x - x'|)$. Simplifying the basis functions of (5.3) by removing the h in the denominator (which had been put in to avoid the singularity in the first place), we prove the following theorem, showing that indeed the set $\{e^{ikr} r^{-j/\alpha}\}$ makes a great basis for expanding $\tilde{K}(r)$.

THEOREM 5.1. *Let $\alpha > \frac{2}{3}$, and let $\tilde{K}_p(r)$ be the best uniform approximation of $\tilde{K}(r)$ in*

$$\text{span} \left\{ \frac{e^{ikr}}{r^{j/\alpha}} : j = 1, \dots, p, \text{ and } r_0 \leq r \leq 1 \right\}.$$

Assume that $r_0 = C/k$ for some $C > 0$ independent of k . Denote by $\tilde{\mathcal{D}}_p$ the operator defined with \tilde{K}_p in place of \tilde{K} . Then, in the operator norm,

$$(5.4) \quad \|\tilde{\mathcal{D}} - \tilde{\mathcal{D}}_p\| \leq C_\alpha p^{1-[3\alpha/2]} \|\tilde{K}\|_\infty$$

for some $C_\alpha > 0$ depending on α .

The proof of Theorem 5.1 is below, and a numerical illustration of the theorem is in section 5.2.

Growing α does not automatically result in a better approximation error, because a careful analysis of the proof would show that C_α grows factorially with α . This behavior translates into a slower onset of convergence in p when α is taken large, as the numerics show in section 5.2, which can in turn be interpreted as the result of “overcrowding” of the basis by very similar-looking matrices.

Notice that \mathcal{D} is not bounded in L^2 , but $\tilde{\mathcal{D}}$ is after the diagonal is cut out. It is easy to see that the operator norm of $\tilde{\mathcal{D}}$ grows like k , for instance, by applying $\tilde{\mathcal{D}}$ to the function e^{-ikx} . The uniform norm of \tilde{K} , however, grows like k^2 (see (5.10)), so the result above shows that an additional factor k is incurred in the error (somewhat akin to numerical pollution) in addition to the factor k gotten from $\|\tilde{\mathcal{D}}\|$.

The important point of the theorem is that the quality of approximation is otherwise independent of k ; i.e., the number p of basis functions does not need to grow like k for the error to be small. In other words, it is unnecessary to “mesh at the wavelength level” to spell out the degrees of freedom that go in the representation of the DtN map’s kernel.

Proof of Theorem 5.1. $K(r)$ has a leading oscillatory factor of e^{ikr} caused by the Hankel function. For this reason it will be easier to begin by analyzing the smoother counterpart of K :

$$(5.5) \quad J(r) = \frac{ik}{2r} H_1^{(1)}(kr) e^{-ikr}.$$

The domain of interest for the r variable is $[r_0, 1]$. Since the proof details will not depend on the constant C in $r_0 = C/k$, we let $r_0 = 1/k$ without loss of generality.

Expanding $K(r)$ in the system of Theorem 5.1 is equivalent to expanding $J(r)$ in polynomials of $r^{-1/\alpha}$ over $[r_0, 1]$. It will be useful to perform the affine rescaling

$$\xi(r) = \frac{2}{r_0^{-1/\alpha} - 1} (r^{-1/\alpha} - 1) - 1 \quad \Leftrightarrow \quad r(\xi) = \left(\frac{\xi + 1}{2} (r_0^{-1/\alpha} - 1) + 1 \right)^{-\alpha}$$

so that the bounds $r \in [r_0, 1]$ turn into $\xi \in [-1, 1]$. Further write $\xi = \cos \theta$ with $\theta \in [0, \pi]$. Our strategy is to expand $J(r(\xi))$ in Chebyshev polynomials $T_n(\xi)$ [9]. By definition, the best p -term approximation of $J(r)$ in polynomials of $r^{-1/\alpha}$ (best in a uniform sense over $[r_0, 1]$) will result in a lower uniform approximation error than that associated with the p -term approximation of $J(r(\xi))$ in the $T_j(\xi)$ system. Let J_p for the p -term approximant of J in our Chebyshev system. Write out the Chebyshev series for $J(r(\xi))$ as

$$(5.6) \quad J(r(\xi)) = \sum_{j=0}^{\infty} c_j T_j(\xi), \quad c_j = \frac{2}{\pi} \int_{-1}^1 \frac{J(r(\xi)) T_j(\xi)}{(1 - \xi^2)^{1/2}} d\xi,$$

with $T_j(\xi) = \cos(j \cos^{-1} \xi)$ and the c_j alternatively written as

$$c_j = \frac{2}{\pi} \int_0^\pi J(r(\cos \theta)) \cos j\theta \, d\theta = \frac{1}{\pi} \int_0^{2\pi} J(r(\cos \theta)) \cos j\theta \, d\theta.$$

The expansion will converge fast because we can integrate by parts in θ and afford to take a few derivatives of J , say m of them, as done in [49]. After noting that the boundary terms cancel out because of periodicity in θ , the coefficients c_j for $j > 0$ are, up to a sign,

$$c_j = \pm \frac{1}{\pi j^m} \int_0^{2\pi} \sin j\theta \frac{d^m}{d\theta^m} J(r(\cos \theta)) \, d\theta, \quad m \text{ odd},$$

$$c_j = \pm \frac{1}{\pi j^m} \int_0^{2\pi} \cos j\theta \frac{d^m}{d\theta^m} J(r(\cos \theta)) \, d\theta, \quad m \text{ even}.$$

It follows that, for $j > 0$ and for all $m > 0$,

$$(5.7) \quad |c_j| \leq \frac{2}{j^m} \max_{\theta} \left| \frac{d^m}{d\theta^m} J(r(\cos \theta)) \right|.$$

Let B_m be a bound on the m th order derivative in (5.7). In order to bound the uniform error $\|J - J_p\|_{L^\infty[r_0, 1]}$ of truncating the Chebyshev series to $J_p = \sum_{j=0}^p c_j T_j$,

combine (5.6) and (5.7), knowing that Chebyshev polynomials are uniformly bounded by 1:

$$\|J - J_p\|_{L^\infty[r_0,1]} \leq \sum_{j=p+1}^{\infty} |c_j| \leq 2B_m \sum_{j=p+1}^{\infty} \frac{1}{j^m} \leq \frac{2B_m}{(m-1)p^{m-1}}, \quad p > 1.$$

The final step is an integral comparison test [7].

The question is now to find a favorable estimate for B_m , from studying successive θ derivatives of $J(r)$ in (5.5). The following bound for the derivatives of Hankel functions, from [18], will be of use. Given any $C > 0$,

$$(5.8) \quad \left| \frac{d^m}{dr^m} \left(H_1^{(1)}(kr) e^{-ikr} \right) \right| \leq C_m (kr)^{-1/2} r^{-m} \quad \text{for } kr \geq C.$$

The change of variables from r to θ results in

$$\begin{aligned} \frac{dr}{d\theta} &= \frac{d\xi}{d\theta} \frac{dr}{d\xi} = (-\sin\theta) \left(-\alpha \left(\frac{\xi+1}{2} (r_0^{-1/\alpha} - 1) + 1 \right)^{-\alpha-1} \frac{(r_0^{-1/\alpha} - 1)}{2} \right) \\ &= (-\sin\theta) \left(-\alpha r^{1+1/\alpha} \frac{k^{1/\alpha} (1 - r_0^{1/\alpha})}{2} \right) = r(kr)^{1/\alpha} \frac{\alpha \sin\theta (1 - r_0^{1/\alpha})}{2}. \end{aligned}$$

Derivatives of higher powers of r are handled by the chain rule, resulting in

$$\frac{d}{d\theta}(r^n) = nr^{n-1} \frac{dr}{d\theta} = nr^n (kr)^{1/\alpha} \frac{\alpha \sin\theta (1 - r_0^{1/\alpha})}{2}.$$

Hence the action of a θ derivative is essentially equivalent to multiplication by $(kr)^{1/\alpha}$. As for higher derivatives of powers of r , it is easy to see by induction that the product rule has them either hit a power of r or a trigonometric polynomial of θ , resulting in a growth of at most $(kr)^{1/\alpha}$ for each derivative:

$$\left| \frac{d^m}{d\theta^m} r^n \right| \leq C_{m,n,\alpha} r^n (kr)^{m/\alpha}.$$

These estimates can now be combined to bound

$$(5.9) \quad \frac{d^m}{d\theta^m} \left(H_1^{(1)}(kr) e^{-ikr} \right).$$

One of two scenarios can occur when applying the product rule in (5.9):

(i) $\frac{d}{d\theta}$ hits

$$\frac{d^{m_1}}{d\theta^{m_1}} \left(H_1^{(1)}(kr) e^{-ikr} \right)$$

for some $m_1 < m$. In this case, one negative power of r results from $\frac{d}{dr}$, and a factor $r(kr)^{1/\alpha}$ results from $\frac{dr}{d\theta}$.

(ii) $\frac{d}{d\theta}$ hits some power of r , or some $\frac{d^{m_2} r}{d\theta^{m_2}}$ for some $m_2 < m$, resulting in a growth of an additional factor $(kr)^{1/\alpha}$.

Thus, there is a $(kr)^{1/\alpha}$ growth factor per derivative in every case. The situation is analogous when dealing with the slightly more complex expression

$$\frac{d^m}{d\theta^m} \left(\frac{1}{r} H_1^{(1)}(kr) e^{-ikr} \right).$$

The number of terms is itself at most factorial in m ; hence

$$\left| \frac{d^m}{d\theta^m} \frac{k}{r} \left(H_1^{(1)}(kr) e^{-ikr} \right) \right| \leq C_{m,\alpha} \frac{k}{r} (kr)^{\frac{m}{\alpha} - \frac{1}{2}} \leq C_{m,\alpha} k^2 (kr)^{\frac{m}{\alpha} - \frac{3}{2}}.$$

Let $m = \lfloor 3\alpha/2 \rfloor$, so that the max over θ is realized when $r = 1/k$, and B_m is on the order of k^2 . It follows that

$$\|J - J_p\|_{L^\infty[r_0,1]} \leq C_\alpha \frac{k^2}{p^{\lfloor 3\alpha/2 \rfloor - 1}}, \quad p > 1, \alpha > 2/3.$$

The kernel of interest, $K(r) = J(r)e^{ikr}$, obeys the same estimate if K_p is the p -term approximation of K in the Chebyshev system modulated by e^{ikr} .

Now, consider the operator norm of $\tilde{\mathcal{D}} - \tilde{\mathcal{D}}_p$ with kernel $\tilde{K} - \tilde{K}_p$, where $\tilde{K}(r) = K(r)\chi_{[r_0,1]}(r)$. Namely,

$$(\tilde{\mathcal{D}} - \tilde{\mathcal{D}}_p)g(x) = \int_0^1 (\tilde{K} - \tilde{K}_p)(|x - y|)g(y) dy.$$

Use the Cauchy–Schwarz inequality to bound

$$\begin{aligned} \|(\tilde{\mathcal{D}} - \tilde{\mathcal{D}}_p)g\|_2 &= \left(\int_{0 \leq x \leq 1} \left| \int_{0 \leq y \leq 1, |x-y| \geq r_0} (K - K_p)(|x - y|)g(y) dy \right|^2 dx \right)^{1/2} \\ &\leq \left(\int_{0 \leq x \leq 1} \int_{0 \leq y \leq 1, |x-y| \geq r_0} |(K - K_p)(|x - y|)|^2 dy dx \right)^{1/2} \|g\|_2 \\ &\leq \left(\int_{0 \leq x \leq 1} \int_{0 \leq y \leq 1, |x-y| \geq r_0} 1 dy dx \right)^{1/2} \|g\|_2 \max_{0 \leq x, y \leq 1, |x-y| \geq r_0} |(K - K_p)(|x - y|)| \\ &\leq \|g\|_2 \|K - K_p\|_{L^\infty[r_0,1]}. \end{aligned}$$

Assembling the bounds gives $\|\tilde{\mathcal{D}} - \tilde{\mathcal{D}}_p\| \leq C_\alpha p^{1 - \lfloor 3\alpha/2 \rfloor} k^2$. It suffices therefore to show that $\|\tilde{K}\|_\infty = \|K\|_{L^\infty[r_0,1]}$ is on the order of k^2 to complete the proof. Letting $\zeta = kr$, the following is obtained:

$$(5.10) \quad \max_{r_0 \leq r \leq 1} |K(r)| = \frac{k^2}{2} \max_{1 \leq \zeta \leq k} \left| \frac{H_1^{(1)}(\zeta)}{\zeta} \right| \geq Ck^2.$$

The last inequality follows from the fact that there exist positive constants d_1, d_2 such that $d_1\zeta^{-3/2} \leq |H_1^{(1)}(\zeta)/\zeta| \leq d_2\zeta^{-3/2}$, shown in [18]. \square

5.2. Numerical confirmation. In order to use Theorem 5.1 to obtain convergent basis matrices, first put back in the oscillations by multiplying $\{r^{-j/\alpha}\}$ with e^{ikr}

to get the following forms¹² (noting again that $c \equiv 1$ means $k = \omega$):

$$\beta_j(x, y) = \frac{e^{i\omega|x-y|}}{|x-y|^{j/\alpha}} \text{ for } x \neq y, 0 \leq x, y \leq 1, \quad \beta_j(x, x) = 0, \quad j = 1, 2, \dots, p.$$

Add to this set the identity matrix in order to capture the diagonal of D , and orthogonalize the resulting collection to get the B_j . Alternatively, we have noticed empirically that using

$$(5.11) \quad \beta_j(x, y) = \frac{e^{i\omega|x-y|}}{(h + |x-y|)^{j/\alpha}}$$

works just as well and is simpler because there is no need to treat the diagonal separately. This different, simpler treatment of the diagonal is used for the exterior problem as well in (3.8) and (3.9) of section 3.3.




FIG. 20. Probing error of the half-space DtN map ($q = 1$, 10 trials, circle markers, and error bars) compared to the approximation error (line), $c(\mathbf{x}) \equiv 1$, $L = 1/4$, $\alpha = 2$, $N = 1024$, $\omega = 51.2$.




FIG. 21. Condition numbers for probing the half-space DtN map, $c(\mathbf{x}) \equiv 1$, $L = 1/4$, $\alpha = 2$, $N = 1024$, $\omega = 51.2$, $q = 1$, 10 trials.

The basis matrices in (5.11) have been used for a numerical confirmation of Theorem 5.1. To obtain the half-space DtN map, instead of solving the exterior problem with a pPML on all sides, solve a problem on a thin strip, with a random Dirichlet boundary condition (for probing) on one of the long edges and a pPML on the other three sides. Figure 20 shows the approximation error, behaving somewhat as in Theorem 5.1. Also plotted are error bars for the probing error, corresponding to ten trials of probing, with $q = 1$. The probing results are about as good as the approximation error, because the relevant condition numbers are all well behaved, as can be seen in Figure 21 for $\alpha = 2$. Back to the approximation error, notice in Figure 20 that increasing α delays the onset of convergence as expected, because of the factor C_α (which is factorial in α) in the statement of Theorem 5.1. Also, for small α , very high inverse powers of r are taken, an ill-conditioned operation. Hence the appearance of a convergence plateau for smaller α is explained by ill-conditioning of the basis matrices, and the absence of data points is because of computational overflow. Finally, increasing α from $1/8$ to 2 gives a higher rate of convergence, as it should because of

¹²In contrast to the forms (3.8) and (3.9) of section 3.3, where the more complicated geometry and media call for taking into account variations across diagonals and more appropriate traveltimes.

the factor $p^{-3\alpha/2}$ in the error (5.4), which gives a rate of convergence of $3\alpha/2$. This is roughly what is obtained numerically. As discussed, further increasing α is not necessarily advantageous since the constant C_α in Theorem 5.1 grows fast in α .

6. Discussion. Probing the DtN map D ultimately makes sense in conjunction with a fast algorithm for its application. In full matrix form, D costs $(4N)^2$ operations to apply. With the help of a compressed representation, this count becomes p times the application complexity of any basis function B_j , which may or may not be advantageous depending on the particular expansion scheme. The better solution for a fast algorithm, however, is to postprocess the compressed expansion from probing into a slightly less compressed but more algorithmically favorable one, such as \mathcal{H} -matrix [30] or butterfly [10].

Note that probing an \mathcal{H} -matrix directly, without further structure, would be highly inefficient. Like in the case of low-rank matrices [33], accurate probing of an \mathcal{H} -matrix requires a number of random applications proportional to the maximum block rank [38], with a big constant. This number could be orders of magnitude larger than the small q needed in our case, when M is a combination of the B_j . Hence the feasibility of probing and the availability of a fast algorithm for matrix-vector multiplication are two different goals that require different expansion schemes.

As for the complexity of solving the Helmholtz equation, compressing the ABC or AL confers the advantage of making the number of nonzeros in the matrix \mathcal{L} of (2.7) independent of the complexity of the ABC. After elimination of the layer, it is easy to see that \mathcal{L} has about $20N^2$ nonzero entries, instead of the $5N^2$ one would expect from a five-point stencil discretization of the Helmholtz equation, because the matrix D (part of a small block of \mathcal{L}) is in general full. Although obtaining a fast matrix-vector product for our approximation of D could reduce the application cost of \mathcal{L} from $20N^2$ to something closer to $5N^2$, it should be noted that the asymptotic complexity does not change—only the constant does, by a factor 4 at best. Hence the discussion about fast algorithms for D is not as crucial as the idea of compressing the ABC in the first place, when the goal is to apply \mathcal{L} fast in an iterative method.

In future work, it would be interesting to probe the Green's function of the Helmholtz equation directly, since its singularity on the diagonal is less steep than that of \mathcal{D} , as discussed in section 3.3. However, its decay at infinity is slower. The advantage would be that once this Green's function was probed, one could simply apply it to the right-hand side f , without having to solve the Helmholtz equation again.

REFERENCES

- [1] M. ABRAMOWITZ AND I. A. STEGUN, *Handbook of Mathematical Functions*, Dover, New York, 1972.
- [2] D. APPELÖ AND T. COLONIUS, *A high-order super-grid-scale absorbing layer and its application to linear hyperbolic systems*, J. Comput. Phys., 228 (2009), pp. 4200–4217.
- [3] S. ASVADUROV, V. DRUSKIN, M. N. GUDDATI, AND L. KNIZHNERMAN, *On optimal finite-difference approximation of PML*, SIAM J. Numer. Anal., 41 (2003), pp. 287–305.
- [4] I. M. BABUŠKA 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 (1997), pp. 2392–2423. Reprinted in SIAM Rev., 42 (2000), pp. 451–484.
- [5] A. BAYLISS, C. I. GOLDSTEIN, AND E. TURKEL, *On accuracy conditions for the numerical computation of waves*, J. Comput. Phys., 59 (1985), pp. 396–404.
- [6] R. BÉLANGER-RIOUX, *Compressed Absorbing Boundary Conditions for the Helmholtz Equation*, doctoral thesis, 2014.

- [7] C. M. BENDER AND S. A. ORSZAG, *Advanced Mathematical Methods for Scientists and Engineers I: Asymptotic Methods and Perturbation Theory*, Springer, New York, 1978.
- [8] J.-P. BÉRENGER, *A perfectly matched layer for the absorption of electromagnetic waves*, J. Comput. Phys., 114 (1994), pp. 185–200.
- [9] J. P. BOYD, *Chebyshev and Fourier Spectral Methods*, 2nd ed., Dover, New York, 2001.
- [10] E. CANDÈS, L. DEMANET, AND L. YING, *A fast butterfly algorithm for the computation of Fourier integral operators*, Multiscale Model. Simul., 7 (2009), pp. 1727–1750.
- [11] W. C. CHEW, J. M. JIN, E. MICHIELSEN, AND J. M. SONG, EDs., *Fast and Efficient Algorithms in Computational Electromagnetics*, Artech House, 2001.
- [12] J. CHIU AND L. DEMANET, *Matrix probing and its conditioning*, SIAM J. Numer. Anal., 50 (2012), pp. 171–193.
- [13] D. COLTON AND R. KRESS, *Integral Equation Methods in Scattering Theory*, Wiley-Interscience, New York, 1983.
- [14] R. COURANT AND D. HILBERT, *Methods of Mathematical Physics. Vol. II: Partial differential equations*, Interscience Publishers, New York, 1962.
- [15] O. CHRISTENSEN, *Frames, Riesz bases, and discrete Gabor/Wavelet expansions*, Bull. Amer. Math. Soc. (N.S.), 38 (2001), pp. 273–291.
- [16] T. A. DAVIS, *Algorithm 832: UMFPACK V4.3—an unsymmetric-pattern multifrontal method*, ACM Trans. Math. Software, 30 (2011), pp. 196–199.
- [17] L. DEMANET, P. D. LETOURNEAU, N. BOUMAL, H. CALANDRA, J. CHIU, AND S. SNELSON, *Matrix probing: A randomized preconditioner for the wave-equation hessian*, Appl. Comput. Harmon. Anal., 32 (2012), pp. 155–168.
- [18] L. DEMANET AND L. YING, *Scattering in flatland: Efficient representations via wave atoms*, Found. Comput. Math., 10 (2010), pp. 569–613.
- [19] J. DIAZ AND P. JOLY, *A time domain analysis of PML models in acoustics*, Comput. Methods Appl. Mech. Engrg., 195 (2006), pp. 3820–3853.
- [20] I. S. DUFF AND J. K. REID, *The multifrontal solution of indefinite sparse symmetric linear*, ACM Trans. Math. Software, 9 (1983), pp. 302–325.
- [21] B. ENGQUIST AND A. MAJDA, *Absorbing boundary conditions for the numerical simulation of waves*, Math. Comp., 31 (1977), pp. 629–651.
- [22] B. ENGQUIST AND O. RUNBORG, *Computational high frequency wave propagation*, Acta Numer., 12 (2003), pp. 181–266.
- [23] B. ENGQUIST AND L. YING, *Sweeping preconditioner for the Helmholtz equation: Hierarchical matrix representation*, Comm. Pure Appl. Math., 64 (2011), pp. 697–735.
- [24] B. ENGQUIST AND L. YING, *Sweeping preconditioner for the Helmholtz equation: Moving perfectly matched layers*, Multiscale Model. Simul., 9 (2011), pp. 686–710.
- [25] Y. ÉRLANGGA, *Advances in iterative methods and preconditioners for the Helmholtz equation*, Arch. Comput. Methods Eng., 15 (2008), pp. 37–66.
- [26] G. B. FOLLAND, *Introduction to Partial Differential Equations*, 2nd ed., Princeton University Press, Princeton, NJ, 1995.
- [27] A. GEORGE, *Nested dissection of a regular finite element mesh*, SIAM J. Numer. Math., 10 (1973), pp. 345–363.
- [28] A. GILLMAN, A. BARNETT, AND P. G. MARTINSSON, *A spectrally accurate direct solution technique for frequency-domain scattering problems with variable media*, BIT, 55 (2015), pp. 141–170.
- [29] A. GILLMAN AND P. G. MARTINSSON, *A direct solver with $O(N)$ complexity for variable coefficient elliptic PDEs discretized via a high-order composite spectral collocation method*, SIAM J. Sci. Comput., 36 (2014), pp. A2023–A2046.
- [30] W. HACKBUSCH, *A sparse matrix arithmetic based on H-matrices. I. Introduction to H-matrices*, Computing, 62 (1999), pp. 89–108.
- [31] T. HAGSTROM AND T. WARBURTON, *A new auxiliary variable formulation of high-order local radiation boundary conditions: Corner compatibility conditions and extensions to first-order systems*, Wave Motion, 39 (2004), pp. 327–338.
- [32] T. HAGSTROM AND T. WARBURTON, *Complete Radiation Boundary Conditions: Minimizing the Long Time Error Growth of Local Methods*, Technical Report, Southern Methodist University.
- [33] N. HALKO, P. G. MARTINSSON, AND J. A. TROPP, *Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions*, SIAM Rev., 53 (2011), pp. 217–288.
- [34] R. HIGDON, *Numerical absorbing boundary conditions for the wave equation*, Math. Comp., 49 (1987), pp. 65–90.
- [35] J. D. JOANNOPOULOS, S. G. JOHNSON, J. N. WINN, AND R. D. MEADE, *Photonic Crystals:*

- Molding the Flow of Light*, 2nd ed., Princeton University Press, Princeton, NJ, 2008.
- [36] J. B. KELLER AND D. GIVOLI, *Exact non-reflecting boundary conditions*, J. Comput. Phys., 82 (1989), pp. 172–192.
 - [37] L. F. KNOCKAERT AND D. DE ZUTTER, *On the complex symmetry of the Poincaré-Steklov operator*, Progr. Electromagnetics Res. B, 7 (2008), pp. 145–157.
 - [38] L. LIN, J. LU, AND L. YING, *Fast construction of hierarchical matrix representation from matrix-vector multiplication*, J. Comput. Phys. 230 (2011), pp. 4071–4087.
 - [39] P.-R. LOH, A. F. OSKOOL, M. IBANESCU, M. SKOROBOGATY, AND S. G. JOHNSON, *Fundamental relation between phase and group velocity, and application to the failure of perfectly matched layers in backward-wave structures*, Phys. Rev. E, 79 (2009), 065601.
 - [40] W. MCLEAN, *Strongly Elliptic Systems and Boundary Integral Equations*, Cambridge University Press, Cambridge, UK, 2000.
 - [41] C. S. MORAWETZ AND D. LUDWIG, *An inequality for the reduced wave operator and the justification of geometrical optics*, Comm. Pure Appl. Math., 21 (1968), pp. 187–203.
 - [42] A. F. OSKOOL, L. ZHANG, Y. AVNIEL, AND S. G. JOHNSON, *The failure of perfectly matched layers, and their redemption by adiabatic absorbers*, Optics Express, 16 (2008), pp. 11376–11392.
 - [43] G. E. PFANDER, *Note on sparsity in signal recovery and in matrix identification*, Open Appl. Math. J., 1 (2007), pp. 21–22.
 - [44] G. E. PFANDER, H. RAUHUT, AND J. TANNER, *Identification of matrices having a sparse representation*, IEEE Trans. Signal Process., 56 (2008), pp. 5376–5388.
 - [45] J. A. SETHIAN, *A fast marching level set method for monotonically advancing fronts*, Proc. Natl. Acad. Sci., 93 (1996), pp. 1591–1595.
 - [46] E. SOMERSALO, M. CHENEY, D. ISAACSON, AND E. ISAACSON, *Layer stripping: A direct numerical method for impedance imaging*, Inverse Problems, 7 (1991), pp. 899–926.
 - [47] E. A. SPENCE, S. N. CHANDLER-WILDE, I. G. GRAHAM, AND V. P. SMYSHLYAEV, *A new frequency-uniform coercive Boundary Integral Equation for acoustic scattering*, Comm. Pure Appl. Math. 64 (2011), pp. 1384–1415.
 - [48] C. C. STOLK, *A rapidly converging domain decomposition method for the Helmholtz equation*, J. Comput. Phys., 241 (2013), pp. 240–252.
 - [49] E. TADMOR, *The exponential accuracy of Fourier and Chebyshev differencing methods*, SIAM J. Numer. Anal., 23 (1986), pp. 1–10.
 - [50] L. N. TREFETHEN AND D. BAU III, *Numerical Linear Algebra*, SIAM, Philadelphia, 1997.
 - [51] G. WHITHAM, *Linear and Nonlinear Waves*, John Wiley and Sons, New York, 1974.