

SPARSIFY AND SWEEP: AN EFFICIENT PRECONDITIONER FOR THE LIPPMANN–SCHWINGER EQUATION*

FEI LIU[†] AND LEXING YING[‡]

Abstract. This paper presents an efficient preconditioner for the Lippmann–Schwinger equation that combines the ideas of the sparsifying and the sweeping preconditioners. Following first the idea of the sparsifying preconditioner, this new preconditioner starts by transforming the dense linear system of the Lippmann–Schwinger equation into a nearly sparse system. The key novelty is a newly designed perfectly matched layer (PML) stencil for the boundary degrees of freedoms. The resulting sparse system gives rise to fairly accurate solutions and hence can be viewed as an accurate discretization of the Helmholtz equation. This new PML stencil also paves the way for applying the moving PML sweeping preconditioner to invert the resulting sparse system approximately. When combined with the standard GMRES solver, this new preconditioner for the Lippmann–Schwinger equation takes only a few iterations to converge for both two-dimensional and three-dimensional (3D) problems, where the iteration numbers are almost independent of the frequency. To the best of our knowledge, this is the first method that achieves near-linear cost to solve the 3D Lippmann–Schwinger equation in high-frequency cases.

Key words. Lippmann–Schwinger equation, acoustic and electromagnetic scattering, quantum scattering, preconditioner

AMS subject classifications. 65F08, 65F50, 65N22, 65R20, 78A45

DOI. 10.1137/17M1132057

1. Introduction. This paper concerns the time-harmonic scattering problem

$$(1) \quad \begin{cases} \left(-\Delta - \frac{\omega^2}{c(x)^2}\right)(u(x) + u_I(x)) = 0, & x \in \mathbb{R}^d, \\ \lim_{r \rightarrow \infty} r^{(d-1)/2} \left(\frac{\partial}{\partial r} - i\omega\right) u(x) = 0, \end{cases}$$

where $u_I(x)$ is the given incoming wave, $u(x)$ is the scattered field to solve, ω is the angular frequency, and $c(x) = \Theta(1)$ is the velocity field such that $c(x) = 1$ outside some bounded region Ω . See Figure 1 for an example. The incoming wave $u_I(x)$ satisfies the homogeneous Helmholtz equation

$$(2) \quad (-\Delta - \omega^2)u_I(x) = 0, \quad x \in \Omega.$$

Let $m(x) = 1 - 1/c(x)^2$ be the perturbation field. Rewriting (1) in terms of $m(x)$ we have

$$(3) \quad (-\Delta - \omega^2 + \omega^2 m(x))u(x) = -\omega^2 m(x)u_I(x), \quad x \in \mathbb{R}^d.$$

Let $G(x)$ be Green’s function of the free space Helmholtz equation

*Submitted to the journal’s Computational Methods in Science and Engineering section May 26, 2017; accepted for publication (in revised form) November 30, 2017; published electronically March 20, 2018.

<http://www.siam.org/journals/sisc/40-2/M113205.html>

Funding: This work was partially supported by the National Science Foundation under award DMS-1521830 and the U.S. Department of Energy’s Advanced Scientific Computing Research program under award DE-FC02-13ER26134/DE-SC0009409.

[†]Institute for Computational and Mathematical Engineering, Stanford University, Stanford, CA 94305 (liufei@stanford.edu).

[‡]Department of Mathematics and Institute for Computational and Mathematical Engineering, Stanford University, Stanford, CA 94305 (lexing@stanford.edu).




FIG. 1. An example of the incoming wave $u_I(x)$ and the scattered field $u(x)$.

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

Convolving both sides of (3) with $G(x)$ gives

$$(4) \quad u(x) + \omega^2 \int_{\Omega} G(x-y)m(y)u(y) \, dy = -\omega^2 \int_{\Omega} G(x-y)m(y)u_I(y) \, dy,$$

which is known as the Lippmann–Schwinger equation written in terms of the scattered field $u(x)$.

Solving the integral equation (4) has several advantages compared to solving (1). First, since $m(x)$ is compactly supported, we only need to solve (4) in Ω . The scattered field $u(x)$ for $x \in \Omega^c$ is explicitly given by (4) once $u(x)$ in Ω is known. More importantly, the resulting wave field $u(x)$ in \mathbb{R}^d automatically satisfies the Sommerfeld radiation condition. On the contrary, for (3) one has to truncate the domain \mathbb{R}^d to some bounded region and impose appropriate boundary conditions to simulate the radiation condition. Second, most local discretizations of (3) suffer from the pollution effect [4] due to inaccurate dispersion relations. Equation (4) avoids this problem by leveraging Green’s function explicitly in the equation.

However, discretizing (4) also raises several issues. First, the resulting linear system is dense. By the Nyquist theorem, a constant number of points per wavelength is needed to capture the oscillations; thus, the number of degrees of freedom N is at least $\Theta(\omega^d)$. In high-frequency cases, N can be rather large where it is impractical to solve general dense linear systems with the direct method. Second, the discretized system can have a very large condition number for nonnegligible perturbations $m(x)$ due to multiple scattering when ω is large. As a result, most standard iterative solvers require a large number of iterations to converge.

Recently, progress has been made to solve the Lippmann–Schwinger equation [2, 7, 8, 1, 28, 18, 24, 27, 31, 33]. [28] proposes a numerical scheme that has spectral accuracy for smooth media by truncating the interactions on the physical domain. [1] presents an adaptive method for the Lippmann–Schwinger equation in two dimensions (2D). [8] solves the 2D Lippmann–Schwinger equation with a technique which is now often referred to as recursive interpolative factorization or recursive skeletonization, where the setup cost is $O(N^{3/2})$ and the solve cost is $O(N \log N)$. [31] approximates

the discretized dense system by a sparse system and applies the nested dissection factorization [15] to the sparse system as a preconditioner to the original dense system. The costs are dominated by merely the nested dissection solver, which are $O(N^{3/2})$ and $O(N \log N)$ for setup and solve in 2D, $O(N^2)$ and $O(N^{4/3})$ for setup and solve in three dimensions (3D), respectively. [33] combines the sparsifying preconditioner [31] with the method of polarized traces [32] to design a preconditioner for the Lippmann–Schwinger equation in 2D, which achieves $O(N)$ setup and $O(N \log N)$ solve costs. As far as we know, [33] is the first to achieve near-linear cost in 2D high-frequency cases.

Meanwhile, a series of domain decomposition methods were developed to solve the Helmholtz equation with Sommerfeld radiation condition [13, 14, 25, 9, 10, 32, 29, 19]. The idea is to divide the domain into slices and impose suitable transmission conditions between these slices. These methods reduce the computational costs to $O(N)$ for setup and $O(N)$ for solve in 2D and $O(N^{4/3})$ for setup and $O(N \log N)$ for solve in 3D, which is a notable improvement over the nested dissection method. A recursive technique [20] further reduces both the setup and the solve costs in 3D to $O(N)$.

This work combines the sparsifying preconditioner in [31] with the sweeping preconditioner in [14] to develop a new preconditioner which solves the Lippmann–Schwinger equation in near-linear cost. The sketch of the method is as follows. We first construct two types of compact stencil schemes to approximate the discretized dense system by a sparse system and then apply the sweeping factorization to the sparse system. The solving process of the sweeping factorization induces an approximating solution, which defines a preconditioner to the original system. The setup and application costs are $O(N)$ and $O(N)$ in 2D and $O(N^{4/3})$ and $O(N \log N)$ in 3D, respectively. Furthermore, the costs in 3D can be reduced to $O(N)$ for setup and $O(N)$ for application by a recursive sweep similar to [20]. When combined with the standard GMRES solver, the preconditioner only needs a few iterations to converge, where the iteration number is almost independent of the angular frequency ω as shown by the numerical results. To the best of our knowledge, this is the first algorithm to solve the Lippmann–Schwinger equation in near-linear cost in 3D high-frequency cases.

Another highlight of this work is the newly designed compact stencil introduced for the preconditioner. The design approach focuses on fitting the stencils to the wave data given by the analytic expressions, such as Green’s function. This approach is quite different from the state-of-the-art methods [26, 17, 3, 5] to design compact stencils, which focus more on the analytic property of the underlying differential operator. Numerical results show that, when used as a method for solving the Helmholtz equation, this scheme is comparably as accurate as the quasi-stabilized finite element method (QSFEM) method in [5] in terms of the phase error.

The rest of the paper is organized as follows. Sections 2 and 3 present the preconditioners and the numerical results in 2D and 3D, respectively, where the detailed approach is explained in section 2 for the 2D case, and section 3 generalizes it to 3D with necessary modifications. Section 4 presents numerical results to show the validity of the compact stencil sparsifying scheme presented in this work when used as a direct method. Conclusions and future work are given in section 5.

2. Preconditioner in 2D. This section describes the preconditioner for the 2D Lippmann–Schwinger equation. Starting by formalizing the dense linear system obtained from discretization, we transform it into an approximately sparse one by

introducing two types of compact stencils. After that, the sweeping factorization is used to solve the truncated sparse system approximately. The whole process can then be treated as a preconditioner for the original dense system of the Lippmann–Schwinger equation.

2.1. Problem formulation. Without loss of generality, we assume that $\Omega = (0, 1)^2$ and that $m(x)$ is supported in Ω . The task is to discretize the Lippmann–Schwinger equation (4) and solve for $u(x)$ in Ω .

The domain Ω is discretized by a uniform Cartesian grid, which allows for the rapid evaluation of the convolution in (4) by fast Fourier transform. Let n be the number of grid points per unit length, $h := 1/(n+1)$ be the step size, and $N := n^2$ be the number of degrees of freedom.

Denote i as the 2D index point and p_i as the grid point with step size h by

$$i := (i_1, i_2), \quad p_i := ih = (i_1h, i_2h), \quad i_1, i_2 \in \mathbb{Z}.$$

Let \mathcal{I} be the index set of the grid points in Ω and \mathcal{D} be the set of the corresponding grid points given by

$$\mathcal{I} := \{i = (i_1, i_2) : 1 \leq i_1, i_2 \leq n\}, \quad \mathcal{D} := \{p_i : i \in \mathcal{I}\}.$$

We also introduce $\bar{\mathcal{I}}$ as the index set for $\bar{\Omega}$ and $\partial\mathcal{I}$ as the boundary index set by

$$\bar{\mathcal{I}} := \{i = (i_1, i_2) : 0 \leq i_1, i_2 \leq n+1\}, \quad \partial\mathcal{I} := \bar{\mathcal{I}} \setminus \mathcal{I},$$

and correspondingly we have $\bar{\mathcal{D}}$ and $\partial\mathcal{D}$ as

$$\bar{\mathcal{D}} := \{p_i : i \in \bar{\mathcal{I}}\}, \quad \partial\mathcal{D} := \{p_i : i \in \partial\mathcal{I}\}.$$

Let u_i be the numerical solution of (4) at p_i for $i \in \mathcal{I}$. To compute the integral in (4), we use the Nyström method

$$\int_{\Omega} G(p_i - y)m(y)u(y) \, dy \approx \sum_{j \in \mathcal{I}} k_{i-j}m_ju_j,$$

where

$$m_i := m(p_i), \quad i \in \mathcal{I}, \quad k_i := G(p_i)h^2, \quad i \neq (0, 0),$$

and $k_{(0,0)}$ is the weight given by a quadrature correction at the singular point of $G(x)$ at $x = 0$, which achieves $O(h^4 \log(1/h)^2)$ accuracy when $m(x)$ is smooth [12]. This gives the discretized equation

$$(5) \quad u_i + \omega^2 \sum_{j \in \mathcal{I}} k_{i-j}m_ju_j = g_i, \quad i \in \mathcal{I},$$

where

$$g_i := -\omega^2 \sum_{j \in \mathcal{I}} k_{i-j}m_j[u_I]_j, \quad i \in \mathcal{I},$$

and $[u_I]_j := u_I(p_j)$ is the discrete value of the incoming wave. Higher-order quadrature can be achieved by using more extended local quadrature correction [12].

With a slight abuse of the notations, we extend the discrete vectors m and g to the whole 2D grid by zero padding:

$$m_i := 0, \quad g_i := 0, \quad i \in \mathbb{Z}^2 \setminus \mathcal{I}.$$

Introducing matrix K with $K_{i,j} := k_{i-j}$, (5) can be written into a more compact form as

$$(6) \quad (I + \omega^2 KM)u = g,$$

where $M := \text{diag}(m)$.

A subtle difference between (5) and (6) is that (5) is a set of equations for the unknowns with indices $i \in \mathcal{I}$, while (6) can be regarded as an equation set defined on the infinite index set \mathbb{Z}^2 , where the unknown vector u is also extended to the whole 2D grid with the extension value determined by the equation implicitly. We have two observations for (6):

- The solution of (6) agrees with the one of (5) in \mathcal{I} . To get the numerical solution of (4) in Ω , we can solve (6) and then restrict the solution to \mathcal{I} instead of solving (5).
- The solution of (6) does not match the numerical solution of (4) outside Ω since the zero padding of g differs from the discretized value of the right-hand side of (4) in Ω^c . Nonetheless, this is not an issue, as we only care about the solution of (4) in Ω .

One may ask, why do we extend the discrete domain to the infinite grid and consider a problem with infinite size? Moreover, the zero padding pattern of g seems rather irrational, as it creates discontinuities at $\partial\mathcal{D}$. The answer is that we are not going to actually solve the \mathbb{Z}^2 -size problem. The purpose of extending the unknown to a larger domain is to introduce the wave attenuation by a perfectly matched layer (PML) on the extended grid to simulate the Sommerfeld radiation condition as we shall see in subsection 2.2.2. The zero padding of g is to ensure that there is no source outside Ω such that the PML approximation holds.

The reader may notice that, if we just use the discretized value of the right-hand side of (4) defined on the whole plane, the solution will also satisfy the Sommerfeld condition, so it seems meaningless to introduce the zero padding. It is true that the right-hand side of (4) on the whole plane will induce a solution satisfying the radiation condition. However, in some cases, when solving (5), we are only given g defined in \mathcal{I} without knowing the actual incoming wave u_I , and it is computationally impractical to get the extension of g determined by (4). This is especially true when we develop preconditioners where the input only involves the right-hand side in the domain of interest.

With the extended problem (6), we are now ready to build a sparse system to approximate (5).

2.2. Sparsification. In this section, we adopt the idea of the sparsifying preconditioner [31] to build a sparse system which serves as an approximation to (5). The sparse system to be constructed has the same sparsity pattern as a compact stencil scheme; i.e., each equation only involves the unknowns at one grid point and its neighbor points, unlike (5), where each equation is dense in \mathcal{I} .

To be specific, we define μ_i as the neighborhood for the index i :

$$\mu_i := \{j : \|j - i\|_\infty \leq 1\}.$$

Now the task is to build for each point i a local stencil supported only in μ_i . We shall build two types of stencils in what follows. The first type is for the interior points, while the second type is for the points near the boundary which are inside what we call “the PML region.”

The PML [6, 16, 11] is a technique to attenuate the waves exponentially near the boundary of the domain so that the zero Dirichlet boundary conditions can be imposed directly to simulate the radiation boundary condition without bringing in too much error. We will explain the PML usage during the construction of the second type of the stencils.

Let us start with extending our domain Ω . Denote Ω^h as Ω with an h -size extension and $\Omega^{h+\eta}$ as the PML extension of Ω^h with width η given by

$$\Omega^h := (-h, 1+h)^2, \quad \Omega^{h+\eta} := (-h-\eta, 1+h+\eta)^2.$$

Here $\eta = bh$ is the PML width, where $b = O(1)$ is the number of discrete layers in each side of the PML region. η is typically around one wavelength. The PML region $\Omega^{h+\eta} \setminus \Omega^h$ is where we will attenuate the scattered field $u(x)$ in subsection 2.2.2. Note that there is a small h -distance between the domain of interest Ω and the PML region $\Omega^{h+\eta} \setminus \Omega^h$. This small distance is introduced on purpose, and the reason will be clear later.

The corresponding index sets in these regions are

$$\mathcal{I}^h := \{i : 0 \leq i_1, i_2 \leq n+1\}, \quad \mathcal{I}^{h+\eta} := \{i : -b \leq i_1, i_2 \leq n+1+b\}.$$

Similar to the notations $\bar{\mathcal{I}}, \partial\mathcal{I}, \mathcal{D}, \partial\mathcal{D}$, and $\bar{\mathcal{D}}$, we introduce $\bar{\mathcal{I}}^h, \partial\mathcal{I}^h, \mathcal{D}^h, \mathcal{D}^{h+\eta}, \partial\mathcal{D}^h$, etc. as the corresponding grid point sets, boundary sets, closures, and so on. The meanings are straightforward, and we omit the formal definitions. See Figure 2 for an illustration.

We now describe how to design two types of stencils for the unknowns indexed by $\mathcal{I}^{h+\eta}$: first for the ones in \mathcal{I}^h and then for the ones in $\mathcal{I}^{h+\eta} \setminus \mathcal{I}^h$. At the end, we assemble them together to form our sparse system.

2.2.1. Stencils for the interior points in \mathcal{I}^h . Following the approach in [31], we design the first type of the stencils for the neighborhood μ_i , where $i \in \mathcal{I}^h$ (see Figure 2, the 3×3 green grids). Taking out the equations in (6) indexed by μ_i we have

$$(7) \quad u_i + \omega^2 \sum_{j \in \mathcal{I}} K_{i,j} m_j u_j = g_i, \quad i \in \mu_i,$$

which can be written as

$$(8) \quad u_{\mu_i} + \omega^2 (K_{\mu_i, \mu_i} [mu]_{\mu_i} + K_{\mu_i, \mu_i^c} [mu]_{\mu_i^c}) = g_{\mu_i}.$$

Here are some explanations for the notations in (8):

- The subscript μ_i stands for the corresponding vector restricted to the index set μ_i ; for example, $[mu]_{\mu_i}$ is the vector of the elementwise multiplication of m and u restricted to μ_i .
- $\mu_i^c := \mathcal{I} \setminus \mu_i$, which is the complement of μ_i with respect to \mathcal{I} .
- K_{μ_i, μ_i^c} is the submatrix of K with row index set μ_i and column index set μ_i^c .

Let us consider a linear combination of the equations in (8). Suppose α is a column vector supported on μ_i . Multiplying both sides of (8) by α^* gives

$$(9) \quad \alpha^* u_{\mu_i} + \omega^2 (\alpha^* K_{\mu_i, \mu_i} [mu]_{\mu_i} + \alpha^* K_{\mu_i, \mu_i^c} [mu]_{\mu_i^c}) = \alpha^* g_{\mu_i},$$

where α^* is the conjugate transpose of α .




FIG. 2. This figure is an illustration of $\Omega, \Omega^h, \Omega^{h+\eta}, \mathcal{D}, \mathcal{D}^h$, etc., where $n = 18, b = 4$. The inner green box is the boundary of Ω , the outer green box (with blue points on it) is the boundary of Ω^h , and the big blue box (with gray points on it) is the boundary of $\Omega^{h+\eta}$. The set of all the green points is \mathcal{D}^h , and the set of all the blue points is $\mathcal{D}^{h+\eta} \setminus \mathcal{D}^h$. \mathcal{D} is the set of all the green points strictly inside the inner green box, while $\partial \mathcal{D}$ consists of the points exactly located on the inner green box. The gray points form $\partial \mathcal{D}^{h+\eta}$, where we will impose the zero Dirichlet boundary conditions after wave attenuation. Each of the 3×3 green grids corresponds to a neighborhood μ_i for some $i \in \mathcal{I}^h$, where we construct stencils of the first type. Correspondingly, the 3×3 blue grids are for μ_i , with $i \in \mathcal{I}^{h+\eta} \setminus \mathcal{I}^h$, where we build stencils of the second type.

To design a local stencil, we hope that the resulting equation (9) only involves unknowns indexed by μ_i . Observing the left-hand side of (9), we found if $\alpha^* K_{\mu_i, \mu_i^c} \approx 0$, then we can truncate the terms involving $u_{\mu_i^c}$, and the resulting equation will be local. But does there exist an α such that $\alpha^* K_{\mu_i, \mu_i^c} \approx 0$? The answer is yes. The reason is that the elements of K are defined by Green's function $G(x)$, which satisfies

$$(10) \quad (-\Delta - \omega^2)G(x) = 0, \quad x \in \mathbb{R}^2 \setminus \{(0, 0)\}.$$

Each column of the matrix K_{μ_i, μ_i^c} can be treated as Green's function centered at some grid point indexed by $j \in \mu_i^c$ and evaluated at the points indexed by the neighborhood μ_i , which does not involve the singular point of $G(x)$ at $x = 0$. Thus, it is reasonable to expect some local stencil α , which can be thought of as a discretization of the local operator $(-\Delta - \omega^2)$ such that $\alpha^* K_{\mu_i, \mu_i^c} \approx 0$. By the translational invariance of Green's function, to find such α , it suffices to require that $\alpha^* K_{\mu, \mu^c} \approx 0$, where

$$\mu := \mu_0 = \{j : \|j\|_\infty \leq 1\}, \quad \mu^c := \{i : -n \leq i_1, i_2 \leq n\} \setminus \mu,$$

which means that we can translate the index i to the origin and consider an equivalent problem. Here the complement of μ is taken with respect to a larger index set. The reason is that, when we translate different indices i to the origin, the corresponding

complement μ_i^c will also be translated. The larger set is taken as the union of all those translated complements to ensure that the condition is sufficient.

To minimize $\alpha^* K_{\mu, \mu^c}$, we consider the optimization problem

$$(11) \quad \min_{\alpha: \|\alpha\|_2} \|\alpha^* K_{\mu, \mu^c}\|_2.$$

The solution is the left singular vector corresponding to the smallest singular value of K_{μ, μ^c} , which can be solved in $O(N)$.

Once we have α , we compute β by setting

$$(12) \quad \beta^* := \alpha^* K_{\mu, \mu}.$$

Then (9) can be approximated as

$$(13) \quad \alpha^* u_{\mu_i} + \omega^2 \beta^* [mu]_{\mu_i} \approx \alpha^* g_{\mu_i}.$$

This defines the local stencil for each $i \in \mathcal{I}^h$.

Note that, if we do the same thing for $i \notin \mathcal{I}^h$, the right-hand side $\alpha^* g_{\mu_i}$ will be 0 due to the zero padding of g . If we build the stencils for all $i \in \mathbb{Z}^2 \setminus \mathcal{I}^h$ and combine them with the Sommerfeld radiation condition at infinity, it will induce a discrete DtN map at $\partial\mathcal{I}^h$. This linear map, though existing in theory, is dense and expensive to compute. Subsection 2.2.2 circumvents this issue by exploiting PML on the extended domain and introducing the second type of stencils to approximate this dense map efficiently.

Now why do we introduce the h -size padded domain Ω^h and build the first type of stencils for $i \in \mathcal{I}^h$ rather than just for \mathcal{I} ? The reason is that, for $i \in \partial\mathcal{I}$, $\alpha^* g_{\mu_i}$ is not necessarily zero; thus, we cannot assign $\partial\mathcal{I}$ to the second type where the corresponding right-hand side is zero. So we enlarge Ω by h -size and build stencils of the first type for $\mathcal{I}^h = \bar{\mathcal{I}}$. Figure 3 shows this subtlety in one dimension (1D) as an illustration.

2.2.2. Stencils for the PML points in $\mathcal{I}^{h+\eta} \setminus \mathcal{I}^h$. Next, we design the stencils for $i \in \mathcal{I}^{h+\eta} \setminus \mathcal{I}^h$ (see Figure 2, the 3×3 blue grids). Define the auxiliary function




FIG. 3. This figure is an illustration of the supports of the discrete vectors g and f in 1D, where $n = 18, b = 4$, and $f_i := \alpha^* g_{\mu_i}$. We see that by introducing the stencil α , the support of g is enlarged to the support of f by one grid point on each side.

$$\sigma(x) := \begin{cases} -\frac{C}{\omega} \left(\frac{x+h}{\eta} \right)^2, & \text{if } -h-\eta < x \leq -h, \\ 0, & \text{if } -h < x < 1+h, \\ \frac{C}{\omega} \left(\frac{x-1-h}{\eta} \right)^2, & \text{if } 1+h \leq x < 1+h+\eta, \end{cases}$$

where $C \sim \Theta(1)$ is some positive constant. We attenuate the scattered field $u(x)$ in the PML region $\Omega^{h+\eta} \setminus \Omega^h$ by introducing the complex stretching

$$\begin{aligned} x^\sigma &:= (x_1^\sigma, x_2^\sigma) = (x_1 + i\sigma(x_1), x_2 + i\sigma(x_2)), \\ u^\sigma(x) &:= u(x^\sigma) = u(x_1 + i\sigma(x_1), x_2 + i\sigma(x_2)), \\ u_i^\sigma &:= u^\sigma(p_i) = u(p_i^\sigma). \end{aligned}$$

By changing variable from x to x^σ , we know that the function $u^\sigma(x)$ satisfies the modified Helmholtz equation in the PML region:

$$(14) \quad \left(-\sum_{d=1}^2 \left(\frac{\partial_d}{1 + i\sigma'(x_d)} \right)^2 - \omega^2 \right) u^\sigma(x) = 0, \quad x \in \Omega^{h+\eta} \setminus \Omega^h.$$

A simple way to build local stencils for $\mathcal{I}^{h+\eta} \setminus \mathcal{I}^h$ is to discretize (14) explicitly with some local scheme, such as the central difference scheme. Unfortunately, it turns out to be not accurate enough to do so. We adopt a different approach. The idea is similar to what we did in the previous section, where we aim to find some local stencil to annihilate a set of given functions evaluated at the points indexed by μ_i . In the above, we used Green's function $G(x)$ to design the stencil α . Here we use a set of “modified plane waves” to achieve the same goal.

Specifically, we first note that the plane wave function

$$F(x) := \exp(i\omega(r \cdot x)), \quad \|r\|_2 = 1,$$

satisfies the free space Helmholtz equation

$$(15) \quad (-\Delta - \omega^2) F(x) = 0, \quad x \in \mathbb{R}^2.$$

Let $F^\sigma(x) := F(x^\sigma)$ be the complex stretching of $F(x)$. We immediately have that $F^\sigma(x)$ satisfies (14) by definition. If we were to design a local stencil γ for μ_i , where $i \in \mathcal{I}^{h+\eta} \setminus \mathcal{I}^h$, we would hope that $\gamma^* F_{\mu_i}^\sigma \approx 0$, where $F_{\mu_i}^\sigma$ is the function $F^\sigma(x)$ evaluated at the grid points indexed by μ_i . Note that any direction r such that $\|r\|_2 = 1$ induces a “modified plane wave” $F^\sigma(x)$. We hope to solve γ by annihilating as many r as possible. To be precise, let R be a set of directions where the elements are sampled uniformly from the unit circle $\{r : \|r\|_2 = 1\}$ and $F_{\mu_i, R}^\sigma$ be a matrix of size $|\mu_i| \times |R|$, each column of which is a modified plane wave function $F^\sigma(x)$ with a direction $r \in R$, evaluated at the grid points indexed by μ_i . Then we solve γ by

$$(16) \quad \min_{\gamma: \|\gamma\|_2=1} \|\gamma^* F_{\mu_i, R}^\sigma\|_2.$$

Intuitively, it is better to increase the sample size $|R|$ to improve the reliability of the stencil. However, larger sample size also leads to more computational cost. Fortunately, it turns out that not too many samples are needed for a reliable result.




FIG. 4. By the symmetry and the translational invariance of the complex stretching along each dimension, one only needs to compute the PML stencils for the points marked with red color near the top right corner.

It suffices to use only the eight most common directions—north, south, west, east, northwest, northeast, southwest, and southeast—to form R , and γ is given by the vector perpendicular to the eight corresponding vectors on μ_i . Note that the solution is unique up to a coefficient ± 1 since we have eight independent modified plane waves and the size of the neighborhood μ_i is nine.

In the PML region, we need to compute different stencils for different neighborhoods due to the lack of translational invariance as a result of the complex stretching. Nevertheless, by the symmetry of the stretching, we only need to compute the stencils near a corner of $\mathcal{I}^{h+\eta} \setminus \mathcal{I}^h$, which takes only $O(b^2)$ work in total. See Figure 4 for an illustration.

We denote γ_i as the stencil for μ_i ; then the corresponding approximating equation is

$$(17) \quad \gamma_i^* u_{\mu_i}^\sigma \approx 0.$$

This defines the local stencil for each $i \in \mathcal{I}^{h+\eta} \setminus \mathcal{I}^h$.

2.2.3. Assemble together. Assembling (13) and (17) together and noting that $u_i^\sigma = u_i$ for $i \in \bar{\mathcal{I}}^h$, we have

$$(18) \quad \begin{cases} \alpha^* u_{\mu_i}^\sigma + \omega^2 \beta^* [mu^\sigma]_{\mu_i} \approx \alpha^* g_{\mu_i}, & i \in \mathcal{I}^h, \\ \gamma_i^* u_{\mu_i}^\sigma \approx 0, & i \in \mathcal{I}^{h+\eta} \setminus \mathcal{I}^h, \end{cases}$$

where α , β , and γ are given in (11), (12), and (16), respectively. Noticing also that u^σ almost satisfies the zero Dirichlet boundary conditions

$$u_i^\sigma \approx 0, \quad i \in \partial \mathcal{I}^{h+\eta},$$

we can introduce the sparse linear system

$$\begin{cases} \alpha^* \tilde{u}_{\mu_i} + \omega^2 \beta^* [m \tilde{u}]_{\mu_i} = \alpha^* g_{\mu_i}, & i \in \mathcal{I}^h, \\ \gamma_i^* \tilde{u}_{\mu_i} = 0, & i \in \mathcal{I}^{h+\eta} \setminus \mathcal{I}^h, \\ \tilde{u}_i = 0, & i \in \partial \mathcal{I}^{h+\eta}, \end{cases}$$

for the unknown \tilde{u} defined on $\mathcal{D}^{h+\eta}$ that serves as an approximation to u^σ . In what follows, we write this system conveniently as

$$(19) \quad H\tilde{u} = f,$$

where the right-hand side f is given by

$$f_i := \begin{cases} \alpha^* g_{\mu_i}, & \text{if } i \in \mathcal{I}^h, \\ 0, & \text{if } i \in \mathcal{I}^{h+\eta} \setminus \mathcal{I}^h. \end{cases}$$

The system (19) is defined on $\mathcal{D}^{h+\eta}$. To get the unknowns on \mathcal{D} , we simply solve (19) and extract the solutions on \mathcal{D} . The result is an approximation to the true solution of (5), and this process can serve as a preconditioner for the linear system (5). In the next section, we present an approach for approximating the solution of (19) efficiently by leveraging the idea of the sweeping preconditioner.

2.3. Sweeping factorization. In this section, we adopt the sweeping factorization to solve the sparse system (19) approximately. The main idea of the sweeping factorization is to divide the domain into slices and eliminate the unknowns slice by slice. An auxiliary PML region is introduced for each slice to build a subproblem to approximate the inverse of the Schur complement during the Gaussian elimination to save computational cost.

To be specific, we first divide the 2D grid into ℓ slices along the x_1 direction. Each slice contains only a few layers. The leftmost slice contains the left PML region, and the rightmost one contains the right PML region (see Figure 5). For simplicity, we assume that each of the middle slices contains b layers and each of the two boundary slices contains $2b$ layers – b normal layers plus b attenuating layers in the PML region. Let $\mathcal{D}_1, \dots, \mathcal{D}_\ell$ be the discrete points in each slice correspondingly, and define $\tilde{u}_{[i]}$ and $f_{[i]}$ as the restrictions of \tilde{u} and f on \mathcal{D}_i , respectively. The sparse system (19) can be written as the block tridiagonal form

$$\begin{bmatrix} H_{[1,1]} & H_{[1,2]} & & & & \\ & H_{[2,1]} & H_{[2,2]} & \ddots & & \\ & & \ddots & \ddots & \ddots & \\ & & & \ddots & \ddots & H_{[\ell-1,\ell]} \\ & & & & H_{[\ell,\ell-1]} & H_{[\ell,\ell]} \end{bmatrix} \begin{bmatrix} \tilde{u}_{[1]} \\ \tilde{u}_{[2]} \\ \vdots \\ \tilde{u}_{[\ell-1]} \\ \tilde{u}_{[\ell]} \end{bmatrix} = \begin{bmatrix} f_{[1]} \\ f_{[2]} \\ \vdots \\ f_{[\ell-1]} \\ f_{[\ell]} \end{bmatrix},$$

where $H_{[i,j]}$'s are the corresponding sparse blocks. Note that we use the bracket subscripts $[\cdot]$ to emphasize that the corresponding unknowns are grouped together in each slice.

We introduce the Schur complement $S_{[i]}$ and its inverse $T_{[i]}$ slice by slice recursively as

$$\begin{aligned} S_{[1]} &= H_{[1,1]}, \quad T_{[1]} = S_{[1]}^{-1}, \\ S_{[i]} &= H_{[i,i]} - H_{[i,i-1]} T_{[i-1]} H_{[i-1,i]}, \quad T_{[i]} = S_{[i]}^{-1}, \quad \text{for } i = 2, \dots, \ell. \end{aligned}$$




FIG. 5. The grid points are divided into five slices along the x_1 direction. Each of the middle slices contains four layers, and each of the two boundary ones contains four more attenuating layers in the PML region.

Then we can solve \tilde{u} by the Gaussian elimination

$$\begin{aligned}\tilde{u}_{[1]} &= T_{[1]} f_{[1]}, \\ \tilde{u}_{[i]} &= T_{[i]} (f_{[i]} - H_{[i,i-1]} \tilde{u}_{[i-1]}), \quad \text{for } i = 2, \dots, \ell, \\ \tilde{u}_{[i]} &= \tilde{u}_{[i]} - T_{[i]} (H_{[i,i+1]} \tilde{u}_{[i+1]}), \quad \text{for } i = \ell - 1, \dots, 1.\end{aligned}$$

The expensive part of the above process is to compute $T_{[i]}$ and apply it to the vectors on \mathcal{D}_i . If, say, we formed $T_{[i]}$ directly, the computation would take $O(b^3 n^3)$ steps and the application $O(b^2 n^2)$ steps. The sweeping factorization reduces the cost by approximating $T_{[i]}$ with a subproblem. To introduce the approximation, we first make a key observation of the operator $T_{[i]}$: Inverting the top left $i \times i$ block of H , one notices that $T_{[i]}$ appears at the bottom right block of the resulting matrix. In other words,

$$H_{[1:i, 1:i]}^{-1} = \begin{bmatrix} H_{[1,1]} & H_{[1,2]} & & & \\ H_{[2,1]} & H_{[2,2]} & & & \\ & & \ddots & & \\ & & & \ddots & \\ & & & & H_{[i-1,i]} \\ & & & & H_{[i,i-1]} & H_{[i,i]} \end{bmatrix}^{-1} = \begin{bmatrix} * & * & \dots & * & * \\ * & * & \dots & * & * \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ * & * & \dots & * & * \\ * & * & \dots & * & T_{[i]} \end{bmatrix}.$$

This means that $T_{[i]}$ is the restriction of $H_{[1:i, 1:i]}^{-1}$ to \mathcal{D}_i . Think of $T_{[i]}$ as an operator from some input vector v to $T_{[i]}v$ on the grid \mathcal{D}_i . Then given v , we can compute $T_{[i]}v$ by solving the equation

$$(20) \quad \begin{bmatrix} H_{[1,1]} & H_{[1,2]} & & & \\ H_{[2,1]} & H_{[2,2]} & \ddots & & \\ & \ddots & \ddots & \ddots & \\ & & \ddots & \ddots & H_{[i-1,i]} \\ & & & H_{[i,i-1]} & H_{[i,i]} \end{bmatrix} \begin{bmatrix} * \\ * \\ \vdots \\ * \\ w \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ \vdots \\ 0 \\ v \end{bmatrix},$$

where w is exactly equal to $T_{[i]}v$. That is to say, given v , we can find $T_{[i]}v$ by padding v with zeros on $\mathcal{D}_{1:(i-1)}$, solving the unknowns on $\mathcal{D}_{1:i}$ by (20), and then extracting the solution on \mathcal{D}_i .

Note that the right-hand side of (20) is zero on $\mathcal{D}_{1:(i-1)}$; thus, the only role of the first $i-1$ blocks of equations in (20) is to induce the radiation condition at the left boundary of \mathcal{D}_i implicitly. To simulate this radiation condition, one can directly put the PML region to the left side of \mathcal{D}_i instead of putting it far away on \mathcal{D}_1 . That is the key idea of the sweeping factorization: move the PML region adjacent to the domain of interest \mathcal{D}_i and approximate the operator $T_{[i]}$ by solving a much smaller system compared to (20) (see Figure 6).

By introducing the modified plane waves, we can build the local stencils for points in the auxiliary PML region on the left of \mathcal{D}_i similar to what was done in subsection 2.2.2. A subtle difference is that the local spatial frequency is perturbed to $\omega\sqrt{1-m(x)}$ instead of ω at location x , and we need to use this local frequency to build the local stencil for each point.

To save computational cost of the stencil construction, we do not use the exact value of the local frequency. Though building local stencil in the PML region with the exact local frequency takes only constant steps per point in theory, the constant




FIG. 6. This figure is an illustration of the moving PML method. Left: $T_{[3]}$ is the restriction of $H_{[1:3,1:3]}^{-1}$ to the discrete domain \mathcal{D}_3 . Right: We move the PML adjacent to \mathcal{D}_3 to induce the approximation operator $\tilde{T}_{[3]}$ as the restriction to \mathcal{D}_3 of the corresponding subproblem on $\mathcal{D}_{2:3}$.

is not small since it involves finding the kernel of a 8×9 matrix. Instead, we consider the square frequency range:

$$[\omega^2(1 - \max\{m(x)\}), \quad \omega^2(1 - \min\{m(x)\})].$$

We choose some samples uniformly from this range interval and build local stencils only for those samples. Then for each point in the PML region, we assign the stencil to be the one from the samples with the closest local square frequency value, a technique introduced earlier in [21]. In practice, only n samples will be enough for an accurate approximation. So it only takes $O(bn)$ steps to build these stencils, which is negligible compared to the problem size $O(n^2)$. An intuition of why we only need n samples is that $O(\omega^2/n)$ is the size of the variation in one neighborhood μ_i on average, so there is no need to make the sampling scale smaller than that.

With the auxiliary PML region on the left of \mathcal{D}_i , we can solve a much smaller system instead of solving (20). In our setting, the set of the auxiliary PML points for \mathcal{D}_i is just \mathcal{D}_{i-1} since the width of the PML region is the same as \mathcal{D}_{i-1} . The auxiliary system can be written as

$$(21) \quad \begin{bmatrix} \tilde{H}_{[i-1, i-1]} & \tilde{H}_{[i-1, i]} \\ H_{[i, i-1]} & H_{[i, i]} \end{bmatrix} \begin{bmatrix} * \\ w \end{bmatrix} = \begin{bmatrix} 0 \\ v \end{bmatrix},$$

where the bottom block of equations is inherited from (20) and the top block is defined by the local stencils of the second type in the auxiliary PML region, the role of which is to simulate the radiation boundary condition on the left of \mathcal{D}_i .

A minor problem here is that the auxiliary PML region for \mathcal{D}_2 consists only the normal layers in \mathcal{D}_1 rather than all the layers, so (21) needs a slight modification for $i = 2$: we restrict the columns of $H_{[2, 1]}$ to the normal layers in \mathcal{D}_1 so that the two blocks are compatible. This problem is inessential, and the patch here is only to make the discussion strictly correct. In practice, the width of the slices and the PML regions can be rather flexible.

Equation (21) defines an approximating operator $\tilde{T}_{[i]} : v \rightarrow w$ for $i \in 1 \dots, \ell$ by restricting the system (21) on $\mathcal{D}_{(i-1):i}$ to \mathcal{D}_i . Note that for $i = 1$, $\tilde{T}_{[1]}$ is exactly equal to $T_{[1]}$ if we treat \mathcal{D}_0 as \emptyset naturally. Compared to (20), (21) is a much smaller quasi-1D problem, which can be solved efficiently with the LU factorization.

2.4. Putting together. We now have all the tools needed to design a linear-complexity preconditioner for the discretized Lippmann–Schwinger equation (5). The setup and application processes of the preconditioner are given by Algorithms 1 and 2, respectively. The slice width b is typically a small integer less than 10; thus, both the setup and the application costs are linear.

We would like to make some comments below for the actual implementation of the algorithm.

1. The algorithm presented above constructs the sweeping factorization along the x_1 direction from left to right. Indeed, since we have radiation conditions on all sides of the domain, we can construct the factorization from both sides and sweep toward the middle slice. The two sweeping fronts can be processed independently until they meet in the middle, where they exchange some local information in the middle slice and then sweep back to the boundaries independently. This is potentially helpful for the parallelization of the algorithm.

Algorithm 1. Setup of the preconditioner for the system (5). Complexity = $O(b^2N)$.

- 1: Compute the stencils α and β . Complexity = $O(N)$.
 - 2: Compute the PML stencils γ for different local frequency samples and different positions of complex stretching. Complexity = $O(bn)$.
 - 3: Divide the domain into ℓ slices as $\mathcal{D}_1, \dots, \mathcal{D}_\ell$.
 - 4: Define the approximating operator $\tilde{T}_{[i]}$ by the sweeping factorization below from steps 5 to 8. Complexity = $O(b^2N)$:
 - 5: **for** $i = 1, \dots, \ell$ **do**
 - 6: Pad \mathcal{D}_i with auxiliary PML points in \mathcal{D}_{i-1} to form the subproblem (21), where the auxiliary PML stencils are built with samples closest to the local square frequency values.
 - 7: Compute the LU factorization of (21), which defines the solution operator. Restricting the solution operator to \mathcal{D}_i induces $\tilde{T}_{[i]}$.
 - 8: **end for**
-

Algorithm 2. Application of the preconditioner. Complexity = $O(bN)$.

- 1: Form the right-hand side f of (19). Complexity = $O(N)$.
 - 2: Solve the linear system (19) approximately by the process below from steps 3 to 9. Complexity = $O(bN)$.
 - 3: $\tilde{u}_{[1]} = \tilde{T}_{[1]}f_{[1]}$
 - 4: **for** $i = 2, \dots, \ell$ **do**
 - 5: $\tilde{u}_{[i]} = \tilde{T}_{[i]}(f_{[i]} - H_{[i,i-1]}\tilde{u}_{[i-1]})$
 - 6: **end for**
 - 7: **for** $i = \ell - 1, \dots, 1$ **do**
 - 8: $\tilde{u}_{[i]} = \tilde{u}_{[i]} - \tilde{T}_{[i]}(H_{[i,i+1]}\tilde{u}_{[i+1]})$
 - 9: **end for**
 - 10: Now \tilde{u} is an approximation to $H^{-1}f$. Extract the solution of \tilde{u} on \mathcal{D} as the output.
-

2. The widths of the slices and the auxiliary PML regions are completely arbitrary. There are two reasons why we set them to be b uniformly in the above. The first is for the simplicity of discussion. The second is that, given the PML width b , it is optimal to set the width of each slice to be also b to minimize the setup and application costs of the preconditioner. In practice, it may not be possible to uniformly divide the domain where each slice contains b layers exactly. In that case, we change the widths of one or two slices accordingly, which has negligible effect to the cost and efficiency of the preconditioner.
3. The constructions of the stencils α, β , and γ , though depending on n and ω , are essentially independent of the velocity field $c(x)$. First, the computation of α and β only involves the free space Green's function $G(x)$, where the velocity field is completely irrelevant. Next, for the local PML stencils γ , they might depend on $c(x)$ slightly, but only on the range, as we see from the sampling process. In practice, $c(x) = \Theta(1)$, so the range is actually bounded for fixed ω . Thus, we can precompute the stencils without given the velocity field. This means that the stencil construction only needs a fixed cost for given problem size, which can be eliminated from the setup process of the algorithm for the input $c(x)$.

2.5. Numerical results. In this section we present the numerical results in 2D. The algorithm is implemented in MATLAB and the tests are performed on a 2.4-GHz server. We force MATLAB to use only one computational thread to test the sequential time cost. The preconditioner is combined with the standard GMRES solver with relative tolerance 10^{-6} and restart value 20. The domain is discretized with $h = \lambda/8$, where $\lambda = 2\pi/\omega$ is the typical wavelength.

We choose $b = 8$ as the width of the slices and the auxiliary PML regions. This corresponds to about one wavelength width for the PML regions and the slices used in the sweeping preconditioner. The sweeping factorization is built with two fronts sweeping toward the middle slice, and the middle slice is padded with auxiliary PMLs on both sides for the corresponding quasi-1D subproblem.

Four velocity fields are tested in 2D:

- (i) a converging Gaussian centered at $(0.5, 0.5)$;
- (ii) a diverging Gaussian centered at $(0.5, 0.5)$;
- (iii) 32 randomly placed converging Gaussians with narrow width;
- (iv) a random velocity field that is equal to 1 at $\partial\Omega$.


The incoming wave $u_I(x)$ for each test is a plane wave shooting downward at frequency ω . The test results are given in Tables 1 to 4, respectively. The notations in the tables are listed as follows:

- ω is the angular frequency;
- N is the number of unknowns;
- T_{setup} is the setup cost of the preconditioner in seconds;
- T_{apply} is the application cost of the preconditioner in seconds;
- N_{iter} is the iteration number;
- T_{solve} is the solve cost of the preconditioner in seconds.

From the numerical tests we observe that both the setup time and the application time scale linearly in N , which are in accordance with the complexity analyses. More

TABLE 1

Numerical results for velocity field (i) in 2D. Top: The velocity field $c(x)$ (left) and the total wave field $u(x) + u_I(x)$ (right) for $\omega/(2\pi) = 64$. Bottom: Table of the numerical results for different problem sizes.



$\omega/(2\pi)$	N	T_{setup}	T_{apply}	N_{iter}	T_{solve}
16	127^2	$1.03\text{e}-01$	$1.47\text{e}-02$	5	$1.96\text{e}-01$
32	255^2	$3.01\text{e}-01$	$4.85\text{e}-02$	5	$4.27\text{e}-01$
64	511^2	$1.23\text{e}+00$	$1.82\text{e}-01$	5	$1.89\text{e}+00$
128	1023^2	$4.80\text{e}+00$	$7.29\text{e}-01$	6	$8.58\text{e}+00$
256	2047^2	$1.95\text{e}+01$	$2.90\text{e}+00$	7	$4.69\text{e}+01$

TABLE 2

Numerical results for velocity field (ii) in 2D. Top: The velocity field $c(x)$ (left) and the total wave field $u(x) + u_I(x)$ (right) for $\omega/(2\pi) = 64$. Bottom: Table of the numerical results for different problem sizes.


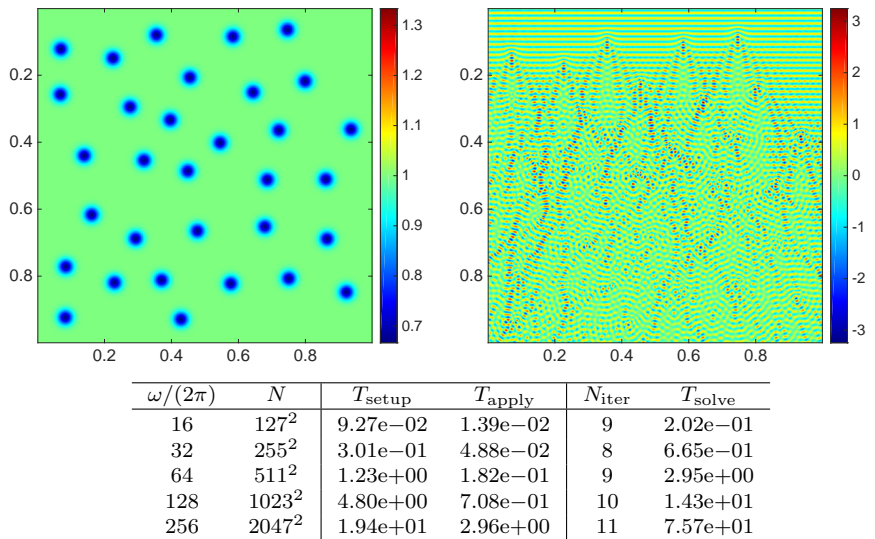


TABLE 3

Numerical results for velocity field (iii) in 2D. Top: The velocity field $c(x)$ (left) and the total wave field $u(x) + u_I(x)$ (right) for $\omega/(2\pi) = 64$. Bottom: Table of the numerical results for different problem sizes.




importantly, the iteration numbers change only slightly as the problem size grows, almost independent of ω .

We notice that the iteration number also depends on the velocity field. For simple fields, such as the diverging Gaussian, it requires less iterations compared to more complicated fields, such as the narrower converging Gaussians. This makes sense

TABLE 4

Numerical results for velocity field (iv) in 2D. Top: The velocity field $c(x)$ (left) and the total wave field $u(x) + u_I(x)$ (right) for $\omega/(2\pi) = 64$. Bottom: Table of the numerical results for different problem sizes.



$\omega/(2\pi)$	N	T_{setup}	T_{apply}	N_{iter}	T_{solve}
16	127^2	$9.05\text{e-}02$	$1.39\text{e-}02$	7	$1.54\text{e-}01$
32	255^2	$2.98\text{e-}01$	$4.88\text{e-}02$	7	$5.80\text{e-}01$
64	511^2	$1.23\text{e+}00$	$1.82\text{e-}01$	8	$2.65\text{e+}00$
128	1023^2	$4.83\text{e+}00$	$7.07\text{e-}01$	9	$1.26\text{e+}01$
256	2047^2	$1.95\text{e+}01$	$2.90\text{e+}00$	9	$6.09\text{e+}01$

intuitively since converging lenses and velocity fields with drastic local variations increase the oscillations and refractions of the wave field; thus, the corresponding systems are harder to solve. In addition, for the sweeping factorization to work well, we need to assume that there are no strong reflections and refractions during the transmission of the waves so that the auxiliary PMLs in the intermediate slices can make correct approximations to the true underlying DtN maps. In practice, moderate amount of wave-ray bendings can be taken care of by a few more iterations as we see in the tests for the multiple diverging Gaussians and the random field. If the velocity field is even worse, for example, if the field has large region of strong discontinuities, then neither will the Nyström method be able to give an accurate discretization scheme nor can the sweeping factorization provide an accurate approximating solution due to the strong reflections caused by the discontinuities. Thus, for our preconditioner to work, we require certain smoothness from the velocity fields. Nonetheless, as we can tell from the numerical examples, the preconditioner works well even when the fields have drastic transitions in narrow regions. So this approach can be widely applied to many use cases.

3. Preconditioner in 3D. This section presents the preconditioner in 3D. As we see from section 2, the approach is essentially dimension independent, and it can be easily generalized to 3D. We will keep the description short, mainly emphasizing the differences compared to the 2D case so that the reader can get the central idea effortlessly. The 3D numerical results for both the recursive approach and the nonrecursive approach of the sweeping factorization are provided in the second part of this section.

3.1. Problem formulation, sparsification, and sweeping factorization.

In this section we formulate the approach in 3D. All the notations in 2D can be easily reused without causing any ambiguities. We will keep them unless otherwise stated.

We assume $\Omega = (0, 1)^3$ contains the support of $m(x)$. The domain is discretized with step size $h = O(1/\omega)$ in each dimension. A similar quadrature correction formula is used for the central weight of Green's function, which gives an accuracy of $O(h^4)$.

For the sparsification process, the first type of stencils α and β can be constructed similarly, where now each neighborhood μ_i has 27 points. For the second type of stencils in the PML region, we use the modified plain waves in 3D, defined similarly as

$$F^\sigma(x) := \exp(i\omega(r \cdot x^\sigma)), \quad \|r\|_2 = 1,$$

where now x^σ and r are in \mathbb{R}^3 and x^σ is stretched to the complex plane from x for all three coordinates. In 2D, the stencil γ is defined as the kernel vector which annihilates the independent waves shooting toward the eight most common directions. This can be done similarly in 3D. We now need a set of 26 directions, which is defined as

$$R := \left\{ \frac{(r_1, r_2, r_3)}{\sqrt{r_1^2 + r_2^2 + r_3^2}} : (r_1, r_2, r_3) \in \{-1, 0, 1\}^3 \setminus \{(0, 0, 0)\} \right\}.$$

In other words, these are the directions shooting from the center of a neighborhood to the 26 boundary neighbor points.

The computational cost of constructing the stencils in 3D seems higher due to more degrees of freedom and larger size of the neighborhoods. But indeed, the relative cost compared to the sweeping factorization is lower than the 2D case, let alone that the stencil computations are independent of the velocity field and can be done by a once-in-a-lifetime preprocessing.

For the sweeping factorization, the domain are now divided into ℓ quasi-2D slices. The auxiliary PMLs are padded to each slice similarly. Each subproblem is quasi-2D, which can be solved efficiently by the nested dissection algorithm with $O(b^3 n^3)$ setup cost and $O(b^2 n^2 \log n)$ application cost. Consisting of $\ell \approx n/b$ subproblems, the whole process has a total setup cost $O(b^2 n^4) = O(b^2 N^{4/3})$ and application cost $O(b n^3 \log n) = O(b N \log N)$. Note that the direct use of the nested dissection algorithm to the 3D sparse system costs $O(N^2)$ for setup and $O(N^{4/3})$ for solve. The sweeping factorization drastically reduces the costs by dimension reduction.

For each of the quasi-2D problems, we can sweep similarly along the x_2 direction, reducing it to ℓ quasi-1D subproblems. This reduces the setup cost to $O(b^4 N)$ and the application cost to $O(b^2 N)$, which are both linear in N but more sensitive to the slice width b . We call this the *recursive* approach [20] and the one in the previous paragraph as *nonrecursive*.

3.2. Numerical results. In this section we present the numerical results in 3D. The test configurations are the same as subsection 2.5 unless otherwise stated. In the 3D tests, we set $b = 4$ for the slice width and PML width.

The four velocity fields tested are the following:


- (i) a converging Gaussian centered at $(0.5, 0.5, 0.5)$;
- (ii) a diverging Gaussian centered at $(0.5, 0.5, 0.5)$;
- (iii) 256 randomly placed converging Gaussians of narrow width;
- (iv) a random velocity field that is equal to 1 at $\partial\Omega$;

The right-hand side is a plain wave shooting downward at frequency ω .

The tests of the nonrecursive approach are given in Tables 5 to 8, and the ones of the recursive approach are in Tables 9 to 12, where the relative costs compared to the nonrecursive approach are also listed as percentages, together with the iteration numbers of the nonrecursive ones in the parentheses for the convenience of comparison.

TABLE 5


Numerical results for velocity field (i) in 3D with the nonrecursive approach. Top: The velocity field $c(x)$ (left) and the total wave field $u(x) + u_I(x)$ (right) in a cross-section view for $\omega/(2\pi) = 32$. Bottom: Table of the numerical results for different problem sizes.



$\omega/(2\pi)$	N	T_{setup}	T_{apply}	N_{iter}	T_{solve}
4	31^3	6.85e+00	4.20e-01	5	2.46e+00
8	63^3	5.74e+01	3.28e+00	5	1.78e+01
16	127^3	5.97e+02	2.74e+01	5	1.57e+02
32	255^3	7.24e+03	2.49e+02	6	1.68e+03

TABLE 6

Numerical results for velocity field (ii) in 3D with the nonrecursive approach. Top: The velocity field $c(x)$ (left) and the total wave field $u(x) + u_I(x)$ (right) in a cross-section view for $\omega/(2\pi) = 32$. Bottom: Table of the numerical results for different problem sizes.



$\omega/(2\pi)$	N	T_{setup}	T_{apply}	N_{iter}	T_{solve}
4	31^3	7.04e+00	4.27e-01	4	1.85e+00
8	63^3	5.85e+01	3.26e+00	5	1.77e+01
16	127^3	5.99e+02	2.70e+01	5	1.54e+02
32	255^3	7.24e+03	2.49e+02	5	1.41e+03

From the numerical tests we see that, as in the 2D cases, the iteration numbers remain essentially independent of the problem size. The preconditioner converges in a few iterations for all the test cases. Another highlight is that the recursive approach requires only zero or one more iteration compared to the nonrecursive approach, which

TABLE 7

Numerical results for velocity field (iii) in 3D with the nonrecursive approach. Top: The velocity field $c(x)$ (left) and the total wave field $u(x) + u_I(x)$ (right) in a cross-section view for $\omega/(2\pi) = 32$. Bottom: Table of the numerical results for different problem sizes.

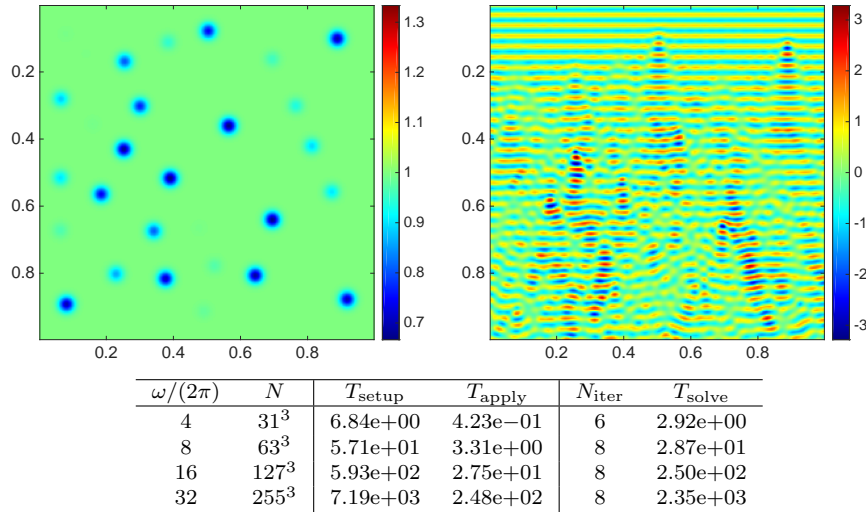
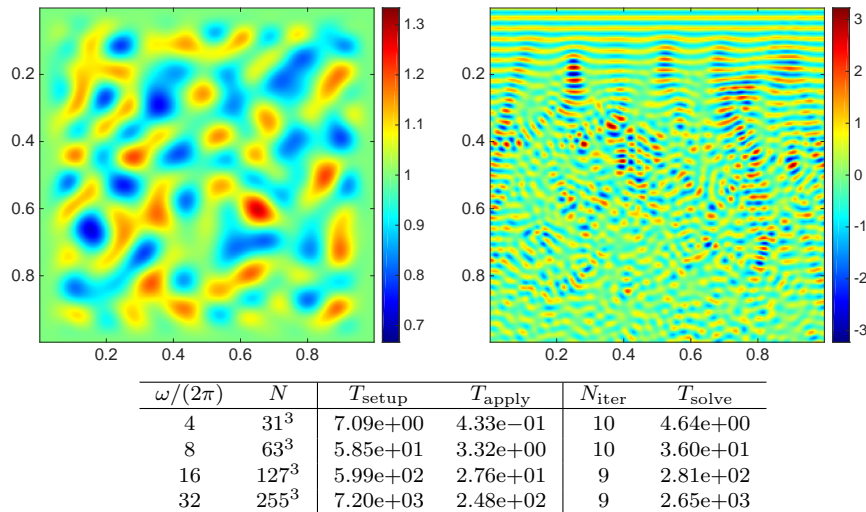


TABLE 8

Numerical results for velocity field (iv) in 3D with the nonrecursive approach. Top: The velocity field $c(x)$ (left) and the total wave field $u(x) + u_I(x)$ (right) in a cross-section view for $\omega/(2\pi) = 32$. Bottom: Table of the numerical results for different problem sizes.



means that the recursive sweeping factorization for the quasi-2D linear systems keeps the total approximation error almost at the same level.

4. Sparsifying scheme as a direct method. In this section, we show that the compact stencils acquired by the sparsifying scheme can be viewed as accurate discretizations of the Helmholtz equation. Specifically, we will solve the 2D homogeneous Helmholtz equation with the compact scheme introduced in the sparsification

TABLE 9
Numerical results for velocity field (i) in 3D with the recursive approach.

$\omega/(2\pi)$	N	T_{setup}	T_{apply}	N_{iter}	T_{solve}
4	31^3	2.79e+00 (41%)	3.44e-01 (82%)	5 (5)	2.01e+00 (82%)
8	63^3	1.62e+01 (28%)	2.25e+00 (69%)	5 (5)	1.29e+01 (73%)
16	127^3	1.10e+02 (18%)	1.64e+01 (60%)	5 (5)	1.01e+02 (65%)
32	255^3	8.23e+02 (11%)	1.24e+02 (50%)	6 (6)	9.33e+02 (56%)

TABLE 10
Numerical results for velocity field (ii) in 3D with the recursive approach.

$\omega/(2\pi)$	N	T_{setup}	T_{apply}	N_{iter}	T_{solve}
4	31^3	2.81e+00 (40%)	3.38e-01 (79%)	5 (4)	1.84e+00 (100%)
8	63^3	1.63e+01 (28%)	2.27e+00 (70%)	5 (5)	1.28e+01 (72%)
16	127^3	1.11e+02 (19%)	1.64e+01 (61%)	5 (5)	1.02e+02 (66%)
32	255^3	8.18e+02 (11%)	1.24e+02 (50%)	6 (5)	9.31e+02 (66%)

TABLE 11
Numerical results for velocity field (iii) in 3D with the recursive approach.

$\omega/(2\pi)$	N	T_{setup}	T_{apply}	N_{iter}	T_{solve}
4	31^3	2.80e+00 (41%)	3.36e-01 (79%)	6 (6)	2.19e+00 (75%)
8	63^3	1.63e+01 (29%)	2.25e+00 (68%)	8 (8)	2.02e+01 (70%)
16	127^3	1.11e+02 (19%)	1.64e+01 (60%)	8 (8)	1.62e+02 (65%)
32	255^3	8.20e+02 (11%)	1.24e+02 (50%)	8 (8)	1.24e+03 (53%)

TABLE 12
Numerical results for velocity field (iv) in 3D with the recursive approach.

$\omega/(2\pi)$	N	T_{setup}	T_{apply}	N_{iter}	T_{solve}
4	31^3	2.81e+00 (40%)	3.37e-01 (78%)	10 (10)	3.67e+00 (79%)
8	63^3	1.64e+01 (28%)	2.24e+00 (68%)	10 (10)	2.52e+01 (70%)
16	127^3	1.12e+02 (19%)	1.64e+01 (60%)	9 (9)	1.82e+02 (65%)
32	255^3	8.22e+02 (11%)	1.24e+02 (50%)	9 (9)	1.39e+03 (52%)

process and compare it with the QSFEM method in [5]. As we shall see from the numerical tests, both methods did comparably well at minimizing the pollution error with only a small number of points per wavelength.

Let us consider

$$(22) \quad (-\Delta - \omega^2)u(x) = f(x), \quad x \in \mathbb{R}^2,$$

where $f(x)$ is a delta source centered at $(0.5, 0.5)$. The exact solution is given by Green's function with a shift of the center (see Figure 7 for an example).

For the sparsifying scheme, we have the discrete equation

$$\alpha^* u_{\mu_i} = \beta^* f_{\mu}$$

for each of the interior point i , where α and β are 9-point stencils given by (11) and (12), respectively, and f is the discrete delta function.




FIG. 7. An example of the 2D Green's function centered at $(0.5, 0.5)$ with $\omega/(2\pi) = 8$.

For the QSFEM method, the 9-point stencil for u is given by

$$A = \begin{bmatrix} A_2 & A_1 & A_2 \\ A_1 & A_0 & A_1 \\ A_2 & A_1 & A_2 \end{bmatrix},$$

where

$$\begin{aligned} A_0 &= 4, \\ A_1 &= 2 \frac{c_1(\kappa)s_1(\kappa) - c_2(\kappa)s_2(\kappa)}{c_2(\kappa)s_2(\kappa)(c_1(\kappa) + s_1(\kappa)) - c_1(\kappa)s_1(\kappa)(c_2(\kappa) + s_2(\kappa))}, \\ A_2 &= 2 \frac{c_2(\kappa) + s_2(\kappa) - c_1(\kappa) - s_1(\kappa)}{c_2(\kappa)s_2(\kappa)(c_1(\kappa) + s_1(\kappa)) - c_1(\kappa)s_1(\kappa)(c_2(\kappa) + s_2(\kappa))}, \\ c_1(\kappa) &:= \cos\left(\kappa \cos \frac{\pi}{16}\right), \quad s_1(\kappa) := \cos\left(\kappa \sin \frac{\pi}{16}\right), \\ c_2(\kappa) &:= \cos\left(\kappa \cos \frac{3\pi}{16}\right), \quad s_2(\kappa) := \cos\left(\kappa \sin \frac{3\pi}{16}\right), \\ \kappa &:= \omega h, \end{aligned}$$

and h is the step size. The right-hand side is the discrete delta function with a scaling.

We solve the 2D homogeneous Helmholtz equation (22) and compare the phase errors against the true solution. Specifically, we write the solutions $u(x)$ as $u(x) = A(x)e^{2\pi i \phi(x)}$ and compare the phase $\phi(x)$ with the one acquired by Green's function. The boundary points are discretized by a slowly turning up PML such that the reflection error is negligible compared to the phase error.

Figure 8 shows the phase errors for a large test case (1024 waves across each dimension) with a small number of points (3 to 5) per wavelength. From the tests we see that the phase error of the sparsifying scheme is comparable to the one of the QSFEM method in [5].

We would like to comment that, as pointed out in [4], 2D compact stencils can be optimized to reduce the pollution error but cannot completely eliminate it. For example, in Figure 8, the phase shifts near the four corners are about $1/6$ for the 3-points-per-wavelength (p.p.w.) test cases, which is not negligible for practical usage. Hence, for large problems, one would eventually have to increase the stencil width or use more points per wavelength.

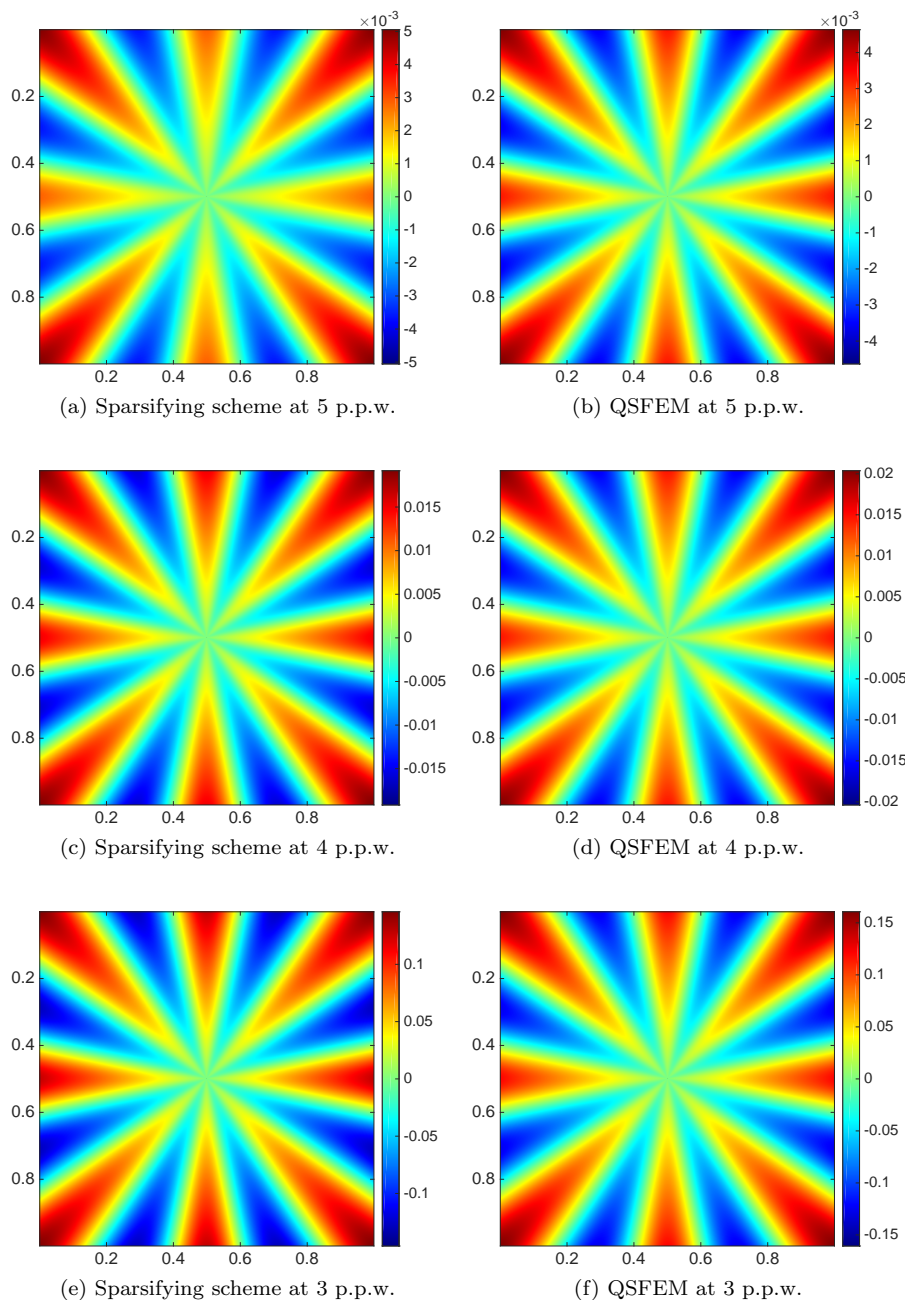


FIG. 8. The phase errors of the two schemes for $\omega/(2\pi) = 1024$. Phase error is defined as $\phi(x) - \phi_G(x)$, where $\phi(x)$ is computed from the numerical solution $u(x) = A(x)e^{2\pi i \phi(x)}$ and $\phi_G(x)$ is acquired by the phase of Green's function (exact solution). From top to bottom are test cases for 5, 4, and 3 p.p.w., respectively. We see that both methods behave similarly in terms of the phase error. For the hardest test cases (3 p.p.w with 1024 waves across each dimension), the phase shifts at the far field (four corners) are about $1/6$. This corresponds to about 2.3×10^{-4} relative phase error. In other words, the distortion of each wavelength of both methods are on the order of 10^{-4} with only three 3 p.p.w.

5. Conclusions and future work. This paper presents the sparsify-and-sweep preconditioner for the Lippmann–Schwinger equation in 2D and 3D. The preconditioner involves two steps. The first step is to sparsify the system by introducing the compact stencil sparsifying scheme. The second step is to apply the sweeping factorization to the sparsified system. Numerical results show that the iteration number is essentially independent of the angular frequency ω .

Though the cost is reduced to linear, potential improvements can be made regarding parallelizations. First, the factorization of the auxiliary subproblems are completely independent and thus can be done in parallel, especially when the recursive approach in 3D is adopted where there are $O(n^2)$ quasi-1D subproblems that can be processed at the same time. Second, the setup and application processes of the nested dissection algorithm can also be parallelized for independent skeleton fronts (see [23], for example). Third, the two sweeping fronts during the application process are also independent and can be processed in a parallel way.

Move future work could include sparsification of dense systems by the data-fitting approach. This approach was first proposed by Ying in [31, 30] for solving highly indefinite systems including time-independent high-frequency wave propagations with radiation conditions or periodic boundary conditions. This paper generalizes it to incorporate the PML approach. There have been some exploration and applications of this sparsification method, such as solving the nonlinear eigenvalue problems in soliton systems [22]. This data-fitting approach to design local schemes is quite different from most classical approaches and could be potentially generalized to other types of integral equations and dense systems.

REFERENCES

- [1] S. AMBIKASARAN, C. BORGES, L.-M. IMBERT-GERARD, AND L. GREENGARD, *Fast, adaptive, high-order accurate discretization of the Lippmann–Schwinger equation in two dimensions*, SIAM J. Sci. Comput., 38 (2016), pp. A1770–A1787, <https://doi.org/10.1137/15M102455X>.
- [2] F. ANDERSSON AND A. HOLST, *A fast, bandlimited solver for scattering problems in inhomogeneous media*, J. Fourier Anal. Appl., 11 (2005), pp. 471–487, <https://doi.org/10.1007/s00041-005-4082-1>.
- [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, <https://doi.org/10.1137/S0036142994269186>.
- [5] I. BABUKA, F. IHLENBURG, E. T. PAIK, AND S. A. SAUTER, *A generalized finite element method for solving the Helmholtz equation in two dimensions with minimal pollution*, Comput. Methods Appl. Mecha. Engrg., 128 (1995), pp. 325–359, [https://doi.org/10.1016/0045-7825\(95\)00890-X](https://doi.org/10.1016/0045-7825(95)00890-X), <http://www.sciencedirect.com/science/article/pii/004578259500890X>.
- [6] J.-P. BERENGER, *A perfectly matched layer for the absorption of electromagnetic waves*, J. Comput. Phys., 114 (1994), pp. 185–200, <https://doi.org/10.1006/jcph.1994.1159>.
- [7] O. P. BRUNO AND E. HYDE, *An efficient, preconditioned, high-order solver for scattering by two-dimensional inhomogeneous media*, J. Comput. Phys., 200 (2004), pp. 670–694, <https://doi.org/10.1016/j.jcp.2004.04.017>.
- [8] Y. CHEN, *A fast, direct algorithm for the Lippmann–Schwinger integral equation in two dimensions*, Adv. Comput. Math., 16 (2002), pp. 175–190, <https://doi.org/10.1023/A:1014450116300>.
- [9] Z. CHEN AND X. XIANG, *A source transfer domain decomposition method for Helmholtz equations in unbounded domain*, SIAM J. Numer. Anal., 51 (2013), pp. 2331–2356, <https://doi.org/10.1137/130917144>.
- [10] Z. CHEN AND X. XIANG, *A source transfer domain decomposition method for Helmholtz equations in unbounded domain Part II: Extensions*, Numer. Math. Theory Methods Appl., 6 (2013), pp. 538–555.

- [11] W. C. CHEW AND W. H. WEEDON, *A 3D perfectly matched medium from modified Maxwell's equations with stretched coordinates*, Microw. Opt. Tech. Lett., 7 (1994), pp. 599–604, <https://doi.org/10.1002/mop.4650071304>.
- [12] R. DUAN AND V. ROKHLIN, *High-order quadratures for the solution of scattering problems in two dimensions*, J. Comput. Phys., 228 (2009), pp. 2152–2174, <https://doi.org/10.1016/j.jcp.2008.11.033>.
- [13] B. ENGQUIST AND L. YING, *Sweeping preconditioner for the Helmholtz equation: Hierarchical matrix representation*, Commun. Pure Appl. Math., 64 (2011), pp. 697–735, <https://doi.org/10.1002/cpa.20358>.
- [14] B. ENGQUIST AND L. YING, *Sweeping preconditioner for the Helmholtz equation: Moving perfectly matched layers*, Multiscale Model. Simul., 9 (2011), pp. 686–710, <https://doi.org/10.1137/100804644>.
- [15] A. GEORGE, *Nested dissection of a regular finite element mesh*, SIAM J. Numer. Anal., 10 (1973), pp. 345–363. Collection of articles dedicated to the memory of George E. Forsythe.
- [16] S. G. JOHNSON, *Notes on Perfectly Matched Layers (PMLs)*, Lecture notes, Massachusetts Institute of Technology, Cambridge, MA, 2008.
- [17] J. KRISTEK, P. MOCZO, AND M. GALIS, *A brief summary of some PML formulations and discretizations for the velocity-stress equation of seismic motion*, Stud. Geophys. Geodaetica, 53 (2009), p. 459.
- [18] F. LANZARA, V. MAZ'YA, AND G. SCHMIDT, *Numerical solution of the Lippmann-Schwinger equation by approximate approximations*, J. Fourier Anal. Appl., 10 (2004), pp. 645–660, <https://doi.org/10.1007/s00041-004-3080-z>.
- [19] F. LIU AND L. YING, *Additive sweeping preconditioner for the Helmholtz equation*, Multiscale Model. Simul., 14 (2016), pp. 799–822.
- [20] F. LIU AND L. YING, *Recursive sweeping preconditioner for the three-dimensional Helmholtz Equation*, SIAM J. Sci. Comput., 38 (2016), pp. A814–A832, <https://doi.org/10.1137/15M1010154>.
- [21] F. LIU AND L. YING, *Localized sparsifying preconditioner for periodic indefinite systems*, Commun. Math. Sci., 15 (2017), pp. 1155–1169, <https://doi.org/10.4310/CMS.2017.v15.n4.a12>.
- [22] J. LU AND L. YING, *Sparsifying preconditioner for soliton calculations*, J. Comput. Phys., 315 (2016), pp. 458–466, <https://doi.org/10.1016/j.jcp.2016.03.061>.
- [23] J. POULSON, B. ENGQUIST, S. LI, AND L. YING, *A parallel sweeping preconditioner for heterogeneous 3D Helmholtz equations*, SIAM J. Sci. Comput., 35 (2013), pp. C194–C212.
- [24] J. SIFUENTES, *Preconditioned Iterative Methods for Inhomogeneous Acoustic Scattering Applications*, Ph.D. thesis, Rice University, Houston, TX, 2010.
- [25] C. C. STOLK, *A rapidly converging domain decomposition method for the Helmholtz equation*, J. Comput. Phys., 241 (2013), pp. 240–252, <https://doi.org/10.1016/j.jcp.2013.01.039>.
- [26] C. C. STOLK, *A dispersion minimizing scheme for the 3-D Helmholtz equation based on ray theory*, J. Comput. Phys., 314 (2016), pp. 618–646, <https://doi.org/10.1016/j.jcp.2016.03.023>, <http://www.sciencedirect.com/science/article/pii/S0021999116001753>.
- [27] G. VAINIKKO, *Fast solvers of the Lippmann-Schwinger equation*, in Direct and Inverse Problems of Mathematical Physics, Springer, Berlin, 2000, pp. 423–440.
- [28] F. VICO, L. GREENGARD, AND M. FERRANDO, *Fast convolution with free-space Green's functions*, J. Comput. Phys., 323 (2016), pp. 191–203, <https://doi.org/10.1016/j.jcp.2016.07.028>, <http://www.sciencedirect.com/science/article/pii/S0021999116303230>.
- [29] A. VION AND C. GEUZAIN, *Double sweep preconditioner for optimized Schwarz methods applied to the Helmholtz problem*, J. Comput. Phys., 266 (2014), pp. 171–190.
- [30] L. YING, *Sparsifying preconditioner for pseudospectral approximations of indefinite systems on periodic structures*, Multiscale Model. Simul., 13 (2015), pp. 459–471, <https://doi.org/10.1137/140985159>.
- [31] L. YING, *Sparsifying preconditioner for the Lippmann-Schwinger equation*, Multiscale Model. Simul., 13 (2015), pp. 644–660, <https://doi.org/10.1137/140985147>.
- [32] L. ZEPEDA-NÚÑEZ AND L. DEMANET, *The method of polarized traces for the 2D Helmholtz equation*, J. Comput. Phys., 308 (2016), pp. 347–388, <https://doi.org/10.1016/j.jcp.2015.11.040>.
- [33] L. ZEPEDA-NÚÑEZ AND H. ZHAO, *Fast alternating bidirectional preconditioner for the 2D high-frequency Lippmann-Schwinger equation*, SIAM J. Sci. Comput., 38 (2016), pp. B866–B888, <https://doi.org/10.1137/16M1064660>.