

RECURSIVE SWEEPING PRECONDITIONER FOR THE THREE-DIMENSIONAL HELMHOLTZ EQUATION*

FEI LIU[†] AND LEXING YING[‡]

Abstract. This paper introduces the recursive sweeping preconditioner for the numerical solution of the Helmholtz equation in three dimensions. This is based on the earlier work of the sweeping preconditioner with the moving perfectly matched layers. The key idea is to apply the sweeping preconditioner recursively to the quasi-two-dimensional auxiliary problems introduced in the three-dimensional (3D) sweeping preconditioner. Compared to the nonrecursive 3D sweeping preconditioner, the setup cost of this new approach drops from $O(N^{4/3})$ to $O(N)$, the application cost per iteration drops from $O(N \log N)$ to $O(N)$, and the iteration count increases only mildly when combined with the standard GMRES solver. Several numerical examples are tested and the results are compared with the nonrecursive sweeping preconditioner to demonstrate the efficiency of the new approach.

Key words. Helmholtz equation, perfectly matched layers, preconditioners, high frequency waves

AMS subject classifications. 65F08, 65N22, 65N80

DOI. 10.1137/15M1010154

1. Introduction. Let the domain of interest be the unit cube $D = (0, 1)^3$ for simplicity. The time-independent wave field $u(x)$ satisfies the Helmholtz equation

$$(1) \quad \Delta u(x) + \frac{\omega^2}{c^2(x)} u(x) = f(x) \quad \forall x \in D,$$

where ω is the angular frequency, $c(x)$ is the velocity field with a bound $c_{\min} \leq c(x) \leq c_{\max}$ where c_{\min} and c_{\max} are assumed to be of $\Theta(1)$, and $f(x)$ is the time-independent external force. The typical boundary conditions for this problem are approximations of the Sommerfeld radiation condition, which means that the wave is absorbed by the boundary and there is no reflection coming from it. Other boundary conditions, such as the Dirichlet boundary condition, can also be specified on part of the boundary depending on the modeling setup.

In this setting, $\omega/(2\pi)$ is the typical wave number of the problem and $\lambda = 2\pi/\omega$ is the typical wavelength. For most applications, the Helmholtz equation is discretized with at least a few points (typically 4 to 20) per wavelength. So the number of points n in each direction is at least proportional to ω . As a result, the total degree of freedom $N = n^3 = \Omega(\omega^3)$ can be very large for high frequency three-dimensional (3D) problems. In addition, the corresponding discrete system is highly indefinite and the standard iterative solvers and/or preconditioners are no longer efficient for such problems. These together make the problem challenging for numerical solution. We refer to the review article [10] by Ernst and Gander for more details on this.

*Submitted to the journal's Methods and Algorithms for Scientific Computing section February 25, 2015; accepted for publication (in revised form) January 8, 2016; published electronically March 15, 2016. This work was partially supported by the National Science Foundation under award DMS-1328230 and the U.S. Department of Energy's Advanced Scientific Computing Research program under award DE-FC02-13ER26134/DE-SC0009409.

<http://www.siam.org/journals/sisc/38-2/M101015.html>

[†]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).

Recently in [8], Engquist and Ying developed a sweeping preconditioner using the moving perfectly matched layers (PMLs) and obtained essentially linear solve times for 3D high frequency Helmholtz equations. A key step of that approach is to approximate the 3D problem with a sequence of $O(n)$ PML-padded auxiliary quasi-2D problems, each of which can be solved efficiently with a sparse direct method such as the nested dissection algorithm. As an extension, this paper applies the sweeping idea recursively to further reduce each auxiliary quasi-2D problem into a sequence of PML-padded quasi-1D problems, each of which can be solved easily with the sparse LDU factorization for banded systems. As a result, the setup cost of the preconditioner improves from $O(N^{4/3})$ to $O(N)$ and the application cost reduces from $O(N \log N)$ to $O(N)$.

There is a vast literature on iterative methods and preconditioners for the Helmholtz equation, among which the multigrid method combined with a shifted operator is one of the most effective approaches. For instance, in [2], Calandra et al. proposed a new two-grid method and achieved an improvement of the iteration number, which grows linearly in ω even in high frequency ranges. Our approach is different from the complex shift based method, and typically the iteration number grows at most logarithmically in ω , while the computational cost per iteration is higher than for those based on the multigrid method. For a complete discussion of the iterative methods for the Helmholtz equation, we refer to the review articles [9] by Erlangga and [10] by Ernst and Gander. The discussion below only touches on the methods that share similarity with the sweeping preconditioners. The analytic ILU factorization [11] is the first to use incomplete LDU factorizations for preconditioning the Helmholtz equation. Compared to the moving PML sweeping preconditioner, the method uses the absorbing boundary condition, which is less effective compared to the PML, and hence the iteration count grows much more rapidly.

Since the sweeping preconditioners [7, 8] were proposed, there have been a number of exciting developments for the numerical solutions of the high frequency Helmholtz equation, including but not limited to [16, 15, 19, 17, 18, 20, 3, 4, 21]. In [16], Stolk proposed a domain decomposition algorithm that utilizes suitable transmission conditions based on the PMLs between the subdomains to achieve a near-linear cost. In [15], Poulson et al. discussed a parallel version of the moving PML sweeping preconditioner to deal with large-scale problems from applications such as seismic inversion. In [19, 17, 18], Tsuji and coauthors extended the moving PML sweeping preconditioner method to other time-harmonic wave equations and more general numerical discretization schemes. In [20], Vion and Geuzaine proposed a double sweep algorithm, studied several implementations of the absorbing boundary conditions, and compared their numerical performance. Finally in [3, 4], Chen and Xiang introduced a sweeping-style domain decomposition method where the emphasis was on the source transferring between the adjacent subdomains. In [21], Zepeda-Núñez and Demanet developed a novel parallel domain decomposition method that uses transmission conditions to define explicitly the up- and down-going waves.

The rest of the paper is organized as follows. We first state the problem and the discretization used in section 2. Section 3 reviews the nonrecursive moving PML sweeping preconditioner proposed in [8]. Section 4 discusses in detail the recursive sweeping preconditioner. Numerical results are presented in section 5. Finally, the conclusion and some future directions are provided in section 6.

2. Problem formulation. Following [8], we assume that the PML [1, 5, 13] is utilized at part of the boundary where the Sommerfeld radiation condition is specified.

The sweeping preconditioner in [8] requires that at least one of the six faces of the domain $D = (0, 1)^3$ is specified with the PML boundary condition. As we shall see soon, the recursive sweeping preconditioner instead requires the PML condition to be specified at least at two nonparallel faces. Without loss of generality, we assume that it is specified at $x_2 = 0$ and $x_3 = 0$. There is no restriction on the type of boundary conditions specified on the other four faces. However, to simplify the discussion, we assume that the Dirichlet condition is used. The PML boundary condition introduces auxiliary functions

$$\sigma(x) = \begin{cases} \frac{C}{\eta} \left(\frac{x - \eta}{\eta} \right)^2, & x \in [0, \eta], \\ 0, & x \in (\eta, 1], \end{cases}$$

and

$$s(x) = \left(1 + i \frac{\sigma(x)}{\omega} \right)^{-1}, \quad s_1(x) \equiv 1, \quad s_2(x) = s(x_2), \quad s_3(x) = s(x_3),$$

where C is an appropriate positive constant independent of ω , and η is the PML width, which is typically around one wavelength. The Helmholtz equation with PML is

$$(2) \quad \begin{cases} \left((s_1 \partial_1)^2 + (s_2 \partial_2)^2 + (s_3 \partial_3)^2 + \frac{\omega^2}{c^2(x)} \right) u(x) = f(x) & \forall x \in D = (0, 1)^3, \\ u(x) = 0 & \forall x \in \partial D. \end{cases}$$

It is typically assumed that the support of $f(x)$ is in $(0, 1) \times (\eta, 1) \times (\eta, 1)$, which means that the force is not located in the PML region. The cube $[0, 1]^3$ is discretized with a Cartesian grid where the grid size is $h = \frac{1}{n+1}$ and n is proportional to ω . The set of all the interior points of the grid is given by

$$P = \{p_{i,j,k} = (ih, jh, kh) : 1 \leq i, j, k \leq n\},$$

and the degree of freedom is $N = n^3$.

Applying the standard seven-point finite difference stencil results in the discretized system

$$(3) \quad \begin{aligned} & \frac{(s_1)_{i,j,k}}{h} \left(\frac{(s_1)_{i+1/2,j,k}}{h} (u_{i+1,j,k} - u_{i,j,k}) - \frac{(s_1)_{i-1/2,j,k}}{h} (u_{i,j,k} - u_{i-1,j,k}) \right) \\ & + \frac{(s_2)_{i,j,k}}{h} \left(\frac{(s_2)_{i,j+1/2,k}}{h} (u_{i,j+1,k} - u_{i,j,k}) - \frac{(s_2)_{i,j-1/2,k}}{h} (u_{i,j,k} - u_{i,j-1,k}) \right) \\ & + \frac{(s_3)_{i,j,k}}{h} \left(\frac{(s_3)_{i,j,k+1/2}}{h} (u_{i,j,k+1} - u_{i,j,k}) - \frac{(s_3)_{i,j,k-1/2}}{h} (u_{i,j,k} - u_{i,j,k-1}) \right) \\ & + \left(\frac{\omega^2}{c^2} \right)_{i,j,k} u_{i,j,k} = f_{i,j,k} \quad \forall 1 \leq i, j, k \leq n, \end{aligned}$$

where the subscript (i, j, k) means that the corresponding function is evaluated at the point $p_{i,j,k} = (ih, jh, kh)$ and the definition of the points here extends to half integers as well. The computational task is to solve (3) efficiently. We note that, unlike the symmetric version adopted in [7, 8], here the nonsymmetric version of the equation is used. Figure 1 provides an illustration of the computational domain and the discretization grid.

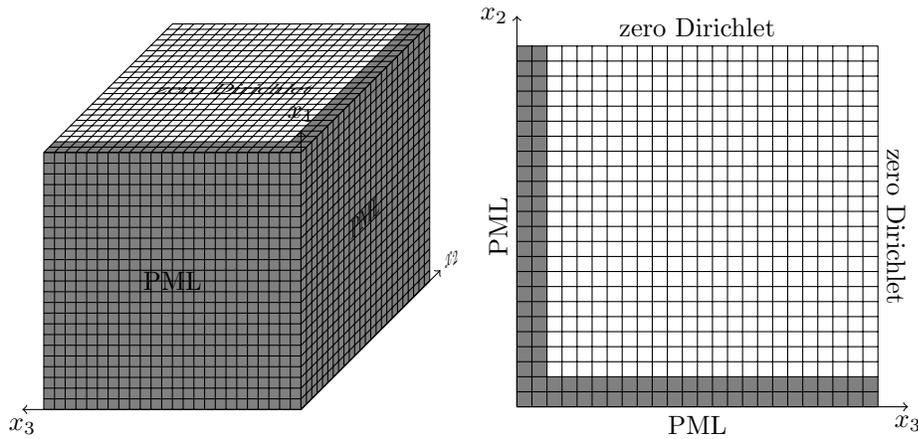


FIG. 1. The domain of interest. Left is a 3D view of the domain. Right is an x_3 - x_2 cross-section view, where each cell stands for a 1D column. The gray area stands for the PML region.

3. Review of the sweeping preconditioner with moving PML. This section gives a brief review of the nonrecursive moving PML sweeping preconditioner proposed in [8] for completeness. More details can be found in the original paper [8]. The starting point of the sweeping preconditioner is a block LDU factorization called the sweeping factorization. To build this factorization, the algorithm sweeps along the x_3 direction starting from the face $x_3 = 0$. The unknowns with subscript index (i, j, k) are ordered with column-major order, i.e., first dimension 1, then dimension 2, and finally dimension 3. We define the vectors

$$u = [u_{1,1,1}, \dots, u_{n,1,1}, \dots, u_{n,n,1}, \dots, u_{n,n,n}]^T,$$

$$f = [f_{1,1,1}, \dots, f_{n,1,1}, \dots, f_{n,n,1}, \dots, f_{n,n,n}]^T.$$

By introducing

$$P_m = \{p_{1,1,m}, \dots, p_{n,1,m}, \dots, p_{n,n,m}\}$$

as the points on the m th plane and also

$$u_{:,:,m} = [u_{1,1,m}, \dots, u_{n,1,m}, \dots, u_{n,n,m}]^T,$$

$$f_{:,:,m} = [f_{1,1,m}, \dots, f_{n,1,m}, \dots, f_{n,n,m}]^T,$$

one can write the system (3) compactly as $Au = f$ with the following block form:

$$(4) \quad \begin{bmatrix} A_{1,1} & A_{1,2} & & & & \\ A_{2,1} & A_{2,2} & \ddots & & & \\ & \ddots & \ddots & & & \\ & & & A_{n-1,n} & & \\ & & & A_{n,n-1} & A_{n,n} & \end{bmatrix} \begin{bmatrix} u_{:,:,1} \\ u_{:,:,2} \\ \vdots \\ u_{:,:,n} \end{bmatrix} = \begin{bmatrix} f_{:,:,1} \\ f_{:,:,2} \\ \vdots \\ f_{:,:,n} \end{bmatrix}.$$

By defining S_k and T_k recursively via

$$S_1 = A_{1,1}, \quad T_1 = S_1^{-1},$$

$$S_m = A_{m,m} - A_{m,m-1}T_{m-1}A_{m-1,m}, \quad T_m = S_m^{-1}, \quad m = 2, \dots, n,$$

the standard block LDU factorization of the block tridiagonal matrix A is

$$A = L_1 \dots L_{n-1} \begin{bmatrix} S_1 & & \\ & \ddots & \\ & & S_n \end{bmatrix} U_{n-1} \dots U_1,$$

where L_m and U_m are the corresponding unit lower and upper triangular matrices with the only nonzero off-diagonal blocks

$$L_m(P_{m+1}, P_m) = A_{m+1,m} T_m, \quad U_m(P_m, P_{m+1}) = T_m A_{m,m+1}, \quad m = 1, \dots, n-1.$$

It is not difficult to see that computing this factorization takes $O(N^{7/3})$ steps. Once it is available, u can be computed in $O(N^{5/3})$ steps by

$$u = \begin{bmatrix} u_{:,1} \\ \vdots \\ u_{:,n} \end{bmatrix} = A^{-1} f = U_1^{-1} \dots U_{n-1}^{-1} \begin{bmatrix} T_1 & & \\ & \ddots & \\ & & T_n \end{bmatrix} L_{n-1}^{-1} \dots L_1^{-1} f.$$

The main disadvantage of the above algorithm is that S_m and T_m are in general dense matrices of size $n^2 \times n^2$, so the corresponding dense linear algebra operations are expensive. The sweeping preconditioner overcomes this difficulty by approximating T_m efficiently for P_m with $mh \in (\eta, 1]$, i.e., for P_m not in the PML region at the face $x_3 = 0$. The key point is to consider the physical meaning of T_m . From now on let us assume $\eta = bh$, which implies that there are b layers in the PML region at $x_3 = 0$. Restricting the factorization to the upper left $m \times m$ block of A where $m = b+1, \dots, n$ gives

$$\begin{bmatrix} A_{1,1} & A_{1,2} & & & \\ A_{2,1} & A_{2,2} & \ddots & & \\ & \ddots & \ddots & & \\ & & & A_{m-1,m} & \\ & & & A_{m,m-1} & A_{m,m} \end{bmatrix} = L_1 \dots L_{m-1} \begin{bmatrix} S_1 & & & & \\ & S_2 & & & \\ & & \ddots & & \\ & & & \ddots & \\ & & & & S_m \end{bmatrix} U_{m-1} \dots U_1,$$

where L_t and U_t are redefined by restricting to their upper left $m \times m$ blocks. Inverting both sides leads to

$$\begin{bmatrix} A_{1,1} & A_{1,2} & & & \\ A_{2,1} & A_{2,2} & \ddots & & \\ & \ddots & \ddots & & \\ & & & A_{m-1,m} & \\ & & & A_{m,m-1} & A_{m,m} \end{bmatrix}^{-1} = U_1^{-1} \dots U_{m-1}^{-1} \begin{bmatrix} T_1 & & & & \\ & T_2 & & & \\ & & \ddots & & \\ & & & \ddots & \\ & & & & T_m \end{bmatrix} L_{m-1}^{-1} \dots L_1^{-1}.$$

The left-hand side is the discrete half-space Green's function with Dirichlet zero boundary condition at $x_3 = (m+1)h$ and a straightforward calculation shows that the lower right block of the right-hand side is T_m . Therefore, T_m is the discrete half-space Green's function restricted to the m th layer. Note that the PML at $x_3 = 0$ is used to simulate an absorbing boundary condition. If we assume that there is little reflection during the transmission of the wave, we can approximate T_m by placing the PML right next to the m th layer since the domain of interest is only the m th layer (see Figure 2). This is the key idea of the moving PML sweeping preconditioner, where the operator T_m is numerically approximated by putting the PML right next to the domain of interest and solving a much smaller system to save the computational cost.

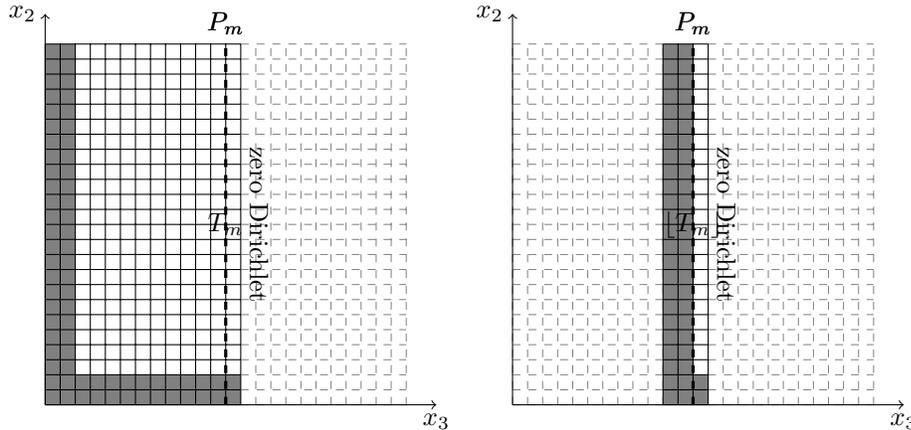


FIG. 2. Left: T_m is the restriction to P_m (the dashed grid) of the half-space Green's function on the solid grid. Right: By moving the PML right next to the layer P_m , the operator T_m is approximated by solving the equation on a much smaller grid.

More precisely, we introduce an auxiliary problem on the domain $D_m = [0, 1] \times [0, 1] \times [(m - b)h, (m + 1)h]$:

$$\begin{cases} \left((s_1 \partial_1)^2 + (s_2 \partial_2)^2 + (s_3^m \partial_3)^2 + \frac{\omega^2}{c^2(x)} \right) v(x) = g(x) & \forall x \in D_m, \\ v(x) = 0 & \forall x \in \partial D_m, \end{cases}$$

where $s_3^m(x) = s(x_3 - (m - b)h)$. The domain D_m is discretized with the partial grid

$$P_{(m-b+1):m} := \{P_t : m - b + 1 \leq t \leq m\}.$$

Applying the same central finite difference scheme gives rise to the corresponding discretized system, denoted as

$$H_m v = g, \quad m = b + 1, \dots, n.$$

To approximate T_m , we numerically define operator $[T_m] : \alpha \in \mathbb{C}^{n^2} \rightarrow \beta \in \mathbb{C}^{n^2}$ by the following procedure:

1. Introduce a vector g defined on $P_{(m-b+1):m}$ by setting α to the layer P_m and zero everywhere else.
2. Solve the discretized auxiliary problem $H_m v = g$ on $P_{(m-b+1):m}$ with g from step 1.
3. Set β as the restriction on P_m of the solution v from step 2.

The discretized system is a quasi-2D system as b is typically a small constant, so the system can be solved efficiently by the nested dissection method [12, 6, 14].

The first b layers, which are in the PML region of the original problem (2), need to be handled with a slight difference. Define

$$\begin{aligned} u_{:,1:b} &= [u_{:,1}^T, \dots, u_{:,b}^T]^T, \\ f_{:,1:b} &= [f_{:,1}^T, \dots, f_{:,b}^T]^T. \end{aligned}$$

Then the system $Au = f$ can be written as

$$\begin{bmatrix} A_{1:b,1:b} & A_{1:b,b+1} & & & \\ A_{b+1,1:b} & A_{b+1,b+1} & \ddots & & \\ & \ddots & \ddots & A_{n-1,n} & \\ & & A_{n,n-1} & A_{n,n} & \end{bmatrix} \begin{bmatrix} u_{:,1:b} \\ u_{:,b+1} \\ \vdots \\ u_{:,n} \end{bmatrix} = \begin{bmatrix} f_{:,1:b} \\ f_{:,b+1} \\ \vdots \\ f_{:,n} \end{bmatrix}.$$

For the first b layers, we simply define $[T_{1:b}]$ as the inverse operator of $H_b := A_{1:b,1:b}$. However, it is essential that $[T_{1:b}]$ is stored in a factorized form by applying the nested dissection method to H_b , since $H_b v = g$ is also a quasi-2D problem.

Based on the above discussion, the setup algorithm of the moving PML sweeping preconditioner is given in Algorithm 1.

Algorithm 1. Construction of the moving PML sweeping preconditioner of the system (3). Complexity = $O(b^3 n^4) = O(b^3 N^{4/3})$.

Construct the nested dissection factorization of H_b , which defines $[T_{1:b}]$.

for $m = b + 1, \dots, n$ **do**

Construct the nested dissection factorization of H_m , which defines $[T_m]$.

end for

Once the factorization is completed, $[T_{1:b}]$ and $[T_m]$ can be applied using the nested dissection factorization. The application process of the sweeping preconditioner is given in Algorithm 2.

Algorithm 2. Computation of $u \approx A^{-1}f$ using the factorization from Algorithm 1. Complexity = $O(b^2 n^3 \log n) = O(b^2 N \log N)$.

$u_{:,1:b} = [T_{1:b}]f_{:,1:b}$

$u_{:,b+1} = [T_{b+1}](f_{:,b+1} - A_{b+1,1:b}u_{:,1:b})$

for $m = b + 1, \dots, n - 1$ **do**

$u_{:,m+1} = [T_{m+1}](f_{:,m+1} - A_{m+1,m}u_{:,m})$

end for

for $m = n - 1, \dots, b + 1$ **do**

$u_{:,m} = u_{:,m} - [T_m](A_{m,m+1}u_{:,m+1})$

end for

$u_{:,1:b} = u_{:,1:b} - [T_{1:b}](A_{1:b,b+1}u_{:,b+1})$

4. Recursive sweeping preconditioner. Recall that the PML is also applied to the face $x_2 = 0$. Therefore, each quasi-2D auxiliary problem is itself a discretization of the Helmholtz equation with the PML specified on one side. Following the treatment in [8] for the 2D Helmholtz equation, it is natural to apply the same sweeping idea once again along the x_2 direction, instead of the nested dissection algorithm used in the previous section.

4.1. Inner sweeping. Recall that the quasi-2D subproblems of the non-recursive sweeping preconditioners are $H_m v = g, m = b, \dots, n$. Since they have essentially the same structure, it is sufficient to consider a single system $\tilde{A}v = g$, where \tilde{A} can be any of the H_m 's. Here the accent mark emphasizes that the problem under consideration is quasi-2D. To formalize the sweeping preconditioner along the

x_2 direction, we define, up to a translation,

$$\tilde{P} = \{p_{i,j,k} = (ih, jh, kh) : 1 \leq i, j \leq n, 1 \leq k \leq b\}$$

to be the discretization grid. For each $m = 1, \dots, n$, let

$$\begin{aligned} \tilde{P}_m &= \{p_{1,m,1}, \dots, p_{1,m,b}, \dots, p_{n,m,b}\}, \\ v_{:,m,:} &= [v_{1,m,1}, \dots, v_{1,m,b}, \dots, v_{n,m,b}]^T, \\ g_{:,m,:} &= [g_{1,m,1}, \dots, g_{1,m,b}, \dots, g_{n,m,b}]^T. \end{aligned}$$

For the first b layers in the x_2 direction, we also define

$$\begin{aligned} \tilde{P}_{1:b} &= \{\tilde{P}_1, \dots, \tilde{P}_b\}, \\ v_{:,1:b,:} &= [v_{:,1,:}^T, \dots, v_{:,b,:}^T]^T, \\ g_{:,1:b,:} &= [g_{:,1,:}^T, \dots, g_{:,b,:}^T]^T. \end{aligned}$$

In this section, we reorder the vectors v, g by grouping the third dimension first and applying the column-major ordering to dimensions 1 and 2:

$$\begin{aligned} v &= [v_{:,1,:}^T, \dots, v_{:,n,:}^T]^T, \\ g &= [g_{:,1,:}^T, \dots, g_{:,n,:}^T]^T. \end{aligned}$$

With this ordering, the corresponding system $\tilde{A}v = g$ is written as

$$\begin{bmatrix} \tilde{A}_{1:b,1:b} & \tilde{A}_{1:b,b+1} & & & \\ \tilde{A}_{b+1,1:b} & \tilde{A}_{b+1,b+1} & \ddots & & \\ & \ddots & \ddots & \tilde{A}_{n-1,n} & \\ & & \tilde{A}_{n,n-1} & \tilde{A}_{n,n} & \end{bmatrix} \begin{bmatrix} v_{:,1:b,:} \\ v_{:,b+1,:} \\ \vdots \\ v_{:,n,:} \end{bmatrix} = \begin{bmatrix} g_{:,1:b,:} \\ g_{:,b+1,:} \\ \vdots \\ g_{:,n,:} \end{bmatrix}.$$

For the block LDU factorization of \tilde{A} , we define

$$\begin{aligned} \tilde{S}_{1:b} &= \tilde{A}_{1:b,1:b}, \quad \tilde{T}_{1:b} = \tilde{S}_{1:b}^{-1}, \\ \tilde{S}_{b+1} &= \tilde{A}_{b+1,b+1} - \tilde{A}_{b+1,1:b} \tilde{T}_{1:b} \tilde{A}_{1:b,b+1}, \quad \tilde{T}_{b+1} = \tilde{S}_{b+1}^{-1}, \\ \tilde{S}_m &= \tilde{A}_{m,m} - \tilde{A}_{m,m-1} \tilde{T}_{m-1} \tilde{A}_{m-1,m}, \quad \tilde{T}_m = \tilde{S}_m^{-1}, \quad m = b+2, \dots, n; \end{aligned}$$

then \tilde{A} can be factorized as

$$\tilde{A} = \tilde{L}_{1:b} \tilde{L}_{b+1} \dots \tilde{L}_{n-1} \begin{bmatrix} \tilde{S}_{1:b} & & & \\ & \tilde{S}_{b+1} & & \\ & & \ddots & \\ & & & \tilde{S}_n \end{bmatrix} \tilde{U}_{n-1} \dots \tilde{U}_{b+1} \tilde{U}_{1:b},$$

where the nonzero off-diagonal blocks of the unit lower and upper triangular matrices \tilde{L}_m and \tilde{U}_m are given by

$$\begin{aligned} \tilde{L}_{1:b}(\tilde{P}_{b+1}, \tilde{P}_{1:b}) &= \tilde{A}_{b+1,1:b} \tilde{T}_{1:b}, \quad \tilde{U}_{1:b}(\tilde{P}_{1:b}, \tilde{P}_{b+1}) = \tilde{T}_{1:b} \tilde{A}_{1:b,b+1}, \\ \tilde{L}_m(\tilde{P}_{m+1}, \tilde{P}_m) &= \tilde{A}_{m+1,m} \tilde{T}_m, \quad \tilde{U}_m(\tilde{P}_m, \tilde{P}_{m+1}) = \tilde{T}_m \tilde{A}_{m,m+1}, \quad m = b+1, \dots, n-1. \end{aligned}$$

Then the solution v can be computed by

$$v = \begin{bmatrix} v_{:,1:b,:} \\ v_{:,b+1,:} \\ \vdots \\ v_{:,n,:} \end{bmatrix} = \tilde{A}^{-1}g = \tilde{U}_{1:b}^{-1}\tilde{U}_{b+1}^{-1}\dots\tilde{U}_{n-1}^{-1} \begin{bmatrix} \tilde{T}_{1:b} \\ \tilde{T}_{b+1} \\ \dots \\ \tilde{T}_n \end{bmatrix} \tilde{L}_{n-1}^{-1}\dots\tilde{L}_{b+1}^{-1}\tilde{L}_{1:b}^{-1}g.$$

By comparing the factorization of the upper left $(m - b + 1) \times (m - b + 1)$ block of \tilde{A} , where $m = b + 1, \dots, n$, we have

$$\begin{bmatrix} \tilde{A}_{1:b,1:b} & \tilde{A}_{1:b,b+1} & & \\ \tilde{A}_{b+1,1:b} & \tilde{A}_{b+1,b+1} & \dots & \\ & \dots & \dots & \tilde{A}_{m-1,m} \\ & & \tilde{A}_{m,m-1} & \tilde{A}_{m,m} \end{bmatrix} = \tilde{L}_{1:b}\tilde{L}_{b+1}\dots\tilde{L}_{m-1} \begin{bmatrix} \tilde{S}_{1:b} & & & \\ & \tilde{S}_{b+1} & & \\ & & \dots & \\ & & & \tilde{S}_m \end{bmatrix} \tilde{U}_{m-1}\dots\tilde{U}_{b+1}\tilde{U}_{1:b},$$

where \tilde{L}_t and \tilde{U}_t are redefined as their restrictions to their top left $(m - b + 1) \times (m - b + 1)$ blocks. Inverting both sides gives

$$\begin{bmatrix} \tilde{A}_{1:b,1:b} & \tilde{A}_{1:b,b+1} & & \\ \tilde{A}_{b+1,1:b} & \tilde{A}_{b+1,b+1} & \dots & \\ & \dots & \dots & \tilde{A}_{m-1,m} \\ & & \tilde{A}_{m,m-1} & \tilde{A}_{m,m} \end{bmatrix}^{-1} = \tilde{U}_{1:b}^{-1}\tilde{U}_{b+1}^{-1}\dots\tilde{U}_{m-1}^{-1} \begin{bmatrix} \tilde{T}_{1:b} \\ \tilde{T}_{b+1} \\ \dots \\ \tilde{T}_m \end{bmatrix} \tilde{L}_{m-1}^{-1}\dots\tilde{L}_{b+1}^{-1}\tilde{L}_{1:b}^{-1}.$$

Thus, by repeating the argument in section 3, the matrix \tilde{T}_m is the restriction to the layer \tilde{P}_m of the discrete half-space Green's function. It can be approximated by $\lfloor \tilde{T}_m \rfloor$, which is defined by solving a quasi-1D problem obtained by placing a moving PML right next to $x_2 = mh$ (see Figure 3). Each auxiliary quasi-1D problem in this inner sweeping step can be solved by the sparse block LDU factorization efficiently, with ordering the system by grouping dimensions 3 and 2 first and dimension 1 last.

More specifically, for each m , we introduce the auxiliary problem on the domain $\tilde{D}_m = [0, 1] \times [(m - b)h, (m + 1)h] \times [0, (b + 1)h]$:

$$\begin{cases} \left((s_1\partial_1)^2 + (s_2^m\partial_2)^2 + (s_3\partial_3)^2 + \frac{\omega^2}{c^2(x)} \right) w(x) = q(x) & \forall x \in \tilde{D}_m, \\ w(x) = 0 & \forall x \in \partial\tilde{D}_m, \end{cases}$$

where $s_2^m(x) = s(x_2 - (m - b)h)$. The domain \tilde{D}_m is discretized with the grid

$$\tilde{P}_{(m-b+1):m} := \{ \tilde{P}_t : m - b + 1 \leq t \leq m \},$$

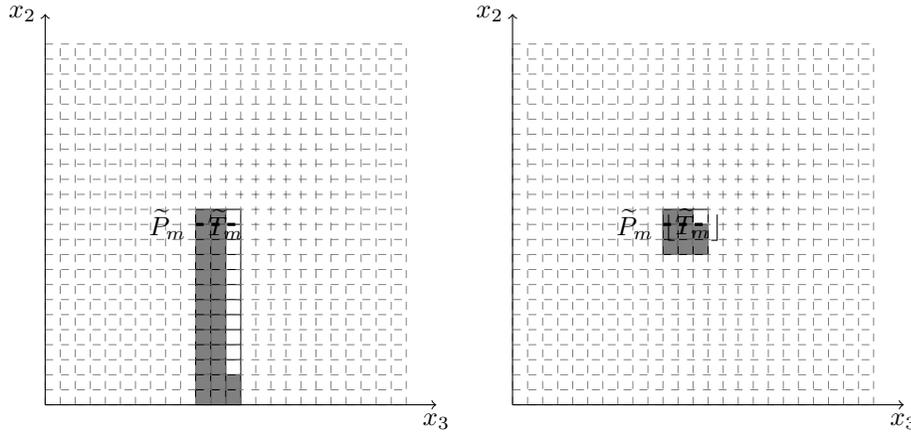


FIG. 3. Left: \tilde{T}_m is the restriction to \tilde{P}_m (the dashed grid) of the Green's function on the quasi-2D solid grid. Right: By moving the PML right next to \tilde{P}_m , the operator \tilde{T}_m is approximated by solving the problem on a quasi-1D grid.

and the same central difference numerical scheme is used here. We denote the corresponding discretized system as $\tilde{H}_m w = q$. Similar to the process described in section 3, we define the operator $[\tilde{T}_m] : \alpha \in \mathbb{C}^{nb} \rightarrow \beta \in \mathbb{C}^{nb}$ by the following procedure:

1. Introduce a vector q defined on the grid $\tilde{P}_{(m-b+1):m}$ by setting α to the layer \tilde{P}_m and zero everywhere else.
2. Solve the auxiliary quasi-1D problem $\tilde{H}_m w = q$ on $\tilde{P}_{(m-b+1):m}$ with q from step 1.
3. Set β as the restriction on \tilde{P}_m of the solution w from step 2.

For the first b layers, $[\tilde{T}_{1:b}]$ is simply defined as the inverse operator of $\tilde{H}_b := \tilde{A}_{1:b,1:b}$, which is essentially the same as $\tilde{T}_{1:b}$, but implemented by using the sparse block LDU factorization of \tilde{H}_b . Summarizing all this, the setup and application algorithm of the inner moving PML sweeping preconditioner are given in Algorithms 3 and 4, respectively.

Algorithm 3. Construction of the inner moving PML sweeping preconditioner of the quasi-2D problem $\tilde{A}v = g$. Complexity = $O(b^6 n^2)$.

Construct the sparse block LDU factorization of \tilde{H}_b , which defines $[\tilde{T}_{1:b}]$.

for $m = b + 1, \dots, n$ **do**

 Construct the sparse block LDU factorization of \tilde{H}_m , which defines $[\tilde{T}_m]$.

end for

4.2. Putting together. As we pointed out earlier, the matrix \tilde{A} can be any one of $H_m, m = b, \dots, n$, where Algorithms 3 and 4 can be applied. Notice that solving the subproblems exactly with the nested dissection algorithm results in the approximation $[T_m]$ to T_m . This extra level of approximation defines a further approximation, which shall be denoted by $\llbracket T_m \rrbracket : \alpha \in \mathbb{C}^{n^2} \rightarrow \beta \in \mathbb{C}^{n^2}$ (to be precise, for the first b layers, it is $\llbracket T_{1:b} \rrbracket : \alpha \in \mathbb{C}^{n^2 b} \rightarrow \beta \in \mathbb{C}^{n^2 b}$). The steps for carrying out $\llbracket T_m \rrbracket$ are similar to the ones for $[T_m]$ except that one uses Algorithms 3 and 4 to solve the quasi-2D problems approximately (instead of the nested dissection method that solves them exactly).

Algorithm 4. Computation of $v \approx \tilde{A}^{-1}g$ using the factorization from Algorithm 3. Complexity = $O(b^4n^2)$.

```

 $v_{:,1:b,:} = \lfloor \tilde{T}_{1:b} \rfloor g_{:,1:b,:}$ 
 $v_{:,b+1,:} = \lfloor \tilde{T}_{b+1} \rfloor (g_{:,b+1,:} - \tilde{A}_{b+1,1:b} v_{:,1:b,:})$ 
for  $m = b + 1, \dots, n - 1$  do
   $v_{:,m+1,:} = \lfloor \tilde{T}_{m+1} \rfloor (g_{:,m+1,:} - \tilde{A}_{m+1,m} v_{:,m,:})$ 
end for
for  $m = n - 1, \dots, b + 1$  do
   $v_{:,m,:} = v_{:,m,:} - \lfloor \tilde{T}_m \rfloor (\tilde{A}_{m,m+1} v_{:,m+1,:})$ 
end for
 $v_{:,1:b,:} = v_{:,1:b,:} - \lfloor \tilde{T}_{1:b} \rfloor (\tilde{A}_{1:b,b+1} v_{:,b+1,:})$ 

```

Given all these preparations, the setup algorithm of the recursive sweeping preconditioner can be summarized compactly in Algorithm 5 and the application algorithm is given in Algorithm 6.

Algorithm 5. Construction of the recursive moving PML sweeping preconditioner of the linear system (3). Complexity = $O(b^6n^3) = O(b^6N)$.

```

Construct the inner moving PML sweeping preconditioner of  $H_b$  by Algorithm 3.
This gives  $\llbracket T_{1:b} \rrbracket$ .
for  $m = b + 1, \dots, n$  do
  Construct the inner moving PML sweeping preconditioner of  $H_m$  by Algorithm 3.
  This gives  $\llbracket T_m \rrbracket$ .
end for

```

Algorithm 6. Computation of $u \approx A^{-1}f$ using the factorization from Algorithm 5. Complexity = $O(b^4n^3) = O(b^4N)$.

```

 $u_{:,:,1:b} = \llbracket T_{1:b} \rrbracket f_{:,:,1:b}$ 
 $u_{:,:,b+1} = \llbracket T_{b+1} \rrbracket (f_{:,:,b+1} - A_{b+1,1:b} u_{:,:,1:b})$ 
for  $m = b + 1, \dots, n - 1$  do
   $u_{:,:,m+1} = \llbracket T_{m+1} \rrbracket (f_{:,:,m+1} - A_{m+1,m} u_{:,:,m})$ 
end for
for  $m = n - 1, \dots, b + 1$  do
   $u_{:,:,m} = u_{:,:,m} - \llbracket T_m \rrbracket (A_{m,m+1} u_{:,:,m+1})$ 
end for
 $u_{:,:,1:b} = u_{:,:,1:b} - \llbracket T_{1:b} \rrbracket (A_{1:b,b+1} u_{:,:,b+1})$ 

```

In the outer loop of Algorithm 6, the unknowns are eliminated layer by layer in the x_3 direction. In the application of $\llbracket T_m \rrbracket$, there is the inner loop in which the unknowns in each quasi-2D problem are eliminated in the x_2 direction. The whole algorithm serves as a preconditioner for the original linear system (3). Notice that, in the recursive sweeping preconditioner, the quasi-2D problems are solved only approximately. Therefore, the overall accuracy might not be as good as the non-recursive method. But as we will show in the next section, the performance of the preconditioner is only mildly affected.

The above algorithms are described in a way to present the main ideas clearly. In the actual implementations, a couple of modifications are taken in order to maximize the efficiency:

1. For each auxiliary problem, in both the inner loop and the outer loop, several layers are processed together instead of one layer.
2. For the PML introduced in the auxiliary problems, the number of layers in the auxiliary PML region does not have to match the number of layers b used for the boundary PML at $x_2 = 0$ and $x_3 = 0$. In fact, the thickness of the auxiliary PML is typically thinner for the sake of efficiency.
3. The problem we described above has zero Dirichlet boundary conditions on the other four faces of the cube. If instead the PMLs are put on all the faces, then the sweeping preconditioner sweeps with two fronts from two opposite faces respectively and they meet in the middle with a subproblem with PML on both sides instead of only one side, as described in [8].

5. Numerical results. This section presents the numerical results to illustrate the performance of the recursive sweeping preconditioner. All algorithms are implemented in MATLAB and the tests are performed on a 2.0-GHz computer with 256 GB memory. We force MATLAB to use only one computational thread to test the sequential time cost.

5.1. Comparison with the nonrecursive method. In this subsection, we adopt the standard seven-point finite difference scheme and compare the recursive approach with the nonrecursive one in [8]. The GMRES algorithm is used as the iterative solver with the relative residual equal to 10^{-3} . As we shall see, since all tests converge within a few iterations, we can use the nonrestarted version of the GMRES algorithm. Therefore, the iteration counts reported in the following tables are the so-called inner iteration count, which is essentially equal to the number of applications of the preconditioner. Because of the effectiveness of the preconditioner, the value of the restart number is irrelevant as long as it is greater than 5.

Each quasi-2D problem is solved approximately by applying the inner sweeping preconditioner only once in order to maximize the efficiency. The velocity fields and forcing terms are kept the same as those used in [8]. The PMLs are put on all six sides of the cube $[0, 1]^3$ to simulate the Sommerfeld radiation condition.

The three velocity fields used here are as follows (see Figure 4):

- (i) a converging lens with a Gaussian profile at the center of the domain,
- (ii) a vertical waveguide with a Gaussian cross section,
- (iii) a random velocity field.

For each velocity field, the tests are performed for two external forces:

- (a) a Gaussian point source centered at $(1/2, 1/2, 1/4)$,
- (b) a Gaussian wave packet with wavelength comparable to the typical wavelength of the domain. The packet centers at $(1/2, 1/4, 1/4)$ and points to the direction $(0, 1/\sqrt{2}, 1/\sqrt{2})$.

In each test, we vary the typical wave number $\omega/(2\pi)$, study the behavior of the recursive preconditioner, and compare the results with the nonrecursive preconditioner.

In our numerical implementation, each wavelength is discretized with $q = 8$ points. The width of the PML at the boundary of the cube is $9h$, and the width of the auxiliary PML for the middle layers is $bh = 5h$. The number of layers processed in each auxiliary problem is 4.

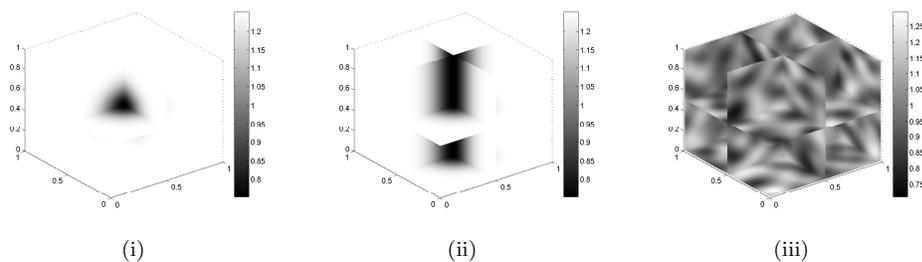
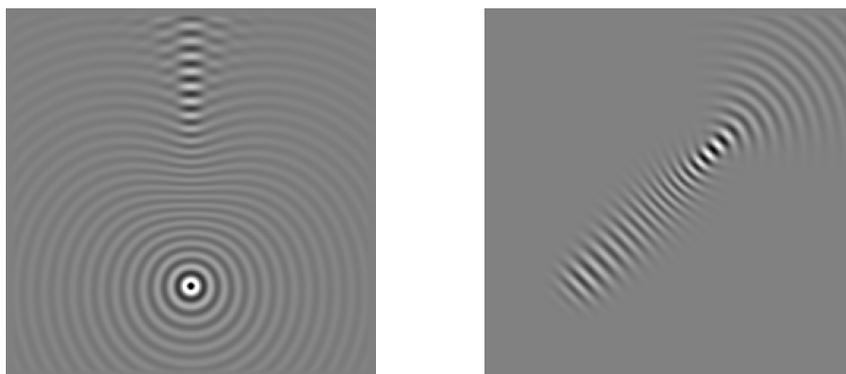


FIG. 4. The three velocity fields tested.

FIG. 5. The solutions of velocity field (i) in Figure 4 with $\omega/(2\pi) = 32$. Left: the wave field generated by force (a) at $x_1 = 0.5$. Right: the wave field generated by force (b) at $x_1 = 0.5$.TABLE 1
The results for velocity field (i) in Figure 4 with varying ω .

$\frac{\omega}{2\pi}$	N	T_{setup}			$f(x)$	N_{iter}		T_{solve}		
		NR	R	Ratio		NR	R	NR	R	Ratio
8	63^3	1.12e+02	1.63e+01	14%	(a)	3	3	1.09e+01	1.08e+01	100%
					(b)	4	4	1.41e+01	1.37e+01	97%
16	127^3	1.67e+03	1.29e+02	8%	(a)	3	4	1.15e+02	1.30e+02	113%
					(b)	4	5	1.54e+02	1.64e+02	106%
32	255^3	2.55e+04	1.10e+03	4%	(a)	4	4	1.85e+03	1.29e+03	69%
					(b)	4	5	1.89e+03	1.62e+03	86%

Figures 5, 6, and 7 plot the solution profiles under the two external forces for the three velocity fields, respectively. The numerical results are summarized in Tables 1, 2, and 3, where T_{setup} is the time used to construct the preconditioner in seconds, T_{solve} is the time used to solve the system in the preconditioned GMRES solver in seconds, and N_{iter} is the number of preconditioner applications in the iterative solving process. “NR” stands for the original nonrecursive method, while “R” stands for the recursive method introduced in this paper. The “ratio” is the time cost of the recursive method over the nonrecursive method. The numerical implementation of the nonrecursive method is slightly improved as compared to [8], by incorporating a more accurate PML discretization. Therefore, the results here for the nonrecursive

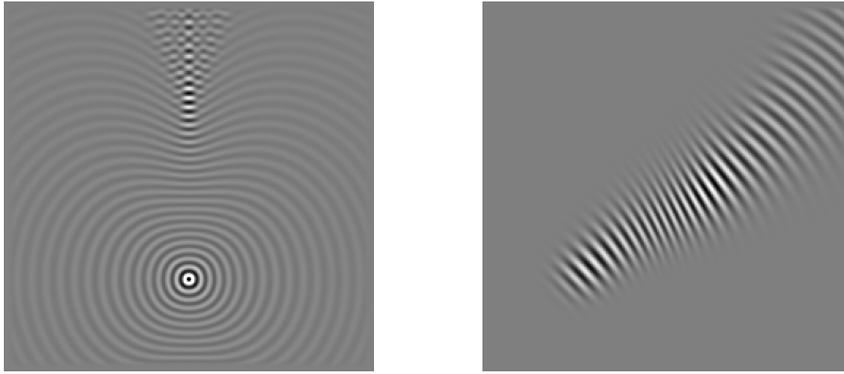


FIG. 6. The solutions of velocity field (ii) in Figure 4 with $\omega/(2\pi) = 32$. Left: the wave field generated by force (a) at $x_1 = 0.5$. Right: the wave field generated by force (b) at $x_1 = 0.5$.

TABLE 2
The results for velocity field (ii) in Figure 4 with varying ω .

$\frac{\omega}{2\pi}$	N	T_{setup}			$f(x)$	N_{iter}		T_{solve}		
		NR	R	Ratio		NR	R	NR	R	Ratio
8	63^3	1.12e+02	1.62e+01	14%	(a)	3	3	1.07e+01	1.06e+01	99%
					(b)	3	4	1.05e+01	1.37e+01	130%
16	127^3	1.65e+03	1.26e+02	8%	(a)	4	4	1.53e+02	1.29e+02	84%
					(b)	3	4	1.15e+02	1.30e+02	113%
32	255^3	2.52e+04	1.05e+03	4%	(a)	5	5	2.33e+03	1.63e+03	70%
					(b)	4	4	1.87e+03	1.32e+03	70%

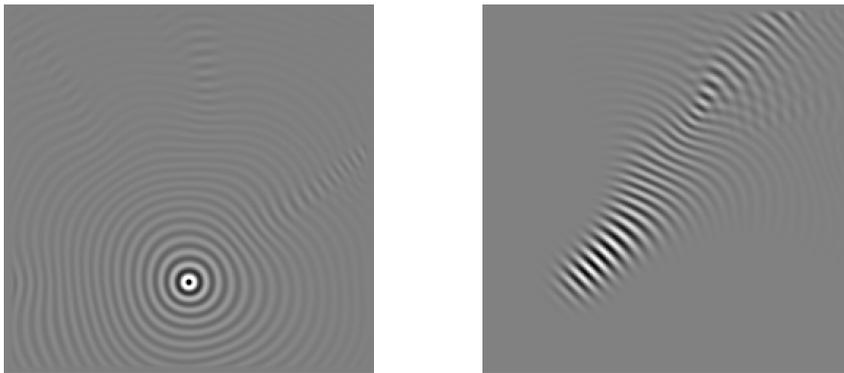


FIG. 7. The solutions of velocity field (iii) in Figure 4 with $\omega/(2\pi) = 32$. Left: the wave field generated by force (a) at $x_1 = 0.5$. Right: the wave field generated by force (b) at $x_1 = 0.5$.

method are better compared to the ones in [8].

Based on the results presented here, we can make the following observations:

1. The setup time cost of the recursive preconditioner is significantly lowered compared to the nonrecursive one. This becomes clearer as the problem size gets larger. The scaling difference between the $O(N)$ setup time cost of the

TABLE 3
The results for velocity field (iii) in Figure 4 with varying ω .

$\frac{\omega}{2\pi}$	N	T_{setup}			$f(x)$	N_{iter}		T_{solve}		
		NR	R	Ratio		NR	R	NR	R	Ratio
8	63^3	1.12e+02	1.62e+01	14%	(a)	4	4	1.42e+01	1.40e+01	98%
					(b)	4	4	1.41e+01	1.37e+01	98%
16	127^3	1.65e+03	1.28e+02	8%	(a)	4	4	1.53e+02	1.28e+02	84%
					(b)	5	5	1.92e+02	1.61e+02	84%
32	255^3	2.52e+04	1.05e+03	4%	(a)	5	5	2.31e+03	1.62e+03	70%
					(b)	5	5	2.36e+03	1.65e+03	70%

TABLE 4
Results for velocity field (ii) in Figure 4 with varying b for $\omega/(2\pi) = 16$.

$\omega/(2\pi)$	N	b	T_{setup}	$N_{\text{iter}}^{(a)}$	$T_{\text{solve}}^{(a)}$	$N_{\text{iter}}^{(b)}$	$T_{\text{solve}}^{(b)}$
16	127^3	5	4.04e+01	4	1.27e+02	4	1.28e+02
16	127^3	6	6.07e+01	4	1.68e+02	4	1.69e+02
16	127^3	7	9.00e+01	3	1.61e+02	4	2.17e+02
16	127^3	8	1.34e+02	3	2.18e+02	3	2.21e+02

TABLE 5
Results for velocity field (ii) in Figure 4 with varying b for $\omega/(2\pi) = 32$.

$\omega/(2\pi)$	N	b	T_{setup}	$N_{\text{iter}}^{(a)}$	$T_{\text{solve}}^{(a)}$	$N_{\text{iter}}^{(b)}$	$T_{\text{solve}}^{(b)}$
32	255^3	5	1.60e+02	5	1.48e+03	4	1.19e+03
32	255^3	6	2.45e+02	4	1.61e+03	4	1.64e+03
32	255^3	7	3.79e+02	4	2.11e+03	4	2.13e+03
32	255^3	8	5.72e+02	3	2.16e+03	4	2.93e+03

TABLE 6
Results for velocity field (ii) in Figure 4 with varying b for $\omega/(2\pi) = 64$.

$\omega/(2\pi)$	N	b	T_{setup}	$N_{\text{iter}}^{(a)}$	$T_{\text{solve}}^{(a)}$	$N_{\text{iter}}^{(b)}$	$T_{\text{solve}}^{(b)}$
64	511^3	5	7.80e+02	6	1.63e+04	4	1.25e+04
64	511^3	6	1.29e+03	5	2.05e+04	4	1.56e+04
64	511^3	7	1.56e+03	4	2.05e+04	4	2.15e+04
64	511^3	8	2.28e+03	4	2.87e+04	4	3.05e+04

recursive method and the $O(N^{4/3})$ cost of the nonrecursive method can be seen clearly from the results.

2. The iteration number of the recursive preconditioner increases only slightly compared to the nonrecursive one. In some cases the recursive approach needs one more iteration.
3. The application time of the recursive sweeping preconditioner is faster than the nonrecursive one when the problem size gets larger. This is also consistent with the $O(N)$ versus $O(N \log N)$ scaling difference between the recursive approach and the nonrecursive one.

5.2. Test with varying the auxiliary PML width. Next, we vary the auxiliary PML width bh to test the sensitivity of the recursive approach to the PML width. Each subproblem processes four consecutive layers. The velocity field (ii) in Figure 4 is used here so the translational invariance along the x_3 direction can be exploited to save the setup and memory cost to scale the problem size to $\omega/(2\pi) = 64$. The test

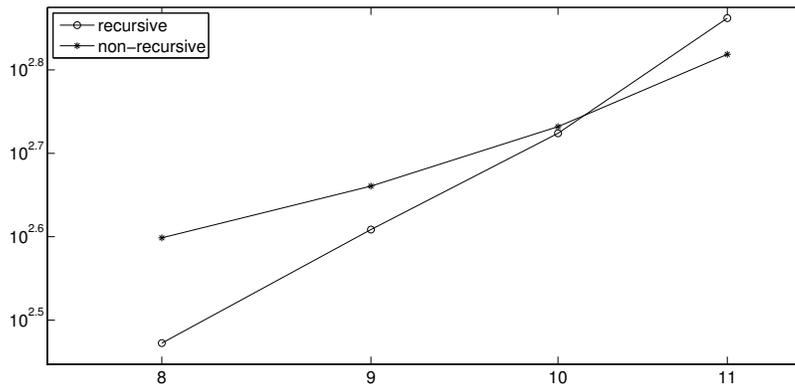


FIG. 8. A comparison of the solve time per iteration of the recursive approach and the non-recursive approach with varying b when $\omega/(2\pi) = 32$. The x -axis is $b+3$, which is the typical number of layers (including the auxiliary PML layers) processed in each subproblem, and the y -axis is the solve time per iteration in seconds. The figure shows that the recursive one is more sensitive to the auxiliary PML width.

results for different b 's are given in Tables 4 to 6. The superscripts (a) and (b) stand for the forces (a) and (b), respectively.

For the problems with the typical wave numbers up to 64, these test results show that the iteration number scales roughly logarithmically with the wave number and the width $5h$ is optimal among all the test cases in terms of both accuracy and efficiency. In terms of physical quantities, this implies that for problems up to $(64\lambda)^3$ in size, around a half wavelength for the PML width is enough. For even larger problems, we expect that at most a logarithmic increase of the PML width should be sufficient for the algorithm to converge in a logarithmic number of iterations.

We would like to point out that the complexities of the nonrecursive and recursive approaches depend quite differently on b . For the nonrecursive approach, the complexities of the construction and application algorithms are $O(b^3 N^{4/3})$ and $O(b^2 N \log N)$, respectively. For the recursive approach, the complexities are $O(b^6 N)$ and $O(b^4 N)$, respectively. Therefore, depending on the problem size N , the nonrecursive approach is more effective if the value of b is sufficiently large. Figure 8 illustrates this behavior for a fixed N and b ranging from 5 to 8.

5.3. Comparison of the spectrum of the preconditioned system in two dimensions. This subsection compares the spectrum of the preconditioned system of the recursive approach with the nonrecursive one. Since the 3D systems are too large for the entire spectrum computation, 2D systems are studied here instead. In the implementation for the 2D problems, the basic idea is the same: the nonrecursive approach solves the quasi-1D subproblems exactly, while the recursive approach solves each quasi-1D subproblem approximately with breaking it down to $O(n)$ quasi-0D (small rectangle) subproblems. The recursive approach does not help with reducing the complexity in this case, since we already have the LDU factorization to solve the quasi-1D subproblem efficiently. The purpose of this subsection is to use the comparison result in two dimensions as an analogy to the 3D case. The velocity field is chosen to be a constant field with velocity equal to 1 everywhere. The PMLs are put on all sides of the unit square. The wave number $w/(2\pi)$ is 16, the number of points per wavelength is 8, the width of the boundary PML is $9h$, the width of the auxiliary PML

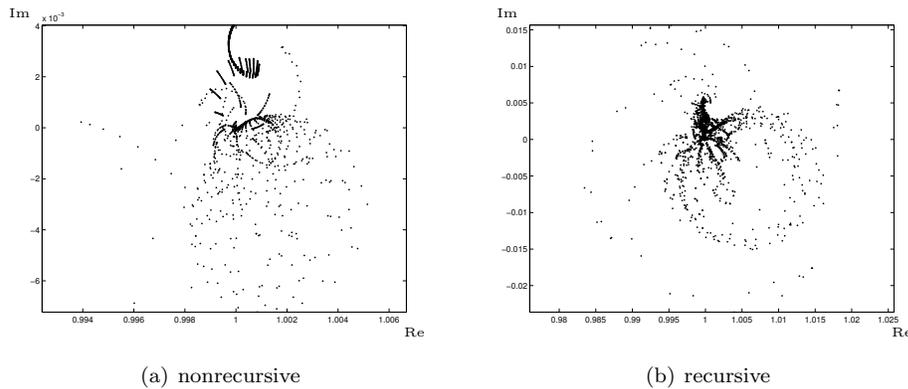


FIG. 9. The spectra of the preconditioned systems in two dimensions.

TABLE 7

Results of the compact stencil method for velocity field (i) in Figure 4 with varying ω .

$\omega/(2\pi)$	N	b	T_{setup}	$N_{\text{iter}}^{(a)}$	$T_{\text{solve}}^{(a)}$	$N_{\text{iter}}^{(b)}$	$T_{\text{solve}}^{(b)}$
8	63^3	5	3.35e+01	3	1.52e+01	4	2.00e+01
16	127^3	5	2.81e+02	4	2.51e+02	4	2.54e+02
32	255^3	5	2.34e+03	4	2.53e+03	4	2.57e+03

TABLE 8

Results of the compact stencil method for velocity field (ii) in Figure 4 with varying ω .

$\omega/(2\pi)$	N	b	T_{setup}	$N_{\text{iter}}^{(a)}$	$T_{\text{solve}}^{(a)}$	$N_{\text{iter}}^{(b)}$	$T_{\text{solve}}^{(b)}$
8	63^3	5	3.35e+01	3	1.51e+01	4	1.99e+01
16	127^3	5	2.78e+02	3	1.89e+02	4	2.52e+02
32	255^3	5	2.33e+03	4	2.51e+03	4	2.53e+03

TABLE 9

Results of the compact stencil method for velocity field (iii) in Figure 4 with varying ω .

$\omega/(2\pi)$	N	b	T_{setup}	$N_{\text{iter}}^{(a)}$	$T_{\text{solve}}^{(a)}$	$N_{\text{iter}}^{(b)}$	$T_{\text{solve}}^{(b)}$
8	63^3	5	3.32e+01	4	2.00e+01	4	2.00e+01
16	127^3	5	2.78e+02	4	2.52e+02	5	3.16e+02
32	255^3	5	2.34e+03	4	2.51e+03	5	3.21e+03

is $5h$, and the number of layers processed in each subproblem is 4. The results, given in Figure 9, show that the spectra of both the nonrecursive one and the recursive one are centered well at $1 + 0i$, though the spectrum of the recursive approach is slightly more spread out. When combined with a standard GMRES solver, the recursive approach results in slightly more iterations compared to the nonrecursive one.

5.4. Test with a compact stencil scheme. In this subsection, we implement a fourth-order compact stencil discretization of the Laplace operator to show that the recursive sweeping algorithm can be easily extended to more general and accurate discretization schemes. The numerical results are given in Tables 7 to 9.

The results demonstrate that the iteration numbers are improved slightly in some cases. The reason is that the 27-point stencil gives a smaller prefactor of the truncation error in the PML and the auxiliary PML region. As a result, its dispersion relationship

is much closer to the dispersion relationship of the Helmholtz equation. On the other hand, both the setup and solve times increase compared to the seven-point scheme due to more interactions between the unknowns.

6. Conclusion and future work. In this paper, we introduced a new recursive sweeping preconditioner for the 3D Helmholtz equation based on the moving PML sweeping preconditioner proposed in [8]. The idea of the sweeping preconditioner is used recursively for the auxiliary quasi-2D problems. Both the setup cost and the application cost of the preconditioner are reduced to strict linear complexity. The iteration number remains essentially independent of the problem size when combined with the standard GMRES solver. Numerical results show that the computational cost is reduced especially in the setup stage of the algorithm.

Several questions remain open and some potential improvements can be made. First, we use the PML to simulate the Sommerfeld condition. Many other simulations of the absorbing boundary condition can be implemented and the recursive sweeping idea can be used as long as the stencil of the simulation is local. Second, the numerical schemes used in this paper are finite difference schemes, which require more smoothness of the PML function $s(x)$ than the finite element method. Other numerical schemes such as the finite element method can be implemented and better numerical results in terms of the iteration number are expected.

Parallel processing can also be introduced to the current recursive method. First, when sweeping from both sides of the domain, either in the outer loop of the algorithm or in the inner loop, the processing of the two fronts can be paralleled so in total it could be four times faster with parallelization theoretically. Second, the quasi-1D problems are solved by the block LDU factorization in the current setting. If instead we use the 1D nested dissection algorithm for the quasi-1D problems, then it can be easily paralleled and the total cost will remain essentially the same. Last, one can notice that the setup process of the algorithm is essentially $O(n^2)$ quasi-1D subproblems which are independent from each other, so this process can be done in parallel, and compared to the original method, which contains only $O(n)$ quasi-2D independent subproblems, the potential advantages of parallelization in the setup stage is more obvious here.

There are also several advantages of the recursive sweeping method that concern flexibility. First, as mentioned above, the setup process contains $O(n^2)$ quasi-1D independent subproblems. So if the velocity field is modified on a subdomain which involves only limited subproblems, then the factorization can be updated with only a slight modification on these involved subproblems. Compared to the original method, where the subproblems are $O(n)$ quasi-2D plates, the recursive method is more flexible on updating the factorization. This could be advantageous in seismic imaging where the velocity field is tested and modified frequently. Second, when the factorization for the $O(n^2)$ subproblems is done, there are naturally two ways of using the factorization. One is, as mentioned in this paper, sweeping along the x_3 direction in the outer loop and sweeping along the x_2 direction in the inner loop. Another choice is to do the opposite, which is sweeping along the x_2 direction in the outer loop and along the x_3 direction in the inner loop. Each of these two choices shows some “bias” since the residual of the system is accumulated in some “chosen” order. So one may ask whether it is possible to combine the two choices together to make the solve process more flexible such that the total solve time can be even less. This is another interesting question to be examined.

Acknowledgments. We thank Lenya Ryzhik for providing computing resources and thank Laurent Demanet and Paul Childs for helpful discussions.

REFERENCES

- [1] J.-P. BERENGER, *A perfectly matched layer for the absorption of electromagnetic waves*, J. Comput. Phys., 114 (1994), pp. 185–200, doi:10.1006/jcph.1994.1159.
- [2] H. CALANDRA, S. GRATTON, X. PINEL, AND X. VASSEUR, *An improved two-grid preconditioner for the solution of three-dimensional Helmholtz problems in heterogeneous media*, Numer. Linear Algebra Appl., 20 (2013), pp. 663–688, doi:10.1002/nla.1860.
- [3] 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, doi:10.1137/130917144.
- [4] 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, doi:10.4208/nmtma.2013.1217nm.
- [5] W. C. CHEW AND W. H. WEEDON, *A 3D perfectly matched medium from modified Maxwell's equations with stretched coordinates*, Microw. Opt. Techn. Let., 7 (1994), pp. 599–604, doi:10.1002/mop.4650071304.
- [6] I. S. DUFF AND J. K. REID, *The multifrontal solution of indefinite sparse symmetric linear equations*, ACM Trans. Math. Software, 9 (1983), pp. 302–325, doi:10.1145/356044.356047.
- [7] B. ENGQUIST AND L. YING, *Sweeping preconditioner for the Helmholtz equation: Hierarchical matrix representation*, Comm. Pure Appl. Math., 64 (2011), pp. 697–735, doi:10.1002/cpa.20358.
- [8] B. ENGQUIST AND L. YING, *Sweeping preconditioner for the Helmholtz equation: Moving perfectly matched layers*, Multiscale Model. Simul., 9 (2011), pp. 686–710, doi:10.1137/100804644.
- [9] Y. A. ERLANGGA, *Advances in iterative methods and preconditioners for the Helmholtz equation*, Arch. Comput. Methods Eng., 15 (2008), pp. 37–66, doi:10.1007/s11831-007-9013-7.
- [10] O. G. ERNST AND M. J. GANDER, *Why it is difficult to solve Helmholtz problems with classical iterative methods*, in Numerical Analysis of Multiscale Problems, Lect. Notes Comput. Sci. Eng. 83, Springer, Heidelberg, 2012, pp. 325–363, doi:10.1007/978-3-642-22061-6_10.
- [11] M. J. GANDER AND F. NATAF, *AILU for Helmholtz problems: A new preconditioner based on the analytic parabolic factorization*, J. Comput. Acoust., 9 (2001), pp. 1499–1506, doi:10.1016/S0764-4442(00)01632-3.
- [12] A. GEORGE, *Nested dissection of a regular finite element mesh*, SIAM J. Numer. Anal., 10 (1973), pp. 345–363, doi:10.1137/0710032.
- [13] S. G. JOHNSON, *Notes on Perfectly Matched Layers (PMLs)*, Lecture notes, MIT, Cambridge, MA, 2008.
- [14] J. W. H. LIU, *The multifrontal method for sparse matrix solution: theory and practice*, SIAM Rev., 34 (1992), pp. 82–109, doi:10.1137/1034004.
- [15] 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, doi:10.1137/120871985.
- [16] C. C. STOLK, *A rapidly converging domain decomposition method for the Helmholtz equation*, J. Comput. Phys., 241 (2013), pp. 240 – 252, doi:10.1016/j.jcp.2013.01.039.
- [17] P. TSUJI, B. ENGQUIST, AND L. YING, *A sweeping preconditioner for time-harmonic Maxwell's equations with finite elements*, J. Comput. Phys., 231 (2012), pp. 3770–3783, doi:10.1016/j.jcp.2012.01.025.
- [18] P. TSUJI, J. POULSON, B. ENGQUIST, AND L. YING, *Sweeping preconditioners for elastic wave propagation with spectral element methods*, ESAIM Math. Model. Numer. Anal., 48 (2014), pp. 433–447, doi:10.1051/m2an/2013114.
- [19] P. TSUJI AND L. YING, *A sweeping preconditioner for Yee's finite difference approximation of time-harmonic Maxwell's equations*, Front. Math. China, 7 (2012), pp. 347–363, doi:10.1007/s11464-012-0191-8.
- [20] 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, doi:10.1016/j.jcp.2014.02.015.
- [21] L. ZEPEDA-NÚÑEZ AND L. DEMANET, *The method of polarized traces for the 2D Helmholtz equation*, arXiv:1410.5910 [math.NA], 2014.