



L-Sweeps: A scalable, parallel preconditioner for the high-frequency Helmholtz equation

Matthias Taus^{a,b,*}, Leonardo Zepeda-Núñez^{c,d,e}, Russell J. Hewett^{f,g}, Laurent Demanet^a

^a Department of Mathematics, Massachusetts Institute of Technology, United States of America

^b Institute of Analysis and Scientific Computing, Vienna University of Technology, Austria

^c Computational Research Division, Lawrence Berkeley National Laboratory, United States of America

^d Department of Mathematics, University of California Berkeley, United States of America

^e Department of Mathematics, University of Wisconsin-Madison, United States of America

^f Total E&P Research and Technology, United States of America

^g Department of Mathematics, Virginia Tech, United States of America



ARTICLE INFO

Article history:

Received 5 September 2019

Received in revised form 6 May 2020

Accepted 7 May 2020

Available online 10 July 2020

Keywords:

Helmholtz equation

Fast high-frequency solvers

High performance computing

Wave propagation

ABSTRACT

We present the first fast solver for the high-frequency Helmholtz equation that scales optimally in parallel for a single right-hand side. The L-sweeps approach achieves this scalability by departing from the usual propagation pattern, in which information flows in a 180° degree cone from interfaces in a layered decomposition. Instead, with L-sweeps, information propagates in 90° cones induced by a Cartesian domain decomposition (CDD). We extend the notion of accurate transmission conditions to CDDs and introduce a new sweeping strategy to efficiently track the wave fronts as they propagate through the CDD. The new approach decouples the subdomains at each wave front, so that they can be processed in parallel, resulting in better parallel scalability than previously demonstrated in the literature. The method has an overall $\mathcal{O}((N/p)\log\omega)$ empirical runtime for $N = n^d$ total degrees-of-freedom in a d -dimensional problem, frequency ω , and $p = \mathcal{O}(n)$ processors. We introduce the algorithm and provide a complexity analysis for our parallel implementation of the solver. We corroborate all claims in several two- and three-dimensional numerical examples involving constant, smooth, and discontinuous wave speeds.

© 2020 Elsevier Inc. All rights reserved.

1. Introduction

In this manuscript we consider the Helmholtz equation with variable wave speed and constant density on an open domain Ω_{bulk} with absorbing boundary conditions,

$$\begin{aligned}
 -\Delta u - \omega^2 m u &= f \quad \text{in } \Omega_{\text{bulk}}, \\
 + \text{A.B.C on } \partial\Omega_{\text{bulk}}, &
 \end{aligned}
 \tag{1}$$

* Corresponding author.

E-mail address: matthias.taus@gmail.com (M. Taus).

where for $x \in \Omega_{\text{bulk}}$, $m(x) = 1/c^2(x)$ is the squared slowness for the p-wave speed c , u is the solution wavefield, ω is the characteristic frequency, and f is the source density. For simplicity, we consider Ω_{bulk} to be a square or a cube. While we consider only problems of this form in the ensuing developments, the method we propose is a framework that can be applied to other time-harmonic formulations that model more complex physics with ease.

The Helmholtz equation arises in modeling many physical phenomena, including electromagnetic and subsurface wave propagation, and it is the cornerstone of many modern algorithmic pipelines of inverse problems, especially in exploration geophysics [87,68]. Within those pipelines, the finest recoverable detail is determined by the highest frequency for which (1) can be solved. Thus, the efficient solution of time-harmonic wave equations at high-frequency is extremely important in scientific and industrial applications. Consequently, we consider discretizations of (1) in the high-frequency regime, which means that the characteristic mesh-size of the spatial discretization has to scale as $1/\omega$.

Solving the high-frequency Helmholtz equation is notoriously difficult, mainly due to two issues:

- the accuracy of the numerical solution deteriorates as the frequency increases. To solve this issue, usually called pollution error, one needs to either over-sample the discretization, or to use specialized discretization schemes [44,45,60],
- the spectrum of the resulting matrix deteriorates as the frequency increases [61,74], thus rendering the numerical solution of the linear system prohibitively expensive for large frequencies.

These issues are a primary driver of current research in this area [29,58,82]. In this manuscript we focus on the second issue and, for clarity and conciseness in the development, we present our new method using a simple low-order finite difference discretization. However, to address the first issue, our method can easily be extended to discretizations that mitigate pollution error¹ [30,82,93] increasing the overall cost only by polylogarithmic factors.

A number of methods for solving the resulting linear systems are available, including direct methods (e.g., [88]), domain decomposition methods (e.g., [84]), and preconditioned iterative methods, including multigrid and shifted-Laplacian methods (e.g., [72,39]). We consider a solver or a preconditioner to be *sequentially scalable* if, up to logarithmic factors, it can compute a solution in $\mathcal{O}(N)$ run-time in a sequential computational environment, where N is the total number of degrees-of-freedom in the discrete problem. We consider a solver or a preconditioner to be *parallel scalable* if, up to logarithmic factors, it can compute a solution in $\mathcal{O}(N/p)$ run-time in a parallel computational environment with $p = \mathcal{O}(n)$ processors for a single right-hand side. Here, n is the number of degrees of freedom in one spatial direction, i.e., $n = \mathcal{O}(N^{\frac{1}{d}})$ where d is the problem dimension. In this paper, we present what we believe to be the first *parallel scalable* preconditioner for the high-frequency Helmholtz equation.

Currently, no scalable direct method, sequential or parallel, is available for the high-frequency regime. Standard domain decomposition methods (DDMs) localize the problem to subdomains and transfer information between subdomains. Domain decomposition methods can be applied as a direct solver (e.g., in the context of optimal Schwarz methods [35,33]), but the resulting solver is not scalable. Domain decomposition methods can yield scalable preconditioners, both in parallel and sequentially. However, the resulting preconditioned iterative solver is not scalable because it requires $\mathcal{O}(\omega)$ iterations. This shows an important aspect for the construction of solvers based on preconditioning: a quality preconditioner has to be scalable but also the resulting iterative method has to converge with limited growth in iterations as ω increases. In this regard, we call a preconditioner *effective* if the resulting preconditioned system can be solved in $\mathcal{O}(\log \omega)$ iterations. Classical DDMs exhibit sub-optimal behavior for two reasons. First, artificial and spurious reflections are induced by imprecise information transfer between subdomains. Second, long-range wave-material interactions are not tracked consistently.

Sweeping preconditioners were introduced [26,86,16,91,32] to alleviate these drawbacks, while preserving the advantages of classical DDMs. Sweeping preconditioners make use of layered domain decompositions, accurate transmission conditions, and a layer-by-layer sweeping strategy. In particular, the layered domain decomposition provides scalability by controlling growth in computational cost and memory footprint. The accurate transmission conditions allow information to flow between the subdomains without numerical artifacts, e.g., artificial reflections. Finally, the sweeping strategy consistently tracks and propagates long-range wave-material interactions. The resulting approach can be interpreted as an approximate block-LU factorization, where the blocks correspond to the local problem in each subdomain. In particular, using sparse direct solvers on blocks that arise from sufficiently thin layers yields a method with quasi-linear (i.e., linear with polylogarithmic factors) sequential complexity for a single right-hand side [26,25,92,55,77,16,86,59]. In the presence of many right-hand sides, the layered domain decomposition allows for optimal parallelization via pipelining [95,86].

Efforts to improve the performance of sweeping preconditioners have focused on obtaining better and more accurate factorizations [86], reducing the overall cost by restricting the unknowns to the interfaces [91,92], and accelerating the computation on local subproblems (blocks in the approximate factorization) with compression or parallelization [67]. Several approaches aim specifically to sparsify those blocks, thus decreasing the sequential costs [55,92]. However, while leveraging parallelism to accelerate the solve for local blocks decreases the run-times, it does not result in a parallel scalable solver for a single right-hand side.

¹ The method of polarized traces, which is the parent method of the new approach, has already been adapted to minimize pollution error, and as it will be made apparent in the sequel, the current method can be adapted in a similar way.

The main bottleneck of current sweeping preconditioners is a lack of parallel scalability for a single right-hand side. The difficulty arises due to the inherently sequential nature of the Helmholtz problem, which is independent of the domain decomposition strategy. This sequential nature is dictated by the hyperbolic behavior of wave equations and is manifested in the need to accurately resolve the long-range interactions. The key to accurately resolving these interactions is a consistent information transfer between subdomains, which has precluded more general domain decompositions (i.e., beyond layered subdomains) and consequently inhibited parallelization.

In this work, we address these issues by departing from standard layered domain decompositions and introducing a Cartesian domain decomposition (CDD) along with a new sweeping strategy that is only viable on such domain decompositions. Our novel sweeping preconditioner consistently and efficiently tracks the wavefield across subdomains. The preconditioner can be interpreted as an approximate LU factorization, where parallelism arises because the diagonal blocks themselves have a block-diagonal structure. For a single right-hand side, the resulting algorithm has a $\mathcal{O}(N/p)$ run-time, up to logarithmic factors, where, independent of geometric dimension, $p = \mathcal{O}(n)$ is the number of processors. The new algorithm therefore provides the first parallel scalable solver for the Helmholtz equation at high-frequency for a single right-hand side.

Our approach can be applied to two- and three-dimensional problems. In two dimensions, the algorithm exhibits good weak parallel scalability. That is, as the frequency increases, and consequently the problem is refined, we refine the CDD so that the local problems in each subdomain have constant size. In three dimensions, we apply the same strategy as for the two-dimensional case in two of the spatial dimensions and extend the subdomains along the third spatial dimension resulting in beam-shaped quasi-one-dimensional local problems. In either case, we employ off-the-shelf direct solvers to solve the local problems and obtain parallel complexities of $\mathcal{O}(n \log \omega)$ and $\mathcal{O}(n^2 \log \omega)$ in two and three dimensions, respectively. The 3D complexity can be reduced to $\mathcal{O}(n \log \omega)$ by employing existing parallel direct solvers [69,52] for the quasi-one-dimensional problems in each subdomain, but we do not exploit these tools here.

The new algorithm therefore results in an $\mathcal{O}(n \log \omega)$ complexity for a single right-hand side, regardless of the geometric dimension. This new algorithm requires sweeps across the CDD in both cardinal and diagonal directions, and we can exploit parallelism in all directions orthogonal to the current sweep direction. While this new strategy increases parallelization, the sweeps are still inherently sequential and cannot be parallelized. Thus, we do not anticipate further reduction in complexity for a single right-hand, due exclusively to modification of the domain decomposition. In the presence of $\mathcal{O}(n)$ right-hand-sides, it has been shown, however, that the sequential nature of the sweeps can be mitigated by pipelining multiple right-hand sides [95]. Using pipelining, solutions for all $\mathcal{O}(n)$ right-hand-sides can be computed with $\mathcal{O}(n \log \omega)$ parallel complexity, i.e., the average parallel complexity per right-hand-side is $\mathcal{O}(\log \omega)$.

As is customary in the development of new preconditioners or solution strategies, we focus on the parallel scalability of the algorithm and we assume that the accuracy of the solution stems from the chosen discretization. However, we give further exposition on achieving a parallel scalable and accurate algorithm in Section 6.

1.1. Related work

Our method is inspired by the method of polarized traces, which is well-established in the literature and, in contrast to other proposed preconditioners, has been proven applicable for problems with high-order discretizations [93,78] and highly heterogeneous (and even discontinuous) wave speed distributions [92].

Standard linear algebra techniques such as nested dissection [40] and multi-frontal solvers [24,89], coupled with \mathcal{H} -matrices [9], have been applied to the Helmholtz problem [42,19,88,1]. While these methods take advantage of compressed linear algebra to gain more efficiency (e.g., [6]), in the high-frequency regime they still suffer from the same sub-optimal asymptotic complexity as standard multi-frontal methods (e.g., [22,2,18]).

Multigrid methods (e.g., [50,13,28,72,62,28,3]), once thought to be inefficient for the Helmholtz problem, have been successfully applied [14,47,76]. Approaches stemming from the complex-shifted Laplacian [28] can be advantageous if properly tuned. However, in general, they either require an expensive solver for the shifted problem or require a large number of iterations to reach convergence, depending on the scaling between the complex shift and the frequency [39]. Although these algorithms do not result in a lower computational complexity, they are highly parallelizable, resulting in low run-times.

Within the geophysical community, the analytic incomplete LU (AILU) method was explored in [65,64]. A variant of Kaczmarz preconditioners [43] has been studied and applied to time-harmonic wave equations by [53]. Another class of methods, called hybrid direct-iterative methods, have been explored by [73]. Although these solvers have, in general, relatively low memory consumption they tend to require many iterations to converge, thus hindering practical run-times.

Domain decomposition methods for solving partial differential equations (PDEs) have a long history [71,54]. The first application of domain decomposition to the Helmholtz problem was proposed by [23], which inspired the development of various domain decomposition algorithms, which are now classified as Schwarz algorithms.² However, the convergence rate of such algorithms is strongly dependent on the boundary conditions prescribed at the interfaces between subdomains [35]. The subsequent introduction of the optimized Schwarz framework in [33], which uses optimized boundary conditions to obtain good convergence, has inspired several competing approaches, including, but not limited to [34,11,37,38,36].

² For a review on classical Schwarz methods see [15,84]; and for other applications of domain decomposition methods for the Helmholtz equations, see [20,41,57,17,56,10,4].

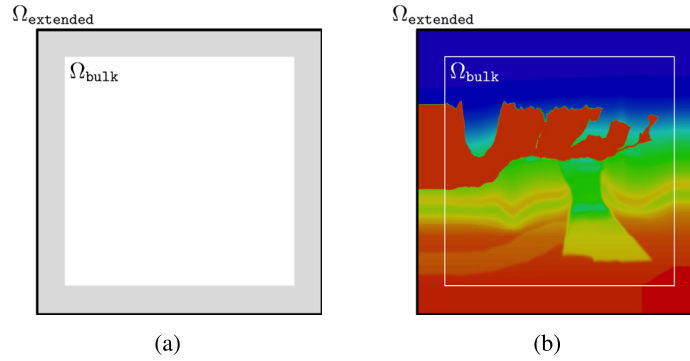


Fig. 1. (a) Ω_{bulk} (white) and Ω_{extended} (white and gray) and (b) part of the BP 2004 [8] wave speed (Ω_{bulk}) extended to Ω_{extended} via the normal extension.

Absorbing boundary conditions for domain decomposition schemes for elliptic problems are introduced by [27] and the first application of such techniques to the Helmholtz problem traces back to the AILU factorization [31]. The sweeping preconditioner, introduced in [25,26], was the first algorithm to show that those ideas could yield algorithms with quasi-linear complexity, leading to several related algorithms with similar claims such as the source transfer preconditioner [16], the rapidly converging domain decomposition [75] and its extensions [77], the double sweep preconditioner [86,63], and the method of polarized traces [91]. For an extensive review on sweeping-type methods we direct the reader to [32].

To our knowledge, all domain decomposition methods and sweeping preconditioners have been based on layered domain decompositions, with one exception. In [51], a $q \times q$ CDD is considered in the context of the source transfer method. Their proposed strategy requires $p = q^2$ processors to obtain the sub-optimal parallel complexity of $\mathcal{O}((N/\sqrt{p}) \log N)$, in both 2D and 3D. The primary difference between this method and our method is the sweeping strategy. The strategy employed in [51] is inspired by classical domain decomposition methods, while our method is inspired by sweeping preconditioners. Consequently, our method inherits a superior sweeping strategy.

1.2. Model problem and discretization

We consider problems formulated within the framework of (1), where we choose to model absorbing boundary conditions using perfectly matched layers (PMLs) [7,48]. To preserve the solution in all of Ω_{bulk} , using a PML requires the domain to be extended, which in turn also implies that the solution, material property, and source spaces must also be extended. The extended domain, Ω_{extended} , contains all of Ω_{bulk} , as illustrated in Fig. 1a and a new boundary value problem is formulated on Ω_{extended} ,

$$\begin{aligned} -\operatorname{div}(\Lambda \nabla u) - \omega^2 m u &= f & \text{in } \Omega_{\text{extended}}, \\ u &= 0 & \text{on } \partial \Omega_{\text{extended}}, \end{aligned} \quad (2)$$

where, for brevity, we re-use the symbols u , m , and f to represent the extended solution wavefield, slowness, and source distribution. As is customary for PMLs in domains with varying wave speed, we assume that the source density, f , is extended by zero into Ω_{extended} and the squared slowness, m , is extended into Ω_{extended} along the normal direction of $\partial \Omega_{\text{bulk}}$ in a constant fashion, as illustrated in Fig. 1b. The PML-extended Helmholtz equation (2) is reduced to (1) in Ω_{bulk} by imposing that Λ is the identity matrix in Ω_{bulk} . In the PML region, $\Omega_{\text{extended}} \setminus \Omega_{\text{bulk}}$, Λ is a complex valued diagonal matrix which depends on the PML formulation, and m and f are complex-valued functions, obtained from imposing the PML. Details of the precise formulation used in our developments are provided in Appendix A. Equation (2) is effectively a complex-valued boundary-value problem in Ω_{extended} with homogeneous Dirichlet boundary conditions. In the remainder of the discussion, we consider (2) to be the canonical problem and therefore, for simplicity of notation, denote Ω_{extended} as Ω .

When (2) is discretized, we obtain the linear algebraic system

$$\mathbf{A} \mathbf{u} = \mathbf{f} \quad (3)$$

where \mathbf{A} is the model-dependent system matrix, \mathbf{u} is the solution vector, and \mathbf{f} is the vector of the source density f . In this paper we restrict our discussion of absorbing boundary conditions to PMLs and discretization to finite-difference methods. These restrictions are merely to simplify the exposition: other transparent boundary conditions, such as absorbing layers or sponge layers, and other discretizations, such as higher-order finite differences and those derived from finite element methods, may be used in this framework. For 2D problems we use standard 5-point finite difference stencils and in 3D, we use 9-point stencils. As a result, the discretization has N global degrees-of-freedom and $n = N^{\frac{1}{2}}$ and $n = N^{\frac{1}{3}}$ degrees-of-freedom in each spatial dimension for two- and three-dimensional problems respectively.

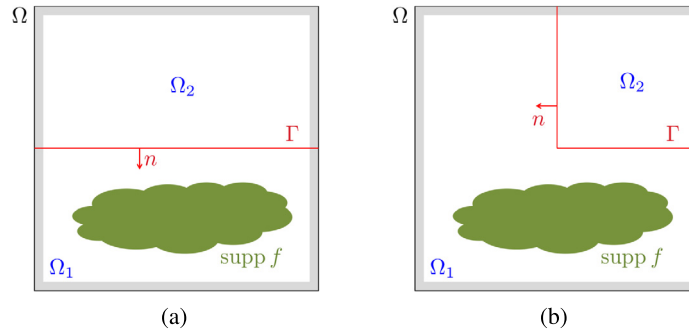


Fig. 2. A schematic representation of the truncated (a) half- and (b) quarter-space problem.

For the PMLs we use a cubic PML profile function. As we increase the frequency of the problem, we do not increase the width of the PML. Instead, we choose the PML width so that the number of wavelengths is constant in the PML region, and increase the absorption constant logarithmically with the frequency. This is motivated by an analysis of the PML [12], where this choice is rigorously justified for a more complex PML-profile. In our work, we employ the same strategy for the cubic PML-profile and obtain satisfactory results. Details on the construction of the linear system and the PMLs are given in Appendix A.

1.3. Continuous polarization

The method of polarized traces was introduced as a solver [91] and then as a preconditioner [94] for the linear system (3). At its core, the method of polarized traces spatially subdivides the discrete degrees-of-freedom in Ω into layers and computes an approximate solution to the global wavefield by sweeping over the layers and solving a local discrete half-space problem in each layer. Following [91], the solutions of the half-space problems are called polarized wavefields. In this work, we make extensive use of this concept and therefore provide a brief review, in the continuous setting, in this section. In Section 1.4, we present a similar treatment in the discrete context.

Consider the boundary-value problem (2) with a source density function f supported only in Ω_{bulk} . Let Γ be an interface dividing Ω into two regions, Ω_1 and Ω_2 , such that the support of f lies entirely within Ω_1 . For example, Γ could be a straight line (Fig. 2a) or an L-shaped line (Fig. 2b).

The wavefield in Ω_2 can be computed from the representation formula

$$u(x) = \int_{\Gamma} \mu(y)G(x, y)ds_y - \int_{\Gamma} \lambda(y) [n(y) \cdot (\Lambda(y)\nabla_y G(x, y))] ds_y \quad x \in \Omega_2, \tag{4}$$

where

$$\lambda = u|_{\Gamma}, \quad \text{and} \quad \mu = [n \cdot (\Lambda \nabla u)]|_{\Gamma}$$

are the Dirichlet and Neumann traces on Γ , and $G(x, y)$ is the Green's function corresponding to the problem (2), i.e., for $x \in \Omega$

$$\begin{aligned} -\operatorname{div}_y (\Lambda(y)\nabla_y G(y, x)) - \omega^2 m(y)G(y, x) &= \delta(x - y) \quad y \in \Omega, \\ G(y, x) &= 0 \quad y \in \partial\Omega. \end{aligned}$$

This formula directly follows from the divergence theorem and properties of the Green's function, assuming that Λ is a diagonal matrix.

Equation (4) requires knowledge of the Dirichlet and Neumann traces of the solution almost everywhere on Γ . Thus, corners in Γ are admissible, as in the quarter-space problem, even though the normal is not uniquely defined. Using (4), the solution u can then be computed on Ω_2 . Extending equation (4) to all of Ω ,

$$U(x) = \int_{\Gamma} \mu(y)G(x, y)ds_y - \int_{\Gamma} \lambda(y) [n(y) \cdot (\Lambda(y)\nabla_y G(x, y))] ds_y \quad x \in \Omega, \tag{5}$$

it can be shown that U vanishes on Ω_1 ,

$$U(x) = 0 \quad \text{for } x \in \Omega_1. \tag{6}$$

Following the terminology of [91], we call (6) the annihilation condition, Γ the polarization interface, and U a polarized wavefield. A proof of the annihilation condition (6) is provided in Appendix B.




Fig. 3. A schematic representation of four sets of degrees-of-freedom Ω_1 , Ω_2 , Γ_1 and Γ_2 .

1.4. Discrete polarization

A discrete counterpart of the polarized wavefield U can be derived by constructing a discrete solution that satisfies the discrete analogue to the annihilation condition in (6). Consider the discretization points corresponding to degrees-of-freedom in (3), as well as an interface Γ which does not intersect with any discretization point, as illustrated in Fig. 3 for both the half-space and quarter-space subdomains. Given a discretization-dependent distance δ (e.g., $\delta = 1$ for a classical 5-point finite difference stencil in 2D), such an interface divides the degrees-of-freedom into four sets, as labeled in Fig. 3:

1. Γ_1 , the set of all degrees-of-freedom physically contained in Ω_1 and δ -adjacent to Γ ;
2. Ω_1 , the set of all degrees-of-freedom physically contained in Ω_1 , excluding Γ_1 ;
3. Γ_2 , the set of all degrees-of-freedom physically contained in Ω_2 and δ -adjacent to Γ ;
4. Ω_2 , the set of all degrees-of-freedom physically contained in Ω_2 , excluding Γ_2 .

Upon reordering the degrees-of-freedom with respect to these four sets, the discrete system (3) can be rewritten as

$$\begin{pmatrix} \mathbf{A}_{\Omega_1\Omega_1} & \mathbf{A}_{\Omega_1\Gamma_1} & \mathbf{0} & \mathbf{0} \\ \mathbf{A}_{\Gamma_1\Omega_1} & \mathbf{A}_{\Gamma_1\Gamma_1} & \mathbf{A}_{\Gamma_1\Gamma_2} & \mathbf{0} \\ \mathbf{0} & \mathbf{A}_{\Gamma_2\Gamma_1} & \mathbf{A}_{\Gamma_2\Gamma_2} & \mathbf{A}_{\Gamma_2\Omega_2} \\ \mathbf{0} & \mathbf{0} & \mathbf{A}_{\Omega_2\Gamma_2} & \mathbf{A}_{\Omega_2\Omega_2} \end{pmatrix} \begin{pmatrix} \mathbf{u}_{\Omega_1} \\ \mathbf{u}_{\Gamma_1} \\ \mathbf{u}_{\Gamma_2} \\ \mathbf{u}_{\Omega_2} \end{pmatrix} = \begin{pmatrix} \mathbf{f}_{\Omega_1} \\ \mathbf{f}_{\Gamma_1} \\ \mathbf{f}_{\Gamma_2} \\ \mathbf{f}_{\Omega_2} \end{pmatrix}, \quad (7)$$

where $\mathbf{A}_{\Gamma_1\Omega_1}$ denotes the slice of the matrix \mathbf{A} generated from the rows corresponding to Γ_1 and the columns corresponding to Ω_1 . The remaining submatrices are similarly generated. In the same way, the vector \mathbf{u}_{Ω_2} (and similarly all other vectors) denotes the slice of the vector \mathbf{u} with respect to Ω_2 . As in the continuous case, we consider the case without sources in Ω_2 , i.e., $\mathbf{f}_{\Omega_2} = \mathbf{0}$ and $\mathbf{f}_{\Gamma_2} = \mathbf{0}$. Then, due to the invertibility of \mathbf{A} , it is easy to see that the solution \mathbf{U} of the linear system

$$\begin{pmatrix} \mathbf{A}_{\Omega_1\Omega_1} & \mathbf{A}_{\Omega_1\Gamma_1} & \mathbf{0} & \mathbf{0} \\ \mathbf{A}_{\Gamma_1\Omega_1} & \mathbf{A}_{\Gamma_1\Gamma_1} & \mathbf{A}_{\Gamma_1\Gamma_2} & \mathbf{0} \\ \mathbf{0} & \mathbf{A}_{\Gamma_2\Gamma_1} & \mathbf{A}_{\Gamma_2\Gamma_2} & \mathbf{A}_{\Gamma_2\Omega_2} \\ \mathbf{0} & \mathbf{0} & \mathbf{A}_{\Omega_2\Gamma_2} & \mathbf{A}_{\Omega_2\Omega_2} \end{pmatrix} \mathbf{U} = \begin{pmatrix} \mathbf{0} \\ \mathbf{A}_{\Gamma_1\Gamma_2} \mathbf{u}_{\Gamma_2} \\ -\mathbf{A}_{\Gamma_2\Gamma_1} \mathbf{u}_{\Gamma_1} \\ \mathbf{0} \end{pmatrix} \quad (8)$$

satisfies

$$\mathbf{U} = \begin{pmatrix} \mathbf{0} \\ \mathbf{0} \\ \mathbf{u}_{\Gamma_2} \\ \mathbf{u}_{\Omega_2} \end{pmatrix}.$$

Consequently, the discrete counterpart of the polarized wavefield \mathbf{U} is

$$\mathbf{U} = \begin{pmatrix} \mathbf{A}_{\Omega_1\Omega_1} & \mathbf{A}_{\Omega_1\Gamma_1} & \mathbf{0} & \mathbf{0} \\ \mathbf{A}_{\Gamma_1\Omega_1} & \mathbf{A}_{\Gamma_1\Gamma_1} & \mathbf{A}_{\Gamma_1\Gamma_2} & \mathbf{0} \\ \mathbf{0} & \mathbf{A}_{\Gamma_2\Gamma_1} & \mathbf{A}_{\Gamma_2\Gamma_2} & \mathbf{A}_{\Gamma_2\Omega_2} \\ \mathbf{0} & \mathbf{0} & \mathbf{A}_{\Omega_2\Gamma_2} & \mathbf{A}_{\Omega_2\Omega_2} \end{pmatrix}^{-1} \begin{pmatrix} \mathbf{0} \\ \mathbf{A}_{\Gamma_1\Gamma_2} \mathbf{u}_{\Gamma_2} \\ -\mathbf{A}_{\Gamma_2\Gamma_1} \mathbf{u}_{\Gamma_1} \\ \mathbf{0} \end{pmatrix}. \quad (9)$$

From the right-hand side of (8), one can easily see that knowledge of both \mathbf{u}_{Γ_1} and \mathbf{u}_{Γ_2} is required in order to compute \mathbf{U} . By construction, these sets contain information about the discrete wavefield and its normal derivative in the vicinity of Γ . Thus, as with the continuous case, the discrete case also requires information about the Dirichlet and Neumann traces to




Fig. 4. The domain decomposition, where the colored regions show examples of source distributions for each of the four scenarios.

compute the polarized wavefield \mathbf{U} in Ω_2 . In fact, [91] demonstrates, using similar techniques, that a discrete counterpart of the representation formula (4) can be derived.

Finally, our technique is easily extended to more general discretization techniques, as long as they allow for a reordering of the degrees-of-freedom to the block-tridiagonal system (7). Depending on the discretization, change in the selection of the sets Γ_1 and Γ_2 may be required. Similar schemes have been applied for many different discretizations such as high-order finite difference methods [95], finite element methods [90], enriched finite element methods [30], discontinuous Galerkin methods [79], and integral representations [93]. For example, for higher-order finite difference methods the stencils centered at the points in Ω_2 (respecting Γ_2, Γ_1 , or Ω_1) cannot involve discretization points in Γ_1 (respecting $\Omega_1, \Omega_2, \Gamma_2$). This is easily enforced by defining an appropriate δ , e.g., $\delta = 2$ for a 9-point, 5×5 , stencil in 2D. Similarly, in finite element or discontinuous Galerkin methods, the sparsity of the system matrices can be exploited in order to obtain suitable sets of degrees-of-freedom.

1.5. Organization

In Section 2, we introduce the algorithm to compute an approximate global solution u of the boundary value problem (2) with constant squared slowness m . We introduce the algorithm first on the continuous level in Section 2.1 and then extend it to the discrete level in Section 2.2. In Section 3, we show how the algorithm can be used to precondition the linear system (3). This opens the possibility of using the algorithm as part of an optimally parallel scaling solver based on a preconditioned GMRES method. We conclude Section 3 with a complexity analysis of this solver with regards to computational and communication effort. In Section 4, we analyze the effects of heterogeneous wave speeds on the effectiveness of the preconditioner. In Section 5, we provide several numerical examples in two- and three-dimensions for constant and non-constant wave speeds to corroborate all claims. The paper is concluded by a discussion of possible extensions of the method and a summary of the result. Additional proofs and implementation details are provided in the Appendices. Pseudo-code for the proposed algorithms is included as supplementary material, a C++ implementation of the 2D algorithm is available at [80], and experimental setups for the 2D examples are available at [81].

2. L-sweeps: reconstruction of wavefields

In this section, we introduce the algorithm to compute the global solution u of the boundary value problem (2). This solution is obtained from local solutions of local problems defined over a CDD. The concept of polarization introduced in Section 1 plays a crucial role in this procedure. We introduce the algorithm for constant wave speeds at the continuous and the discrete level in Sections 2.1 and 2.2. The procedure can be applied to problems with non-constant wave speeds in an analogous fashion. The effects of these heterogeneous wave speeds, in particular discontinuous ones, are addressed in Section 4.

2.1. Continuous formulation

Consider a decomposition of Ω into a CDD,³ with q rows and r columns, of non-overlapping open subdomains. For example, see Fig. 4, where $q = r = 5$. We first define the local problems associated with each subdomain, Ω_{ij} . To this end, we define extended domains $\Omega_{ij, \text{bulk}}^\varepsilon$ and $\Omega_{ij, \text{extended}}^\varepsilon$. The domain $\Omega_{ij, \text{bulk}}^\varepsilon$ is obtained from extending Ω_{ij} by a ε -layer along interior edges of the CDD, and $\Omega_{ij, \text{extended}}^\varepsilon$ is an extension of $\Omega_{ij, \text{bulk}}^\varepsilon$ to impose absorbing boundary conditions via PMLs.

³ Note that the CDD is chosen so that neglecting the PML regions (shown in gray in Fig. 4) each subdomain has the same size.

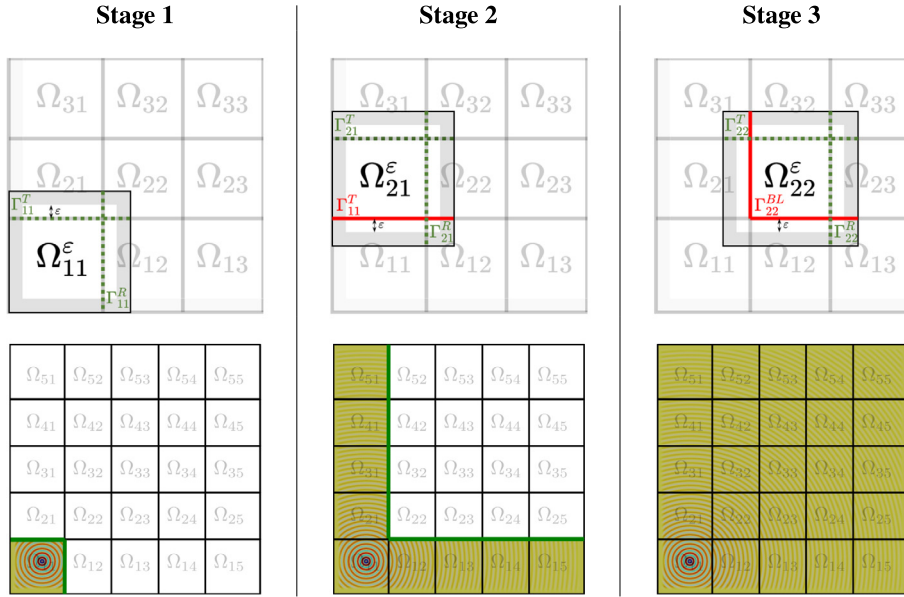


Fig. 5. Illustration of local problems (top row) and the computed solution after each stage of a sweep (bottom row). The green lines in the wavefields depict extracted trace information used in the next stage.

In $\Omega_{ij, \text{extended}}^\varepsilon$ we define the local squared slowness $m_{ij} := m|_{\Omega_{ij, \text{extended}}^\varepsilon}$. For a given source density f_{ij} in $\Omega_{ij, \text{extended}}^\varepsilon$, we define the local problem

$$\begin{aligned}
 -\operatorname{div}(\Lambda_{ij} \nabla u_{ij}) - \omega^2 m_{ij} u_{ij} &= f_{ij} \quad \text{in } \Omega_{ij, \text{extended}}^\varepsilon, \\
 u_{ij} &= 0 \quad \text{on } \partial \Omega_{ij, \text{extended}}^\varepsilon,
 \end{aligned}$$

where u_{ij} denotes the local solution. Due to the local PML, Λ_{ij} is a complex-valued diagonal matrix, and m_{ij} and f_{ij} are complex-valued functions in the PML region $\Omega_{ij, \text{extended}}^\varepsilon \setminus \Omega_{ij, \text{bulk}}^\varepsilon$. Further details of the PML formulation are provided in Appendix A. For the sake of brevity we denote the PML-adjusted squared slowness and source density still as m_{ij} and f_{ij} , respectively. Note that in contrast to (2), we do not set the squared slowness to be constant using a normal extension on $\partial \Omega_{ij, \text{bulk}}^\varepsilon$ in the PML region, rather we use the squared slowness inherited from the global problem in these regions. This definition makes our approach more accurate, since we use the local problem to compute sections of the global problem.

Note that the construction of the local problems consists of two subsequent extensions of Ω_{ij} : we first add an additional ε -layer around Ω_{ij} , and then further extend the subdomain by a PML region. While the latter is clearly needed to avoid artificial reflections in the local solutions, the former appears to be ad-hoc at this stage, but for reasons that will be explained in the sequel, the first extension by an ε -layer is crucial for a consistent exchange of information between subdomains. The resulting local problems are Dirichlet boundary-value problems defined on the extension $\Omega_{ij, \text{extended}}^\varepsilon$ of Ω_{ij} . As before, we denote the domain $\Omega_{ij, \text{extended}}^\varepsilon$ as Ω_{ij}^ε . Examples of these local problems are illustrated in Fig. 5.

In what follows, we introduce our method for computing global solutions to (2) by considering four scenarios, illustrated in Fig. 4, each increasingly more general:

1. the source density is supported in the interior of a corner subdomain of the CDD,
2. the source density is supported in the interior of an arbitrary subdomain,
3. the source density is supported in the interior of an arbitrary number of subdomains such that its support does not intersect the skeleton of the CDD, and
4. the source density has arbitrary support in Ω_{bulk} .

In the subsequent developments, we consider these scenarios using source distributions constructed from unions of point sources. However, the developments do not depend on an assumption that the source distributions are point-sources – any source density is admitted as long as it satisfies the conditions on its support. In particular, scenario 4 allows for arbitrary source distributions, including those which intersect the CDD skeleton, making our approach widely applicable.

2.1.1. Scenario 1: A source density supported in a corner subdomain

Without loss of generality, we introduce the algorithm for a source density supported in Ω_{11} . Source densities supported in any other corner subdomain can be constructed in an analogous way. The computation of the global solution is performed in three stages:

1. compute an approximation of the global solution in Ω_{11} using the local problem associated with Ω_{11} ,
2. compute the approximate global solution in the first row and column of the CDD, and,
3. compute the global solution in the rest of the subdomains.

Ultimately, using these three stages, we are able to compute the global solution, up to PML induced errors, using three sweeps: one vertical, one horizontal, and one diagonal.

Stage 1: Local solution

The local solution u_{11} defined over Ω_{11}^ε can be computed by solving the local problem associated with Ω_{11} with the source density $f_{11} := f|_{\Omega_{11}^\varepsilon}$. Since f is supported in Ω_{11} and the wave speed is constant, the global solution u restricted to Ω_{11} coincides with u_{11} up to errors induced by the PMLs. We therefore simply set the global solution u to u_{11} in Ω_{11} : $u|_{\Omega_{11}} := u_{11}|_{\Omega_{11}}$.

We define the straight lines containing the top and right boundary of Ω_{11} to be Γ_{11}^T and Γ_{11}^R , as depicted in Fig. 5. On Γ_{11}^T and Γ_{11}^R , we extract the Dirichlet and Neumann traces of u_{11} :

$$\begin{aligned} \lambda_{11}^T &:= u_{11}|_{\Gamma_{11}^T}, \mu_{11}^T := [n \cdot (\Lambda_{11} \nabla u_{11})]|_{\Gamma_{11}^T}, \\ \lambda_{11}^R &:= u_{11}|_{\Gamma_{11}^R}, \mu_{11}^R := [n \cdot (\Lambda_{11} \nabla u_{11})]|_{\Gamma_{11}^R}. \end{aligned} \tag{10}$$

These traces provide the necessary information to compute good approximations of the global solution in the neighboring subdomains Ω_{21} and Ω_{12} in stage 2. The solution after stage 1 is shown in Fig. 5 (bottom left), where the trace information that is used in stage 2 is shown in green.

Stage 2: Global solution in the same row/column

By construction, Γ_{11}^T lies entirely inside Ω_{21}^ε and the source density f is zero on Γ_{11}^T . Thus, the traces λ_{11}^T and μ_{11}^T can be used to compute a polarized wavefield in Ω_{21}^ε using $\Gamma_{21}^B := \Gamma_{11}^T$ as the polarization interface, and

$$u_{21}(x) = \int_{\Gamma_{21}^B} \mu_{11}^T(y) G_{21}(x, y) ds_y - \int_{\Gamma_{21}^B} \lambda_{11}^T(y) [n(y) \cdot (\Lambda_{21} \nabla_y G_{21}(x, y))] ds_y. \tag{11}$$

Here, $G_{21}(x, y)$ is the Green's function corresponding to the local problem defined on Ω_{21}^ε .

Following the same reasoning used for u_{11} in stage 1, the global solution restricted to Ω_{21} coincides with u_{21} , up to errors induced by the PMLs, and we set $u|_{\Omega_{21}} := u_{21}|_{\Omega_{21}}$. We also extract the Dirichlet and Neumann traces λ_{21}^T and μ_{21}^T of u_{21} on Γ_{21}^T in the same way as in (10) and repeat the process to compute an approximation of the global solution in the entire first column of the CDD. In addition, for each of these polarized wavefields, we extract the Dirichlet and Neumann traces λ_{i1}^R and μ_{i1}^R on Γ_{i1}^R . These traces are needed in stage 3 to compute approximations of the global solution in the rest of the subdomains. Similarly, we compute the global solution in the first row of the CDD and extract the traces λ_{1j}^T and μ_{1j}^T on Γ_{1j}^T for further use in stage 3. The solution after stage 2 and the trace information extracted for stage 3 are shown in Fig. 5 (bottom middle).

For the extension of the solution into the neighboring subdomains to be accurate, the local problems associated with Ω_{i1} and $\Omega_{(i+1)1}$ need to coincide in an ε -tube around $\Gamma_{i1}^T = \Gamma_{(i+1)1}^B$. This is ensured by the extension of the local problem by the additional ε -layer.

Stage 3: Global solution in the remaining subdomains

Consider Ω_{22} , where, by construction, $\Gamma_{22}^B := \Gamma_{12}^T$ and $\Gamma_{22}^L := \Gamma_{21}^R$ lie entirely inside Ω_{22}^ε . We combine these lines to form an L-shaped line, Γ_{22}^{BL} , so that Ω_{22} is entirely contained in the quadrant defined by Γ_{22}^{BL} as shown in Fig. 5. In addition, we combine the trace information on Γ_{21}^R and Γ_{12}^T to define Dirichlet and Neumann traces on Γ_{22}^{BL} as

$$\lambda_{22}^{BL}(x) := \begin{cases} \lambda_{21}^R(x) & x \in \Gamma_{22}^{BL} \cap \Gamma_{22}^L, \\ \lambda_{12}^T(x) & x \in \Gamma_{22}^{BL} \cap \Gamma_{22}^B, \end{cases} \tag{12}$$

$$\mu_{22}^{BL}(x) := \begin{cases} \mu_{21}^R(x) & x \in \Gamma_{22}^{BL} \cap \Gamma_{22}^L, \\ \mu_{12}^T(x) & x \in \Gamma_{22}^{BL} \cap \Gamma_{22}^B. \end{cases} \tag{13}$$

The traces can be used to compute a polarized wavefield given by

$$u_{22}(x) = \int_{\Gamma_{22}^{BL}} \mu_{22}^{BL}(y) G_{22}(x, y) ds_y - \int_{\Gamma_{22}^{BL}} \lambda_{22}^{BL}(y) [n \cdot (\Lambda \nabla_y G_{22}(x, y))] ds_y, \tag{14}$$

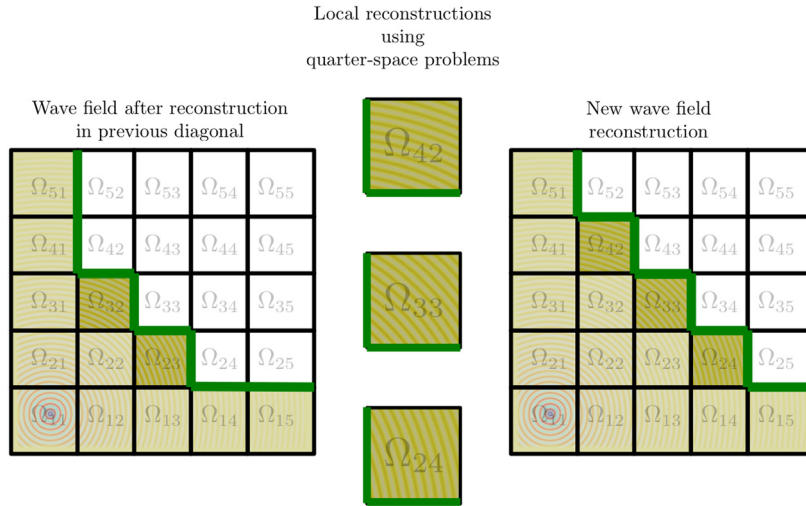


Fig. 6. Summary of the computed solution in one diagonal sweeping step. Updated subdomains are diagonally adjacent in a direction perpendicular to the bottom-left to top-right sweep direction.

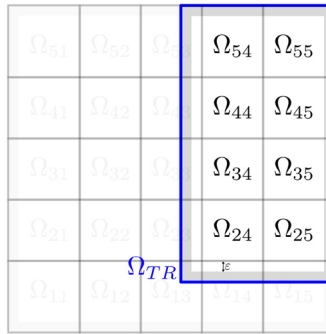


Fig. 7. The definition of the sub-problem by truncation with absorbing boundary conditions for scenario 2.

where $G_{22}(x, y)$ is the Green's function corresponding to the local problem defined on Ω_{22}^ε . In particular, (14) holds due to the fact that the local problems defined on Ω_{12}^ε and Ω_{21}^ε coincide with the local problem in Ω_{22}^ε in the vicinity of Γ_{22}^B and Γ_{22}^L , respectively.

It is clear that the polarized wavefield u_{22} coincides with the global wavefield u in Ω_{22} up to errors induced by the PMLs, and we simply set $u|_{\Omega_{22}} := u_{22}|_{\Omega_{22}}$. Similarly to (10), we extract the Dirichlet and Neumann traces of u_{22} on Γ_{22}^T and Γ_{22}^R and we use them to compute the approximations of the global wavefield in the other subdomains. Following this pattern, the wavefield is computed in the remaining subdomains by sweeping diagonally from the bottom-left corner to the top-right corner. At each step of the sweep, a diagonal perpendicular to the direction of the sweep is updated. This diagonal consists of subdomains only touching each other at a corner. For one diagonal, this procedure is summarized in Fig. 6.

2.1.2. Scenario 2: A source density supported in an arbitrary subdomain

In this scenario, we consider a point source supported in an arbitrary subdomain. Without loss of generality, for illustrative purposes we select the subdomain Ω_{24} . Source densities supported in any other subdomain can be treated analogously. First we restrict the problem to the quadrant in the top right corner of the domain, inclusive of Ω_{24} , which we denote Ω_{TR} and illustrate in Fig. 7, by truncating the global problem with PMLs. The interior boundary of Ω_{TR} is extended by a layer of thickness ε so that the local problems in Ω_{TR} are precisely the same as in Section 2.1.1. Similarly, we define domains Ω_{TL} , Ω_{BL} , and Ω_{BR} in the top-left, bottom-left, and bottom-right corner. These definitions reduce the problem posed on Ω to four sub-problems, each of which have a point source supported in a corner subdomain. Thus, we can readily apply the procedure introduced in Section 2.1.1 to each of them. The computed solutions are shown in Fig. 8.

The global solution may be obtained from summing these four solutions, taking care to avoid counting the contributions from the overlapped region multiple times. Executing this procedure as described requires many redundant computations. We achieve an equivalent method by generalizing the algorithm presented in Section 2.1.1, which we describe in three stages:




Fig. 8. The solutions in each quadrant in scenario 2.




Fig. 9. The computed wavefield and traces (green) after stage 2 of scenario 2.

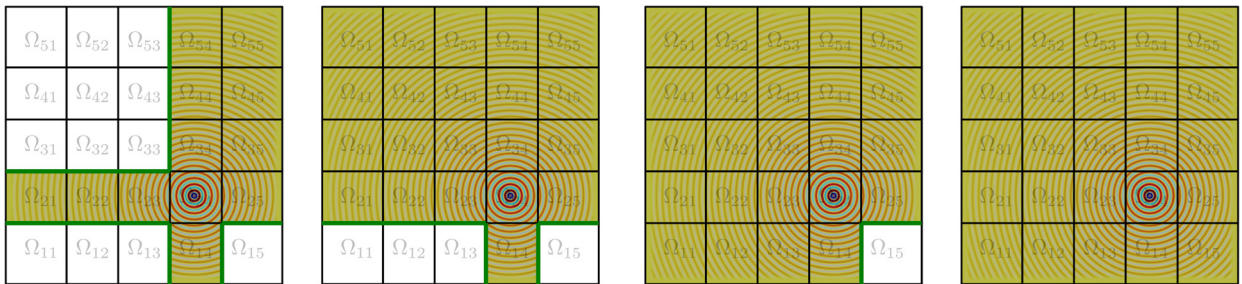


Fig. 10. The computed wavefield and traces (green) after each sweep in stage 3 of scenario 2.

1. compute the local solution in the subdomain where the source density is supported and extract the traces on all interior boundaries of the domain,
2. extend the solution into the subdomains in the same row/column following stage 2 of Section 2.1.1, and extract the traces required for stage 3, and,
3. extend the solution into the rest of the subdomains by computing the local solutions in diagonal rows, perpendicular to the direction of the sweep, similarly to stage 3 of Section 2.1.1.

The result of stage 2 is illustrated in Fig. 9 and stage 3 is illustrated in Fig. 10. In this scenario, the full algorithm requires a total of eight sweeps. Stage 2 requires four sweeps over the domain: up, down, left, right. Stage 3 also requires four sweeps: bottom-left to top-right, top-right to bottom-left, bottom-right to top-left, top-left to bottom-right.

2.1.3. Scenario 3: Arbitrary source distributions not intersecting the CDD skeleton

Global solutions for arbitrary source distributions that do not intersect the skeleton of the CDD can be computed from a union of distinct source densities, each supported in a single subdomain. Therefore, the global solution can be naively computed by applying the procedure from scenario 2 to each of the localized source densities and summing the results. This approach is not computationally efficient because much of the work is redundant.

In this section, we show how to compute the global solution without redundancy, by applying each of the eight sweeps of scenario 2 only once. This allows for an efficient computation of the global wavefield, regardless of the number of subdomains containing components of the source. The algorithm can still be performed in three stages, which are detailed below. For a succinct summary in pseudo-code we direct the reader to the supplementary material.

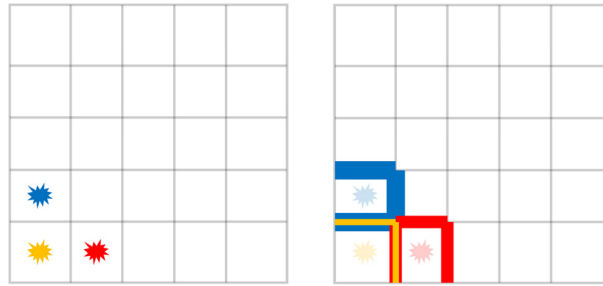


Fig. 11. Left: The setup of the problem in scenario 3, with multiple point sources shown in red, yellow and blue. Right: The extracted trace information is shown using the same colors as their corresponding point sources.

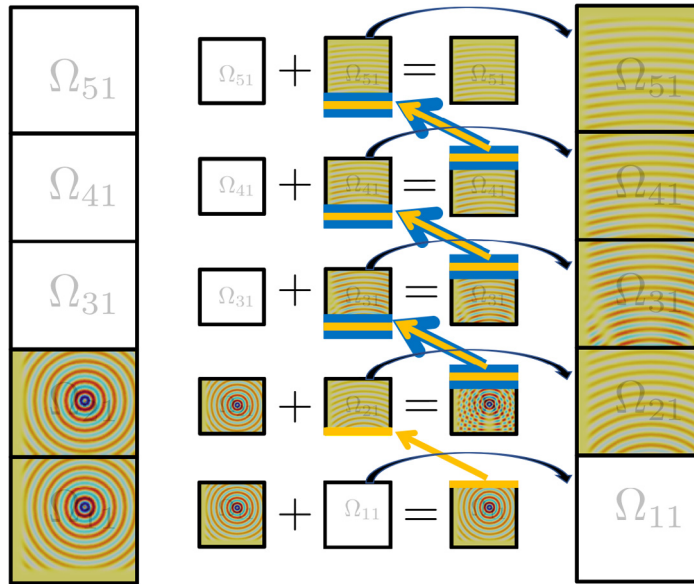


Fig. 12. Summary of the upwards sweep, for scenario 3, in one column in the presence of several point sources.

In the first stage, we restrict the source density to each subdomain. Within each subdomain, we use this restricted source density to compute the corresponding local solution, from which we extract the traces on all interior interfaces Γ_{ij}^l , $l = B, R, T, L$. This stage is illustrated in Fig. 11.

In the second stage, we use the traces extracted in stage 1 to extend the local solutions into the same column and row as subdomains containing sources. This can be done using four sweeps: up, down, left, and right. We consider the upwards sweep in the first column in detail, all other sweeps are performed analogously. In the upward sweep, we use traces on the bottom of each element to compute polarized wavefields and update local solutions. For example, in Fig. 11, the subdomain Ω_{11} has no incoming bottom trace. The outgoing top trace is therefore simply taken from stage one and transferred to Ω_{21} . In Ω_{21} , this trace is the incoming bottom trace and used to compute a polarized wavefield. We then update the local wavefield in Ω_{21} by adding the computed polarized wavefield. The outgoing trace on the top is then extracted from the updated local wavefield and used as the incoming trace in Ω_{31} . Continuing this procedure, we update the local wavefields in the entire first column which concludes the upwards sweep. The procedure is illustrated in Fig. 12. Then, we perform the remaining three cardinal sweeps.

Applying these four sweeps computes an intermediate wavefield. This wavefield is updated in stage 3. As part of the vertical and horizontal sweeps we also extract the vertical and horizontal traces of the subdomains needed in stage 3. The individual contributions from each of the four sweeps are shown in Fig. 13.

We generalize the third stage of the algorithm in a similar way by computing polarized wavefields from the trace information coming into a subdomain, and updating the local solutions by adding these polarized wavefields. We then extract trace information from the updated local wavefields and use them in the neighboring subdomains as incoming traces.

As in Section 2.1.2, stage 3 is realized by sweeping over the CDD from corner to corner, updating subdomains that are diagonally adjacent, perpendicular to the sweep direction. Computed wavefields in three example diagonals for the sweep from the bottom-left to the top-right corner are illustrated in Fig. 14. For the sweep from bottom-left to top-right we only




Fig. 13. The propagation of the trace information in each sweep of stage 2 of scenario 3. The stage 1 traces used in each sweep are shown in the light colors. The traces generated by stage 2 and used in stage 3 are shown in bold colors. The traces are color-coded with respect to the sources they originated from in Fig. 11.



Fig. 14. Illustration of the sweep from the bottom-left to the top-right corner in scenario 3. The colored lines show the trace information as it changes during the sweep. The traces are color-coded with respect to the sources they originated from in Fig. 11.

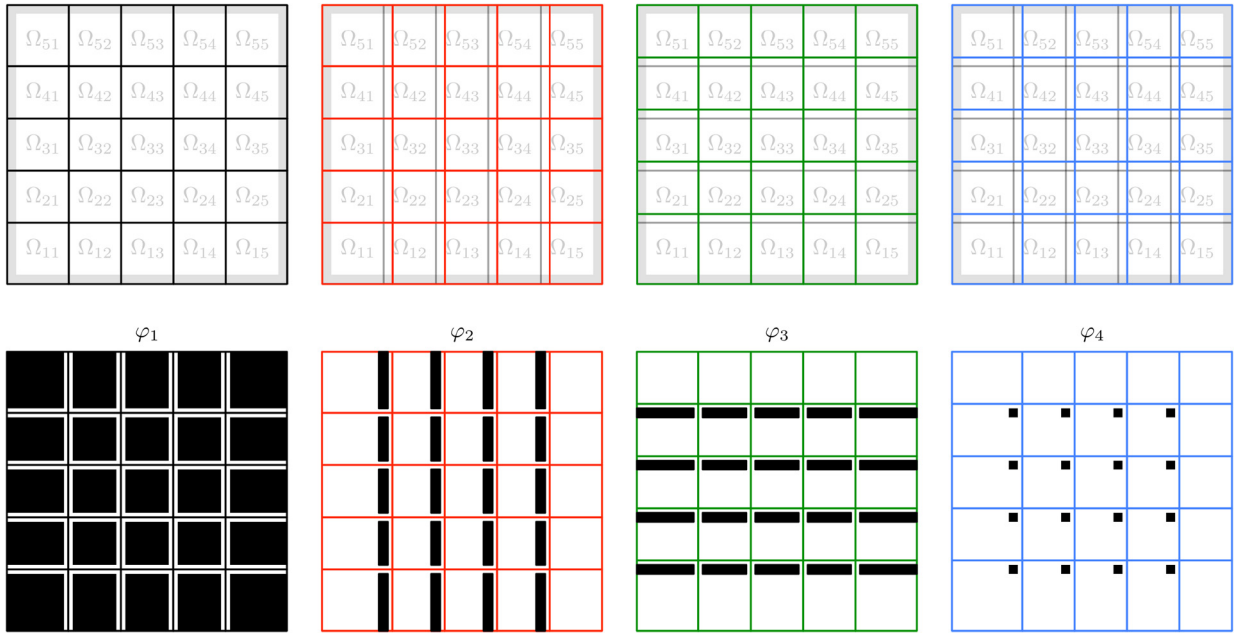


Fig. 15. The original CDD (black) and the three new, shifted CDDs in red, green, and blue (top row) and the corresponding window functions φ_i (bottom row).

use the top traces produced by the right-sweep of stage 2 and the right traces produced by the upwards-sweep of stage 2. The traces used in the other three diagonal sweeps are chosen in a similar way. The same technique can also be employed for the other three diagonal sweeps: top-right corner to bottom-left corner, bottom-right to top-left corner, and top-left to bottom-right corner.

Using these three stages, we compute the global solution in a total of eight sweeps, independently of the source distribution. In the next section, we show that applying this procedure to source densities appropriately windowed on four slightly different domain decompositions, the algorithm can be extended to entirely arbitrary source densities in a straightforward way.

2.1.4. Scenario 4: Arbitrary source distributions

The only remaining restriction on source distribution is that it must not cross the CDD skeleton. This restriction can be overcome by considering three more CDDs obtained from horizontal and vertical shifts in the CDD, as illustrated in the top row of Fig. 15. Using the original CDD and these new shifted CDDs, we define a partition of unity of four window functions φ_i , $i = 1, 2, 3, 4$, each of which vanish on the skeleton of one of the four CDD (see the bottom row of Fig. 15) and we define the corresponding windowed source densities $f^{\varphi_i} := \varphi_i f$. Clearly, f^{φ_i} vanishes on the skeleton of one CDD. Using this CDD, we can therefore apply the procedure of Section 2.1.3 to obtain the global solution u^{φ_i} corresponding to f^{φ_i} . The global solution for a general f is then simply the sum of those solutions:

$$u = \sum_{i=1}^4 u^{\varphi_i}.$$

2.2. Discrete formulation

In this section, we show that the procedure introduced in Section 2.1 can be applied on the discrete level to approximately solve (3). We start by defining the CDD and local problems. We define the CDD so that the skeleton does not intersect with any discretization point. Then, the skeleton clearly divides the global degrees-of-freedom into sets Ω_{ij} . We then define a local problem associated with each Ω_{ij} . This problem is defined on a superset of Ω_{ij} , which we denote Ω_{ij}^δ , and is the discretization of the local problem defined on Ω_{ij}^ε . The interior boundaries are first extended by a (2δ) -layer of degrees-of-freedom and then further extended by a PML region. In the same manner as for the continuous problem, the wave speed for this problem is inherited from the global wave speed. We denote the local system matrices \mathbf{A}_{ij} .

Furthermore, we define interfaces on these local problems to compute discrete polarized wavefields. For the continuous problem, these interfaces are defined as straight lines along the boundaries of the subdomains Ω_{ij} . For the discrete problem, the necessary interfaces for computing discrete polarized wavefields consist of all degrees-of-freedom δ -adjacent to the

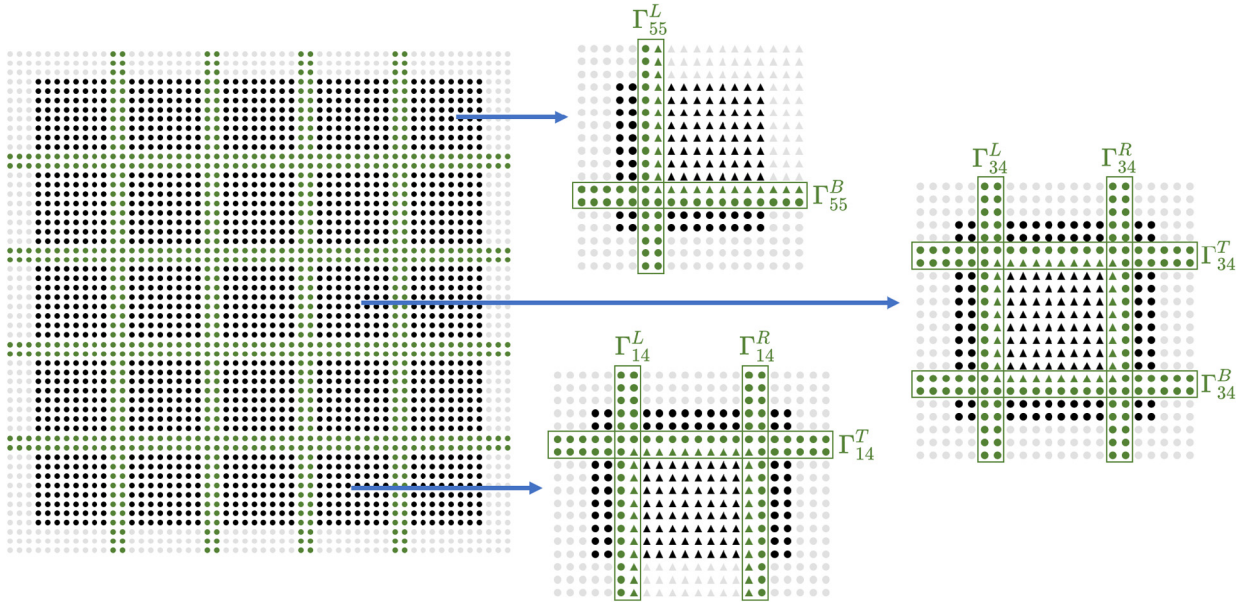


Fig. 16. Illustration of the discrete local problems and the discrete interfaces. Discrete interfaces are colored in green and degrees of freedom corresponding to the sets Ω_{ij} are denoted by triangles.

polarization interface. Therefore, we define the discrete counterparts of the interfaces Γ_{ij}^ℓ , $\ell = L, R, T, B$, to be the δ -adjacent degrees-of-freedom to Γ_{ij}^ℓ . We call those sets Γ_{ij}^ℓ , $\ell = B, R, T, L$. Fig. 16 illustrates these definitions. By construction, Γ_{ij}^ℓ defines the trace information that needs to be transferred between subdomains. In the continuous problem, two neighboring local problems have to coincide in an ε -tube around the polarization interface to accurately extend the wavefield from one subdomain into the other. In the same way, the discrete problem requires two neighboring local problems to coincide in all δ -adjacent degrees-of-freedom to the polarization interface. This justifies the (2δ) -layer between Ω_{ij} and the PML region. One δ -layer is used for Γ_{ij}^ℓ , and one δ -layer is used to ensure that two neighboring local problems coincide in an appropriately sized neighborhood of Γ_{ij}^ℓ .

In what follows we will show how the procedures introduced in Section 2.1 can be treated on the discrete level. Since the procedure introduced in Section 2.1.3 (scenario 3) is a generalization of all previous algorithms, we start with this procedure. In the first stage, we restrict the global source vector \mathbf{f} to the discretization points in Ω_{ij} for each subdomain and define a local source vector \mathbf{f}_{ij} on Ω_{ij}^δ such that $\mathbf{f}_{ij}|_{\Omega_{ij}} = \mathbf{f}|_{\Omega_{ij}}$ and \mathbf{f}_{ij} is zero everywhere else. In the same way as in the continuous case, this local source vector can then be used to compute discrete local solutions \mathbf{u}_{ij} in each subdomain:

$$\mathbf{u}_{ij} = \mathbf{A}_{ij}^{-1} \mathbf{f}_{ij}.$$

Using these local solutions, we define a global wavefield $\tilde{\mathbf{u}}$ such that $\tilde{\mathbf{u}}|_{\Omega_{ij}} = \mathbf{u}_{ij}|_{\Omega_{ij}}$. As with the continuous case, we extract the values of these local solutions on Γ_{ij}^ℓ and use them in stage 2 to update $\tilde{\mathbf{u}}$.

The second and third stage of the algorithm are realized by applying sweeps over the domain in the same way as in Section 2.1.3. The only difference is that the trace information on Γ_{ij}^ℓ is replaced by its discrete counterparts Γ_{ij}^ℓ , and the computation of continuous polarized wavefields is replaced by the computation of discrete polarized wavefields (9) using the local system matrices \mathbf{A}_{ij} and the appropriate interfaces Γ_{ij}^ℓ .

Until now, we have elided discussion of a subtle point in the discrete algorithm: construction of the L-shaped trace information. As an example, we consider the discrete counterpart of the trace Γ_{ij}^{BL} , but any other trace information can be handled analogously. Following the notation introduced in Section 2.1, let us denote the trace information on Γ_{ij}^B by λ_{ij}^B and the trace information on Γ_{ij}^L by λ_{ij}^L . Similarly to Γ_{ij}^{BL} in Section 2.1, Γ_{ij}^{BL} is defined as the L-shaped set of degrees-of-freedom surrounding the upper right quadrant of the set Ω_{ij} , as illustrated in Fig. 17. The trace information λ_{ij}^{BL} is then defined such that $\lambda_{ij}^{BL}|_{\Gamma_{ij}^B} = \lambda_{ij}^B$ and $\lambda_{ij}^{BL}|_{\Gamma_{ij}^L \setminus \Gamma_{ij}^B} = \lambda_{ij}^L|_{\Gamma_{ij}^L \setminus \Gamma_{ij}^B}$. Note that by convention we have defined the trace information of λ_{ij}^{BL} to coincide with λ_{ij}^B in $\Gamma_{ij}^B \cap \Gamma_{ij}^L$. This is an arbitrary choice, as the trace information on $\Gamma_{ij}^B \cap \Gamma_{ij}^L$ can be obtained from any of the two traces λ_{ij}^B or λ_{ij}^L .

Following this procedure, we can compute approximations $\tilde{\mathbf{u}}$ of the global solution that only differ from \mathbf{u} by errors induced by the discretization or PMLs, as long as the global source vector \mathbf{f} is zero on the degrees-of-freedom on the




Fig. 17. Illustration of the construction of the discrete L-shaped trace.




Fig. 18. Illustration of parallelism in diagonal sweeps. Subdomains processed in parallel have same label. Subdomains with like color are assigned to same processor.

skeleton of the CDD, i.e., the union of all Γ_{ij}^{ℓ} . This procedure can be extended to entirely arbitrary source vectors using window functions, as introduced in Section 2.1.4.

3. Implementation and complexity

The procedure introduced in Section 2 produces an approximation of a global solution to (3). This approach therefore defines the approximate solution operator $\tilde{\mathbf{A}}^{-1}$, such that $\tilde{\mathbf{A}}^{-1}(\mathbf{f}) \approx \mathbf{A}^{-1}\mathbf{f}$. Thus, we use $\tilde{\mathbf{A}}^{-1}$ to precondition (3),

$$\tilde{\mathbf{A}}^{-1}(\mathbf{A}\mathbf{u}) = \tilde{\mathbf{A}}^{-1}(\mathbf{f}), \tag{15}$$

and use a Krylov subspace method, such as GMRES [70] or BiCG-stab [85], to solve (15).

The resulting preconditioned iterative solver has the following properties:

- the preconditioner $\tilde{\mathbf{A}}^{-1}$ can be applied with optimal parallel complexity, $\mathcal{O}(N/p)$, and
- Krylov methods applied to (15) converge (empirically) in $\mathcal{O}(\log \omega)$ iterations;

thus making it parallel scalable. In this Section we focus on the first property. In particular, we show the implementation of the preconditioner, while analyzing its complexity with respect to the computational and communication cost. For the second property, we will present extensive numerical evidence in Section 5.

We consider a standard communication model [5,21,66,83]. The model assumes that each process is only able to send or receive a single message at a time, though different messages can be sent and received asynchronously. A message of size M can be communicated with $\alpha + \beta M$ complexity. The latency α represents the minimum complexity with which an arbitrary message from one process to another can be communicated, and is a constant overhead for any communication. The inverse bandwidth β is the complexity with which one unit of data can be communicated.

We implement the algorithm within a distributed memory framework using MPI. For simplicity, we assign each row of subdomains to one MPI rank.⁴ By assigning multiple subdomains to one rank, the preconditioner can be applied optimally in parallel, as described below. Consider a diagonal sweep, for example from the bottom-left corner to the top-right corner, for a $q \times q$ CDD. To maximize parallelism, each subdomain in a diagonal, perpendicular to the sweep direction, has to be processed in parallel. This can be realized by assigning subdomains to MPI ranks in a row-based fashion such that the i -th row of the CDD is processed by rank $i \bmod p$. Thus, each rank is assigned one or several rows of subdomains and does all of their associated computation. An illustration of the assignment of subdomains to rank is provided in Fig. 18, where subdomains assigned to the same processor are shown in the same color. In what follows, we assume this configuration,

⁴ This restriction can be relaxed to exploit asynchronous parallelism models and subdomain pipelining, for extremely large problems.

which allows us to also exploit the maximum parallelism in the application of the preconditioner and obtain a parallel scalable solver. The only restriction in this setup is that the number of ranks p used in this setup is bounded by the number of subdomains in one column, i.e., $p \leq q = \mathcal{O}(n)$.

For analysis purposes, we divide the solver for (3) in three phases:

1. *Setup* the local problems, system matrices, right-hand sides, and the corresponding window functions φ_i ;
2. *Factorize* the local system matrices; and
3. *Solve* the linear system using a Krylov method.

Phase 1: Setup

For two-dimensional problems, the size of the local problems associated with the subdomain Ω_{ij} is $\mathcal{O}(1)$. A crucial requirement to ensure that these local problems are indeed of $\mathcal{O}(1)$ size is that the thickness of the PML region, in wavelengths normal to Ω_{ij} , is held constant with mesh-refinement, i.e., the number of degrees-of-freedom in the PML does not change and the spatial PML-width scales as $O(1/\omega)$. Consequently, the local system matrices \mathbf{A}_{ij} , right-hand sides, and the window functions φ_{ij} can be assembled with $\mathcal{O}(1)$ computational complexity. In a parallel computational environment, each of the $\mathcal{O}(q^2)$ subdomains can be processed independently and no communication is required between subdomains. Thus all rows are processed with a $\mathcal{O}(n)$ parallel computational complexity, because $q = \mathcal{O}(n)$ and $p = \mathcal{O}(q)$. For two-dimensional problems, the computation in the setup phase is therefore realized with optimal $\mathcal{O}(N/p)$ parallel computational complexity. For three-dimensional problems, the size of the local problems associated with each subdomain is $\mathcal{O}(n)$, and the same arguments also yield optimal $\mathcal{O}(N/p)$ parallel complexity. In either case, there is no communication cost during the setup phase.

Phase 2: Factorize

For two-dimensional problems, the system matrices of the local problems have size $\mathcal{O}(1)$ and therefore can be factorized with $\mathcal{O}(1)$ computational complexity. As in stage 1, each subdomain can be processed independently with zero communication, thus we can follow the same argument as in the setup phase to justify that all of the q^2 subdomains can be factorized with $\mathcal{O}(N/p)$ parallel computational complexity. For three-dimensional problems, the system matrices of the local problems have size $\mathcal{O}(n)$ and are quasi-one-dimensional. Thus, standard sparse direct solvers can be used to factorize the matrices with optimal computational complexity [40,24,46], i.e., $\mathcal{O}(n)$. The same arguments as for the two-dimensional case can be employed to show that, in three dimensions, the q^2 subdomains can be factorized with $\mathcal{O}(N/p)$ parallel computational complexity. In either case, there is no communication cost during the factorization phase.

Phase 3: Solve

We use the GMRES method to solve the linear system (3). Each iteration of this method consists of three main operations: the application of \mathbf{A} , the computation between the Ritz vectors, and the application of the preconditioner $\tilde{\mathbf{A}}^{-1}$. For each of the three operations, we assume that the global vectors are provided in a distributed fashion such that each subdomain Ω_{ij} holds the values of the global vector corresponding to Ω_{ij} .

The application of \mathbf{A} can be realized by simply applying the local system matrices to a local vector, as long as each subdomain holds all degrees-of-freedom associated with Ω_{ij} and $\mathbf{\Gamma}_{ij}^\ell$ for $\ell = B, R, T, L$. This requires some communication because the subdomains only store the degrees-of-freedom in Ω_{ij} . Considering the row-based assignments of subdomains to MPI ranks, the sets $\mathbf{\Gamma}_{ij}^B$ have to be communicated to the subdomain $\Omega_{(i-1)j}$, and the sets $\mathbf{\Gamma}_{ij}^T$ have to be communicated to the subdomain $\Omega_{(i+1)j}$. In practice, for two-dimensional problems, each row has to communicate $\mathcal{O}(n)$ values to a neighboring MPI rank. Using the above communication model, this communication has complexity

$$\mathcal{O}(\alpha + \beta n) = \mathcal{O}(n).$$

Due to the row-based rank assignment, up to p of the q rows can simultaneously exchange information, resulting in an overall $\mathcal{O}(qn/p) = \mathcal{O}(N/p)$ parallel communication complexity. The same arguments can be applied for the communication of the sets $\mathbf{\Gamma}_{ij}^T$. Thus, for two dimensional problems, all necessary information can be communicated with optimal $\mathcal{O}(N/p)$ parallel complexity. In three dimensions, each row has to communicate $\mathcal{O}(n^2)$ values to a neighboring MPI rank. Following the same arguments as before all necessary information can be communicated with optimal $\mathcal{O}(N/p)$ parallel complexity. Once all information is communicated, the local system matrices \mathbf{A}_{ij} are applied to obtain the action of the global system matrix \mathbf{A} on the degrees-of-freedom in Ω_{ij} . This application can be realized with $\mathcal{O}(1)$ and $\mathcal{O}(n)$ computational complexity per subdomain in two and three dimensions, respectively. Again, each of the $\mathcal{O}(q^2)$ subdomains can be processed independently resulting in the total $\mathcal{O}(N/p)$ parallel computational and communication complexity to apply the matrix \mathbf{A} for two- and three-dimensional problems.

At iteration k , k inner products between the Ritz vectors need to be computed. Each inner product can be computed efficiently in parallel, given that each subdomain contains the local components of the Ritz vectors. Thus, each inner product is computed with $\mathcal{O}(N/p)$ parallel computational complexity and $\mathcal{O}(p)$ parallel communication complexity. The application of $\tilde{\mathbf{A}}^{-1}$ is analyzed in detail in Section 3.1, where we show that it can be applied with $\mathcal{O}(N/p)$ parallel computational

Table 1

Parallel complexity of the different stages for two-dimensional problems, where N is the total number of discretization points, n is the number of discretization points per dimension, ω is the frequency following $\omega \sim n$, and p is the number of processors following $p = \mathcal{O}(n)$.

Step	Computation		Communication	Total time
	per subdomain	per rank		
Set up	$\mathcal{O}(1)$	$\mathcal{O}(n) = \mathcal{O}(N/p)$	$\mathcal{O}(1)$	$\mathcal{O}(N/p)$
Factorization	$\mathcal{O}(1)$	$\mathcal{O}(n) = \mathcal{O}(N/p)$	$\mathcal{O}(1)$	$\mathcal{O}(N/p)$
Solve	$\mathcal{O}(\log \omega)$	$\mathcal{O}(n \log \omega) = \mathcal{O}((N/p) \log \omega)$	$\mathcal{O}((N/p) \log \omega)$	$\mathcal{O}((N/p) \log \omega)$

Table 2

Parallel complexity of the different stages for three-dimensional problems, where $N = n^3$ are the total number of discretization points, and the number of discretization points per dimension, ω is the frequency following $\omega \sim n$, and p is the number of processors following $p = \mathcal{O}(n)$.

Step	Computation		Communication	Total time
	per subdomain	per rank		
Set up	$\mathcal{O}(n)$	$\mathcal{O}(n^2) = \mathcal{O}(N/p)$	$\mathcal{O}(1)$	$\mathcal{O}(N/p)$
Factorization	$\mathcal{O}(n)$	$\mathcal{O}(n^2) = \mathcal{O}(N/p)$	$\mathcal{O}(1)$	$\mathcal{O}(N/p)$
Solve	$\mathcal{O}(n \log \omega)$	$\mathcal{O}(n^2 \log \omega) = \mathcal{O}((N/p) \log \omega)$	$\mathcal{O}((N/p) \log \omega)$	$\mathcal{O}((N/p) \log \omega)$

and communication complexity. The k -th GMRES iteration can therefore be realized with $\mathcal{O}(kN/p)$ parallel computational complexity and $\mathcal{O}(N/p)$ parallel communication complexity.⁵ Due to the effectiveness of the preconditioner, the number of iterations only grows as $\mathcal{O}(\log \omega)$. This claim is corroborated by the numerical examples in Section 5. Using restarts, the GMRES method can therefore be applied with a total parallel computational and communication complexity of $\mathcal{O}((N/p) \log \omega)$. The computation and communication complexities are summarized in Tables 1 and 2, for two- and three-dimensional problems.

3.1. Application of $\tilde{\mathbf{A}}^{-1}$

The application of $\tilde{\mathbf{A}}^{-1}$ can be divided in four steps:

- Computation of the local solutions,
- Extension of the local solutions into the same row,
- Extension of the local solutions into the same column, and
- Extension of the local solutions into the rest of the subdomains.

As before, we assume that at the beginning and at the end of each step, each rank holds all the values corresponding to the degrees-of-freedom in Ω_{ij} for all of its associated subdomains.

The computation of the local solutions requires to solve a local problem in each subdomain. Each of the local problems can be solved in $\mathcal{O}(1)$ and $\mathcal{O}(n)$ computational complexity for two- and three-dimensional problems, respectively.⁶ All local solutions are computed independently and p local solutions can be computed simultaneously. Therefore, for both two- and three-dimensional problems, step 1 is realized with $\mathcal{O}(N/p)$ parallel computational complexity and there is no communication cost.

The extension of the local solutions into the subdomains contained in the same row requires one left- and one right-sweep in each row. Due to the row-based rank assignment, each of those sweeps must be realized sequentially. However, each row can be processed independently, with no inter-row communication. For example, consider the right-sweep in one row. In each subdomain, a (discrete) polarized wavefield has to be computed (i.e., the local problem has to be solved), the local solution has to be updated, and the traces to be transferred to the right have to be extracted. For two-dimensional problems, each of those stages can be realized with $\mathcal{O}(1)$ computational complexity per subdomain because the size of the subdomains is $\mathcal{O}(1)$. Since each subdomain in the row has to be processed sequentially and there are $q = \mathcal{O}(n)$ subdomains in one row, this procedure can be realized with $\mathcal{O}(n)$ computational complexity per row. Employing this procedure concurrently in each row yields $\mathcal{O}(N/p)$ parallel computational complexity and no communication cost. For three-dimensional problems, each stage can be realized with $\mathcal{O}(n)$ computational complexity per subdomain.⁷ Because each of the $q = \mathcal{O}(n)$ subdomains in one row is processed sequentially, this results in a $\mathcal{O}(n^2)$ computational complexity in every row. Thus,

⁵ To achieve the parallel communication complexity, we implicitly assume that $\mathcal{O}(kp) = \mathcal{O}(N/p)$.

⁶ Similar to the factorization, standard sparse direct solvers are used to solve the quasi-one-dimensional problems arising from three-dimensional problems with optimal computational complexity, i.e., $\mathcal{O}(n)$.

⁷ Again, standard sparse direct solvers are used to solve the quasi-one-dimensional problems arising from three-dimensional problems with optimal complexity, i.e., $\mathcal{O}(n)$.

Table 3Parallel complexity of the different stages for two-dimensional problems for $p = \mathcal{O}(n)$.

Step	Computation		Communication	Total time
	per subdomain	per rank		
Part 1	$\mathcal{O}(1)$	$\mathcal{O}(n) = \mathcal{O}(N/p)$	$\mathcal{O}(1)$	$\mathcal{O}(N/p)$
Part 2	$\mathcal{O}(1)$	$\mathcal{O}(n) = \mathcal{O}(N/p)$	$\mathcal{O}(1)$	$\mathcal{O}(N/p)$
Part 3	$\mathcal{O}(1)$	$\mathcal{O}(n) = \mathcal{O}(N/p)$	$\mathcal{O}(N/p)$	$\mathcal{O}(N/p)$
Part 4	$\mathcal{O}(1)$	$\mathcal{O}(n) = \mathcal{O}(N/p)$	$\mathcal{O}(N/p)$	$\mathcal{O}(N/p)$

Table 4Parallel complexity of the different stages for three-dimensional problems when $p = \mathcal{O}(n)$.

Step	Computation		Communication	Total time
	per subdomain	per rank		
Part 1	$\mathcal{O}(n)$	$\mathcal{O}(n^2) = \mathcal{O}(N/p)$	$\mathcal{O}(1)$	$\mathcal{O}(N/p)$
Part 2	$\mathcal{O}(n)$	$\mathcal{O}(n^2) = \mathcal{O}(N/p)$	$\mathcal{O}(1)$	$\mathcal{O}(N/p)$
Part 3	$\mathcal{O}(n)$	$\mathcal{O}(n^2) = \mathcal{O}(N/p)$	$\mathcal{O}(N/p)$	$\mathcal{O}(N/p)$
Part 4	$\mathcal{O}(n)$	$\mathcal{O}(n^2) = \mathcal{O}(N/p)$	$\mathcal{O}(N/p)$	$\mathcal{O}(N/p)$

processing each row concurrently results in the $\mathcal{O}(N/p)$ parallel computational complexity and no communication cost. Analysis of the leftward sweep follows the same logic and reaches the same conclusions.

Extending the local solutions from one subdomain into the other subdomains within the same column requires one upward- and one downward-sweep. Given that the processor assignment does not allow for each column to be processed independently, the organization of the computation is non-trivial in order to reveal parallelism. As an example, we explain the implementation of an upward sweep. We begin by considering the subdomain in the bottom-left corner, Ω_{11} . Once Ω_{11} is processed and its information sent to Ω_{21} , rank 0 can process the subdomain Ω_{12} and rank 1 can process subdomain Ω_{21} . Following this procedure, the upward sweep is applied by sweeping over the domain, essentially pipelining the rows. Parallelism is achieved by processing diagonally-adjacent sets of subdomains simultaneously. By construction, each diagonal set has at most q subdomains. For two-dimensional problems, each subdomain can be processed with $\mathcal{O}(1)$ computational complexity, independent from all other subdomains. Thus, an entire diagonal set can be processed with $\mathcal{O}(q/p)$ parallel computational complexity. For three-dimensional problems, each subdomain can be processed with $\mathcal{O}(n)$ computational complexity, resulting in a $\mathcal{O}(nq/p)$ parallel computational complexity to process each diagonal. Since there are $\mathcal{O}(q)$ diagonals, and $q = \mathcal{O}(n)$, this results in the total $\mathcal{O}(N/p)$ parallel computational complexity to apply the upward sweep for both two- and three-dimensional problems.

In contrast with previous phases and steps, the upward and downward sweeps also require communication. For each subdomain Ω_{ij} , the traces $\mathbf{\Gamma}_{ij}^B$ have to be communicated from the subdomain $\Omega_{(i-1)j}$. For two-dimensional problems, the information transfers are performed in parallel, and all $\mathcal{O}(n)$ values, from all $\mathbf{\Gamma}_{ij}^B$ s, are communicated with $\mathcal{O}(n/p)$ parallel communication complexity. For three-dimensional problems, the communication volume is $\mathcal{O}(n^2)$ and, using the same arguments, the parallel communication complexity, per diagonal set, is $\mathcal{O}(n^2/p)$. For both two- and three-dimensional problems, considering all $\mathcal{O}(q) = \mathcal{O}(n)$ diagonals, we find the optimal parallel $\mathcal{O}(N/p)$ communication complexity. The analysis of the downward sweep follows the same argument and achieves the same conclusion.

The extension of the local solutions into the remaining subdomains, those not in the same row or column, requires sweeps over the CDD from corner to corner, along both diagonals. In these diagonal sweeps, each step requires processing of sets of subdomains that are diagonally-adjacent along the direction perpendicular to the sweep direction as illustrated in Fig. 18. This process has the same computation and communication patterns as the upward and downward sweeps and thus, following the same analysis, have the same complexities: they are applied with $\mathcal{O}(N/p)$ parallel computational complexity and $\mathcal{O}(N/p)$ parallel communication complexity.

In summary, each step has $\mathcal{O}(N/p)$ parallel computational complexity and at most $\mathcal{O}(N/p)$ parallel communication complexity, thus the preconditioner can be applied with $\mathcal{O}(N/p)$ parallel computational and communication complexity. The computational complexity and communication complexities, for applying the preconditioner, are summarized in Tables 3 and 4.

4. Heterogeneities

As developed in Section 2, the wavefield computed by the proposed algorithm can be interpreted as a sum of global solutions, each of which is induced by a global source density supported in a single subdomain only. We denote a source density supported in the subdomain Ω_{ij} by \mathbf{f}_{ij} and the global wavefield induced by this source density as $\mathbf{u}^{f_{ij}}$. Accordingly, the global wavefield \mathbf{u} can be written as




Fig. 19. Illustration of CDDs with associated wave speeds and a point source. The point source is shown by the red star. The wave speed is $m = 1$ in the light grey regions and $m = \frac{1}{2}$ in the dark grey regions.

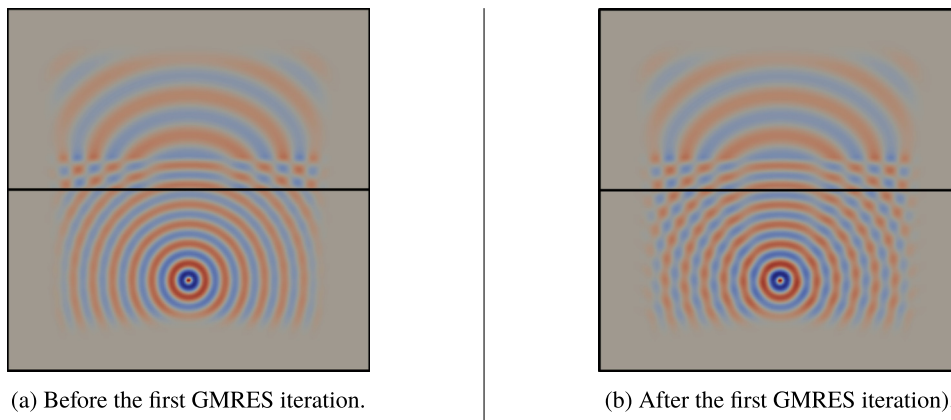


Fig. 20. Illustration of the computed wavefields.

$$\mathbf{u} = \sum_{i=1}^q \sum_{j=1}^r \mathbf{u}^{f_{ij}}.$$

To understand the effect of heterogeneities on the effectiveness of the preconditioner, it is sufficient to consider a wavefield $\mathbf{u}^{f_{ij}}$ due to a single point source in detail. We consider two examples. The first example reveals the effect of reflections on the wavefield constructed by the preconditioner. The second example shows the effects of reflections in more pathological media. Example 1 is a simple domain decomposition with two subdomains with a layered wave speed distribution of two layers where the interface of the wave speed is parallel to (but not on) the interface of the two subdomains. Example 2 is a discontinuous wave speed distribution in a checkerboard pattern aligned with the CDD. The setup of both examples is illustrated in Fig. 19.

In Example 1, excluding local partial wavefield computation, only the upward sweep computes a non-zero partial wavefield. The solution obtained from this sweep is shown in Fig. 20. The wavefield in the bottom subdomain is a solution to a problem with constant wave speed, and the reflection due to the discontinuity in the wave speed is clearly visible in the field in the top subdomain. The solutions in each subdomain can be explained in the following way: the local wavefield for Ω_{11} from stage 1 is computed by solving the local problem in Ω_{11} for the point source. Therefore the discontinuity in Ω_{21} is not visible and the wavefield for the constant wave speed is computed. This results in a poor approximation of the global solution in Ω_{11} . Nevertheless, the top traces λ_{11}^T are extracted. In the upward sweep (as part of stage 2), these traces are used in Ω_{21} to compute a polarized wavefield. This local wavefield is a good approximation of the global solution in Ω_{21} for two reasons:

- The discontinuity in the wave speed is visible to the local problem in Ω_{21} , and
- The trace λ_{11}^T contains all required information from Ω_{11} so that the polarized wavefield computed in Ω_{21} is a good approximation of the global wavefield.

The resulting wavefield is therefore a good approximation of the global wavefield only in Ω_{21} . As shown in Fig. 20a, there is a mismatch between the two local solutions, which produces a large residual concentrated at the interface.

Following the reasoning in [76,16], this residual contains the information necessary to propagate the wavefield locally between subdomains. The solutions obtained from applying the preconditioner to the right-hand side, and the solution after the first GMRES iteration are shown in Fig. 20. It is clear that after the second iteration, the GMRES method has converged to a good approximation of the wavefield.

This example shows that the preconditioner does not construct a good approximation of the global solution in rough media. This is because the preconditioner tracks the physical behavior of some waves propagating through the domain, including waves propagating in a straight line, intra-subdomain reflections, and some refracted waves, but does not account for any inter-subdomain reflections. Nevertheless, since the preconditioner is used as part of an iterative solver for the linear system, the inter-subdomain reflections and unresolved refractions will be treated in subsequent iterations.

To further illustrate this point, we consider Example 2, which is constructed to reveal the effect of inter-subdomain reflections. The wave speed distribution is chosen so that the effect of reflected wavefields have to be propagated through several subdomains. It is therefore expected that several iterations are required to resolve all inter-subdomain reflections and compute an accurate global wavefield. Indeed, in our numerical example we have observed that it takes 27 GMRES iterations to solve this problem to an accuracy of 10^{-6} .

These examples show the sensitivity of the overall solver to the number of reflections induced by the medium. Nevertheless, once the dominant reflections are resolved by the preconditioner, as we have observed in our numerical experiments, the number of iterations only grows as $\mathcal{O}(\log \omega)$ as the frequency (and the degrees-of-freedom) grows, see Section 5.4 and Section 5.5. The proposed solver is therefore scalable, however, the proportionality constant is strongly dependent on the number of reflections induced by the heterogeneous wave speed. While, in theory this can break the preconditioner, the numerical examples in Section 5 show that even in the case of wave-guides where reflections play a crucial role, the solution strategy surprisingly still only results in iteration counts lower than 40. Of course, the same arguments can be applied to refracted waves in the same way and similar effects can be seen.

5. Numerical examples

In this section, we consider several numerical examples to corroborate the claims of this paper. All numerical examples are constructed using variations of a standard setup. The problems are posed on the unit square ($d = 2$) or the unit cube ($d = 3$) and the wave speed is scaled such that the squared slowness, $m(x) \in [m_0, 1] \forall x \in \Omega$ for some $m_0 > 0$. For the characteristic frequency ω , the minimum wavelength is $2\pi/\omega$ and the maximum wavelength is $2\pi/(\sqrt{m_0}\omega)$. All problems are discretized using a uniform second-order finite difference approximation with N total discretization points, i.e., there are $n = N^{\frac{1}{d}}$ points in one direction. The characteristic frequency, ω , is chosen such that there are at least 10 points-per-wavelength, i.e., $\omega = 2\pi n/10$.

In each case, we decompose Ω using a CDD with $q = r$ rows and columns. Each subdomain Ω_{ij} in the CDD is chosen so that, ignoring the discretization points in the PML region, the set Ω_{ij} has 101 discretization points in each direction. The skeleton of the CDD does not intersect with any discretization point, so $n = 101q$ and $N = (101q)^d$. Accordingly, the maximum frequency is $\omega = 2\pi(101q - 1)/10$, which is chosen so that there are 10 discretization points per wavelength. The PML region is chosen so that $\lceil 20/\sqrt{m_0} \rceil$ discretization points are in this region. This ensures that the PML region is at least 2 wavelengths thick. The source distribution of the standard setup consists of four point sources at

$$\begin{aligned} x_1 &= (0.125, 0.125), & x_2 &= (0.125, 0.875), \\ x_3 &= (0.875, 0.125), & x_4 &= (0.875, 0.875). \end{aligned}$$

These point sources are modeled by an approximation with an exponential function:

$$f(x) := \frac{n^2}{\pi} \sum_{i=1}^4 e^{-n^2|x_i-x|^2}.$$

In all examples, the global system is solved using a preconditioned GMRES method [70] with a tolerance of 10^{-6} , using the zero-vector as an initial guess. All local problems are solved using Pardiso 6.0 [49]. Source code, models, and experiment configurations for these examples are available online [80,81].

5.1. Effect of PML-induced and discretization errors

In this section we provide numerical evidence of the claim in Section 2 that, in a homogeneous medium, the accuracy of the preconditioner depends only on the accuracy of the discretization and the quality of the absorbing boundary conditions. We consider the standard setup for a constant squared slowness $m = 1$ and we control the accuracy of the absorbing boundary condition by increasing (or reducing) the PML thickness. Table 5 depicts the dependence of the preconditioner, measured in iterations, on the quality of the absorbing boundary condition, which is tuned by varying the number of

Table 5
Number of iterations necessary to solve a homogeneous problem using increasingly thick PMLs and different frequencies.

N (without PML)	$\omega/2\pi$	$q = r$	Wavelengths in PML region						
			1	2	3	4	5	6	7
202×202	20.1	2	3	2	1	1	1	1	1
404×404	40.3	4	4	3	1	1	1	1	1
808×808	80.7	8	5	3	2	1	1	1	1
1616×1616	161.5	16	6	2	2	1	1	1	1
3232×3232	323.1	32	7	3	2	2	1	1	1
6464×6464	646.3	64	8	4	3	3	2	1	1
12928×12928	1292.7	128	10	5	3	3	3	2	1

wavelengths inside the PML region between 1 and 7. If the accuracy of the absorbing boundary condition is reduced (fewer wavelengths), the preconditioner is less effective. If the accuracy of the absorbing boundary condition is increased (more wavelengths), the preconditioner is more effective. In fact, if the accuracy of the boundary condition is sufficiently high, the preconditioner is perfectly effective resulting in iteration counts independent of ω . It can also be seen that the PML thickness has to be increased with ω in order to achieve this perfect effectiveness. This example demonstrates that in the context of constant wave speeds, the accuracy of the preconditioner depends only on the accuracy absorbing boundary condition.

Nevertheless, if the thickness of the PML region is kept constant with the problem size, the number of required GMRES iterations to solve the global system only grows logarithmically with the wave number ω . It is surprising that this qualitative property holds independently of the thickness of the PML, even for very thin PML regions, for example one wavelength.

This shows that the thickness of the PML has a significant influence on the effectiveness of the preconditioner. On the other hand, the thickness of the PML region is also crucial for the scalability of the preconditioner. For example, if the PML region is chosen to contain n_{pml} discretization points with $n_{\text{pml}} \ll n$, the sizes of the local problems associated with each subdomain grows as $\mathcal{O}(n_{\text{pml}}^2)$ in 2D and as $\mathcal{O}((n + n_{\text{pml}})n_{\text{pml}}^2)$ in 3D. Consequently, the parallel factorization of all subdomains scales as $\mathcal{O}(n^2 n_{\text{pml}}^3 / p)$ in 2D and $\mathcal{O}(n^3 n_{\text{pml}}^6 / p)$ in 3D and the application of the L-sweep preconditioner scales in parallel as $\mathcal{O}(n^2 n_{\text{pml}}^2 \log n_{\text{pml}} / p)$ in 2D and $\mathcal{O}(n^3 n_{\text{pml}}^4 / p)$ in 3D. Therefore, it is crucial to keep the thickness of the PML region as thin as possible to preserve efficiency. This makes the trade-off between efficiency and effectiveness of the preconditioner apparent. For all subsequent numerical examples considered in this paper, we choose a PML thickness of 2 wavelengths, which is empirically sufficient to achieve satisfactory results. In particular, the logarithmic growth of the number of iterations with respect to ω with this choice of PML-thickness allows one to achieve the advocated parallel scaling of the solver.

We also performed the same experiment for finite difference discretizations with fewer than 10 discretization points per wavelength. For very low accuracy discretizations (e.g., four points per wavelengths and more than 300 wavelengths inside the domain), the preconditioner can fail, i.e., the number of iterations to solve the preconditioned linear system grows faster than $\mathcal{O}(\log \omega)$. This is because the preconditioner relies on physical properties of wave propagation which are not captured by such inaccurate discretizations.

Considering these experiments, we conjecture that if the thickness of the PML region is kept constant with the frequency ω and the discretization is sufficiently accurate, the number of iterations required to solve the problem using a GMRES method grows as $\mathcal{O}(\log \omega)$. In addition, if the transparent boundary conditions can be modeled perfectly, the iteration count can be reduced to $\mathcal{O}(1)$ for reasonably smooth media.

5.2. Complexity of the solver

The complexity of the solver is determined by two factors: (i) the time it takes for one iteration, and (ii) the number of iterations needed to solve the problem. In this section, we provide experimental support demonstrating that the run-time per iteration scales as claimed.

We consider the standard setup with constant wave speed, for $q = 2, 4, 8, \dots, 128$ and we measure the factorization time, the time spent in one GMRES iteration, and the total time of the solver, which are summarized in Table 6. All timings are recorded on the NERSC machine ‘Cori’. Cori is composed of 2388 Haswell compute nodes, each node containing two 16-core Intel Xeon Processor E5-2698 v3 running at 2.3 GHz and 128 GB of memory. The nodes communicate via a Cray Aries interconnect with Dragonfly topology. The number of MPI ranks are chosen so that each compute node is assigned one MPI rank and each MPI rank is assigned one row of subdomains in the CDD. The number of MPI ranks is therefore $p = q$.

Table 6 contains the results of this experiment, from which we observe the scalability of the preconditioner. The timings clearly reflect the $\mathcal{O}(N/p)$ scaling as claimed in the prequel. Furthermore, the number of iterations N_{it} grows as $\mathcal{O}(\log \omega)$ which further proves the effectiveness of the preconditioner. Combining both leads to the total time, T_{total} , for the solver which reflects the claimed $\mathcal{O}(N/p) \log \omega$ scaling, as illustrated in Fig. 21. The above result is independent of medium. Thus, in subsequent sections where we consider non-constant wave speeds, it is sufficient to report iteration counts, rather than empirical run-times, to corroborate the claimed complexities of the solver.

Table 6

Timings for the solver on a homogeneous wave speed. N is the number of degrees-of-freedom, ω is the frequency, T_{fact} is the time spent to factorize the local problems, N_{it} is the number of GMRES iterations needed to solve the problem, T_{it} is the average time spent per GMRES iteration, and T_{total} is the total time spent to solve the linear system. All timings are measured in seconds.

N	$\omega/2\pi$	p	T_{fact}	N_{it}	T_{it}	T_{total}
202×202	20.1	2	1.09	2	0.66	2.63
404×404	40.3	4	1.00	3	0.58	2.56
808×808	80.7	8	1.41	3	1.26	6.02
1616×1616	161.5	16	2.80	2	3.39	14.05
3232×3232	323.1	32	4.41	3	5.47	27.47
6464×6464	646.3	64	8.34	4	11.09	67.74
12928×12928	1292.7	128	15.66	5	22.39	160.88




Fig. 21. Timings for 2D experiment with constant wave speed. Solid lines are empirical complexities for one application of the preconditioner (blue) and for the over-all solve (orange). Dashed lines are with measured results for one application of the preconditioner (blue) and for the over-all solve (orange).




Fig. 22. Wave speed distributions for the study of refractions. Random smooth wave speed (left; ‘Random’), smooth BP model background (center; ‘BG1’), and ‘true’ BP model background velocity (right; ‘BG2’).

5.3. Smooth wave speeds

In this section, we demonstrate the impact of *refracted* waves, induced by smooth media, on the performance of the preconditioner and iterative solver. As established in Section 4, the effectiveness of the L-sweeps preconditioner strongly depends on the presence of reflections and refractions in the solution. To evaluate the impact of refracted waves, we consider the standard setup for a random smooth wavefield and two background media arising from seismic imaging. One background medium is obtained by smoothing the background of the BP model [8], the other is the true background velocity provided in [8]. We call the smoothed background medium BG1 and the true background medium BG2. All media are illustrated in Fig. 22.



Fig. 23. Wave speed distributions for the study of the BP model with salt body. BG1 with salt body (left), BG2 with salt body (center), and 'true' BP model (right).

Table 7
Iteration counts for the BP model.

N (without PML)	$\omega/2\pi$	$q = r$	Random	BG1	BG2	BG1 with salt	BG2 with salt	BP model
202×202	20.1	2	3	1	4	7	6	7
404×404	40.3	4	3	2	4	9	9	9
808×808	80.7	8	5	4	6	12	12	12
1616×1616	161.4	16	7	5	6	15	15	15
3232×3232	323.1	32	9	6	7	17	17	16
6464×6464	646.3	64	10	7	7	19	19	19
12928×12928	1292.7	128	12	8	8	21	21	20

Table 7 shows the iteration counts obtained for all three experiments and demonstrates the effectiveness of the preconditioner when solving problems involving a smooth wave speed. In particular, we observe that the number of iterations scales as $\mathcal{O}(\log \omega)$, albeit with higher constants compared to the constant-velocity case.

5.4. BP model

In this section, we consider the wave speed of a standard geophysical benchmark problem, a subset of the BP model [8]. It is well-known that this problem involves many reflections in the resulting wavefields. We therefore use it to study the effect of reflections on the L-sweeps preconditioner.

First, we establish that reflected waves arise primarily due to the salt body (the dark red region in Fig. 23). To demonstrate this, we superimpose the salt body on the two background media considered in Section 5.3, BG1 and BG2, incrementally building to examine the effect of the high contrasts due to the salt body on the performance of the solver. Then, for completeness, we also run the test on the true BP model, as obtained from [8], illustrated in the image on the right of Fig. 23.

We report the iteration counts for these three experiments in Table 7, along-side those from Section 5.3 for comparison. The impact of the salt body on the performance of the solver is clearly reflected in the significant increase in GMRES iterations. This is explained because reflections are introduced in the wavefield due to the high contrast between the salt and the background. We also observe that the exact BP model can be solved in almost the same iteration counts as BG1 and BG2 with the salt body superimposed. Thus, even for the exact BP model, the reflections induced by the salt body dominate the performance of the preconditioner. Finally, let us note that even though reflections in the wavefield result in higher iteration counts, with increasing frequency ω , the number of iterations always grows only as $\mathcal{O}(\log \omega)$ and they are never higher than 21 iterations, even for problems involving 1000 wavelengths inside the domain. This shows the effectiveness of the preconditioner, even in the presence of reflections as they appear in practically relevant problems. The wavefield computed for the BP model and $\omega/2\pi = 323.1$ is shown in Fig. 24.

5.5. Wave-guide

To conclude the studies of two-dimensional problems, we consider a wave-guide with a point source at the entrance, a problem that is designed to stress the capabilities of the preconditioner because solutions contain many reflections. The wave-guide is illustrated in Fig. 25. Other than the change in the source and velocity configuration, the experimental setup follows the standard setup.

To stress the preconditioner, we vary the contrast ratio between the background and wave-guide. Table 8 shows the results of the experiment. Once all reflections are resolved for each problem, the number of iterations grows logarithmically with the frequency ω . In fact, the growth is so slow that it is comparable to the problem involving constant wave speeds. In




Fig. 24. The computed wavefield for the BP model and $\omega/2\pi = 323.1$. Zoom regions (a), (b), and (c) highlight solution quality in areas of high-contrast in the wave speed (at the salt boundary).




Fig. 25. The wave speed distribution of the wave-guide. The point source at the entrance is shown by the green star.

Table 8
Iteration counts for wave-guide problem for different contrast ratios of the wave speed between the wave-guide and the background.

N (without PML)	$\omega/2\pi$	$q = r$	Contrast ratio					
			2	3	4	5	6	
202×202	20.1	2	18	24	24	25	26	
404×404	40.3	4	28	29	29	28	30	
808×808	80.7	8	30	32	34	33	33	
1616×1616	161.5	16	31	33	33	34	35	
3232×3232	323.1	32	32	34	36	36	37	
6464×6464	646.3	64	32	34	35	36	36	

addition, while the number of iterations is clearly higher for this problem, it is interesting to observe that all problems can still be solved in under 40 iterations. It is noteworthy that the number of iterations is apparently independent the contrast ratio. The wavefield computed for a contrast ratio of 6 and $\omega/2\pi = 323.1$ is shown in Fig. 26.

5.6. 3D example

As a final example, we consider a three-dimensional problem. The effectiveness of the preconditioner has been thoroughly established in the previous two-dimensional numerical examples. These results directly translate over to three-dimensional problems. Therefore, we only consider the scalability of the solver for a three-dimensional problem for a constant wave speed.

To this end, we consider the standard setup in a 3D setting. This means that instead of constantly sized two-dimensional problems, the local problems in each subdomain are beam-shaped three-dimensional problems. These beam shaped local problems are obtained from the local 2D problems by extending them in the third dimension such that the global problem




Fig. 26. The computed wavefield for a contrast ratio of 6 and $\omega/2\pi = 323.1$. Zoom regions (a), (b), and (c) highlight solution quality in areas of high-contrast in the wave speed.

Table 9

Timings for the solver on a 3 D homogeneous wave speed. N is the number of degrees-of-freedom, ω is the frequency, T_{fact} is the time spent to factorize the local problems, N_{it} is the number of GMRES iterations needed to achieve a relative residual of 10^{-7} , T_{it} is the average time spent per GMRES iteration, and T_{total} is the total time spent to solve the linear system. All timings are measured in seconds.

N (without PML)	$\omega/2\pi$	p	T_{fact}	N_{it}	T_{it}	T_{total}
$26 \times 26 \times 26$	4.17	2	.04	4	1.34	6.52
$52 \times 52 \times 52$	8.50	4	5.54	6	5.30	37.17
$78 \times 78 \times 78$	12.83	6	12.42	6	12.80	89.76
$104 \times 104 \times 104$	17.17	8	22.91	6	23.27	163.62
$130 \times 130 \times 130$	21.50	10	37.53	7	36.47	292.33
$156 \times 156 \times 156$	25.83	12	52.47	7	51.62	417.08
$182 \times 182 \times 182$	30.17	14	71.71	8	68.92	627.23
$208 \times 208 \times 208$	34.50	16	96.14	7	91.65	743.37
$234 \times 234 \times 234$	38.83	18	124.64	8	116.08	1050.31
$260 \times 260 \times 260$	43.17	20	211.87	7	177.21	1438.12
$312 \times 312 \times 312$	51.83	24	314.93	8	263.16	2457.40
$416 \times 416 \times 416$	69.17	32	418.36	9	377.60	3992.63

has n degrees of freedom in each direction where, as before, $q = r$ is the number of subdomains in one direction of the CDD. The only other differences in problem setup are that we consider a discretization of 6 points per wavelength, assume that each subdomain contains 2 wavelengths, and choose a PML thickness of only one wavelength. Not counting the PML region, this results in problems with $n = 13q$ degrees of freedom in one direction. All of these changes are for computational expedience.

The timings are measured in the same computational environment as in Section 5.2. As before, one row of subdomains is assigned to one node, i.e., $p = q$. Table 9 shows the resulting time T_{fact} needed for the factorization of all the local problems, the number of iterations N_{it} required in the GMRES method, the average time T_{it} required in each iteration, and the total time T_{total} required for the solver.

It can be clearly seen from Table 2 that T_{fact} and T_{it} both show the claimed $\mathcal{O}(N/p)$ scaling. In addition, as in the two-dimensional cases, the number of iterations N_{it} grows as $\mathcal{O}(\log \omega)$ resulting in the almost optimal $\mathcal{O}((N/p) \log \omega)$ scaling for the total time T_{total} needed for the solver as shown in Fig. 27. This corroborates the claimed complexity for three-dimensional problems.

6. Discussion

We have introduced a parallel scalable algorithm to solve the high-frequency Helmholtz equation. To keep length of the paper reasonable, we have omitted some details, including the analysis of the continuum version of the algorithm in

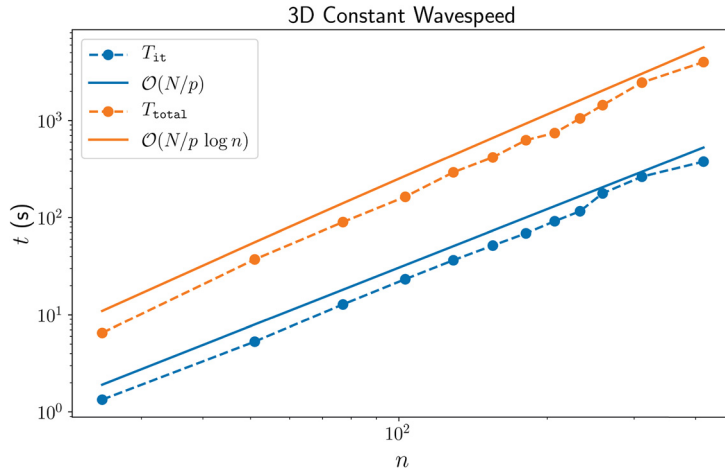


Fig. 27. Timings for the 3D experiment with constant wave speed. Solid lines are empirical complexities for one application of the preconditioner (blue) and for the over-all solve (orange). Dashed lines are with measured results for one application of the preconditioner (blue) and for the over-all solve (orange).

constant media, where the method essentially reduces to a direct solver, and the effect of pollution errors and associated refined discretization on the proposed algorithm. In this section we comment on both cases.

6.1. L-sweeps as a direct solver

In the case of constant wave speeds, it is clear that the medium does not induce any reflections. Thus, one application of the proposed preconditioner results in a global wavefield whose accuracy depends only on the accuracy of the absorbing boundary conditions. We have considered numerical examples that corroborate this fact: for constant wave speeds the proposed L-sweeps solution strategy converges in one iteration as long as the PML is sufficiently thick (see Table 5).

For a particular discretization of the problem, the PML accuracy can be carefully analyzed and chosen so that the error induced by the PML does not exceed the discretization error. As a result, one application of the preconditioner can be interpreted as a direct solver. Similar considerations have been established in [51], and extensions of their results to L-sweeps are straightforward. However, such analysis has to be carried out every time a new discretization is employed and it is nontrivial in the case of heterogeneous media.

6.2. L-sweeps in the presence of other discretizations

In the interest of a more focused presentation, we have considered low-order finite difference discretizations but an extension to other discretizations is straightforward. Such extensions have already been considered for the method of polarized traces (see e.g. [92,78]). In fact, since the proposed preconditioner only exploits physical properties of waves, it can be expected that better discretizations will lead to better performance of the preconditioner (i.e. lower iteration counts to solve the resulting linear system). For the method of polarized traces, this behavior can be seen by comparing the iteration counts for the BP 2004 model in [91] for finite difference methods with the ones in [78] for the hybridizable discontinuous Galerkin method. An extension of these results to L-sweeps is straightforward.

6.3. L-sweeps in the presence of pollution error

In this paper, we have kept the number of discretization points constant within one wavelength as we increased the frequency. It is well-known that this choice results in decreased accuracy for high-frequencies due to pollution error. This issue is inherited from the underlying discretization. We argue that the parallel scalability of the algorithm is not affected by pollution error.

One way to mitigate pollution errors is to consider higher-order discretizations. It has been shown that in heterogeneous media, pollution errors can be avoided if higher-order hybridizable discontinuous Galerkin methods are applied, provided that the local polynomial degree is chosen to increase logarithmically with ω [82,58]. This regime has already been considered in the context of the method of polarized traces where it has been shown that the combination of both techniques results in a pollution-free and sequentially scalable solver⁸ [82]. Like the method of polarized traces, L-sweeps are independent of the discretization. An extension of these results to L-sweeps resulting in a pollution-free and parallel scalable solver⁹ is straightforward.

⁸ The solver has a complexity of $\mathcal{O}(\omega^d)$ up to poly-logarithmic factors.

⁹ The solver can be expected to have a complexity of $\mathcal{O}(\omega^d/p)$ up to poly-logarithmic factors.

7. Summary

We have introduced the first solver for the high-frequency Helmholtz equation that scales as $\mathcal{O}((N/p)\log\omega)$ in a distributed-memory parallel computational environment, for a single right-hand side. The new solver constructs a preconditioner based on a CDD. This approach reveals parallelism which is exploited to obtain the optimal parallel complexity. The performance of the preconditioner is similar to the well-established method of polarized traces, i.e., the number of iterations grows as $\mathcal{O}(\log\omega)$.

We have shown that the preconditioner is effective once all reflections and refractions in the wavefield are resolved. While this is feasible for many practical applications, the performance of the proposed solver deteriorates in the presence of excessive reflected or refracted waves. This limitation arises, for example, on problems containing resonant cavities. We do not think that sweeping strategies are the correct approach to solve these problems, and due to the relevance of these problems, we view this case as a topic for further research.

Future work of the current method includes improvements for three-dimensional problems. As it is presented in this paper, the method results in a $\mathcal{O}(n^2) = \mathcal{O}(N^{\frac{2}{3}})$ parallel complexity for three-dimensional problems. However, all beam-shaped local problems are treated by one MPI rank only. It is well established that these quasi-one-dimensional problems can be treated in a distributed memory parallel computational environment by the use of multi-frontal methods [69,2,52]. In this case, further parallelism can be exploited and an almost $\mathcal{O}(n) = \mathcal{O}(N^{\frac{1}{3}})$ parallel complexity can be achieved for three-dimensional problems.

Using this extension, it can be shown that the high-frequency Helmholtz equation can be optimally parallelized in all spatial directions but one. The remaining direction is due to the inherently serial nature of the sweeps. In this regard, we conjecture that this work achieves the boundary of the scalability of sweeping preconditioners for a single right-hand side and that further improvements require new ideas. Nevertheless, in the presence of $\mathcal{O}(n)$ right-hand sides, sweeping preconditioners can be further parallelized by pipelining over the right-hand sides. In that regard, the parallel complexity can be further improved to an average $\mathcal{O}(1)$, up to logarithmic factors, parallel complexity per right-hand side. All of these extensions are currently under investigation.

Finally, while we only consider a row-based processor assignment in this work, in highly heterogeneous parallel computational environments other processor assignments, in particular processor assignments tailored to the heterogeneous computational environment, can be applied straightforwardly with our approach. This may lead to further improvements of the efficiency of the algorithm. An investigation of these aspects would also be very interesting.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgement

The authors thank Total SA for support and for permission to release the example code. LD is also supported by AFOSR grant FA9550-17-1-0316. The BP model is provided courtesy of BP and Frédéric Billette. We thank NERSC for computation resources. MT thanks the Institute of Applied Mathematics at Graz University of Technology for hosting him for part of this research.

Appendix A. Perfectly Matched Layers (PMLs) and finite difference discretizations

For a given squared slowness m and source density f defined on Ω_{extended} , following the technique of PMLs, a new complex-valued squared slowness and source density can be derived so that the PMLs act as absorbing boundary conditions. The coefficients in Ω_{extended} can be derived from the variable transformation [48,7]

$$\frac{\partial}{\partial x_k} \rightarrow \alpha_k(x) \frac{\partial}{\partial x_k}$$

for $k = 1, \dots, d$ where α_k is the complex-valued function

$$\alpha_k(x) = \frac{1}{1 + i \frac{\sigma_k(x)}{\omega}},$$

where $i = \sqrt{-1}$ is the imaginary unit and σ_k is the PML profile function.

To describe the PML profile, assume that Ω_{bulk} is the unit square or the unit cube. Then the extended domain Ω_{extended} is the square ($d = 2$) or cube ($d = 3$), $[-\delta_{\text{pml}}, 1 + \delta_{\text{pml}}]^d$, with the PML-width $\delta_{\text{pml}} > 0$. The PML profile σ_k is then chosen to be the cubic function

$$\sigma_k(x) := \begin{cases} \frac{C}{\delta_{\text{pml}}} \left(\frac{-x_k}{\delta_{\text{pml}}}\right)^3 & \text{for } x_k < 0 \\ 0 & \text{for } 0 \leq x_k \leq 1 \\ \frac{C}{\delta_{\text{pml}}} \left(\frac{x_k-1}{\delta_{\text{pml}}}\right)^3 & \text{for } x_k > 1 \end{cases}$$

where $C > 0$ is the absorption constant chosen to be $C = \ln \omega$. Employing this variable transformation in equation (2) gives rise to the diagonal matrix Λ with diagonal entries α_k^2/β , and the complex-valued squared slowness and source density can be written as

$$\frac{m}{\prod_{k=1}^d \alpha_k}, \quad \text{and} \quad \frac{f}{\prod_{k=1}^d \alpha_k},$$

respectively.

The resulting system is usually called the symmetric formulation of the Helmholtz equation, which in 2D takes the form:

$$-\left(\nabla \cdot \Lambda \nabla + \frac{\omega^2 m(\mathbf{x})}{\alpha_1(\mathbf{x})\alpha_2(\mathbf{x})}\right) u(\mathbf{x}) = \frac{f(\mathbf{x})}{\alpha_1(\mathbf{x})\alpha_2(\mathbf{x})}, \tag{16}$$

and

$$\Lambda(\mathbf{x}) = \begin{bmatrix} s_1(\mathbf{x}) & 0 \\ 0 & s_2(\mathbf{x}) \end{bmatrix}, \tag{17}$$

$s_1 = \alpha_1/\alpha_2$, $s_2 = \alpha_2/\alpha_1$, with homogeneous Dirichlet boundary conditions.

We discretize Ω_{bulk} as an equispaced regular grid of stepsize h , and of dimensions $n_x \times n_y$. For the extended domain Ω_{extended} , we extend this grid by $n_{\text{pml}} = \delta_{\text{pml}}/h$ points in each direction, obtaining a grid of size $(2n_{\text{pml}} + n_x) \times (2n_{\text{pml}} + n_y)$. Define $\mathbf{x}_{p,q} = (x_p, y_q) = (ph, qh)$. We use the 5-point stencil Laplacian to discretize (16). For the interior points $\mathbf{x}_{i,j} \in \Omega$, we have

$$(\mathbf{A}\mathbf{u})_{p,q} = -\frac{1}{h^2} (\mathbf{u}_{p-1,q} - 2\mathbf{u}_{p,q} + \mathbf{u}_{p+1,q}) - \frac{1}{h^2} (\mathbf{u}_{p,q-1} - 2\mathbf{u}_{p,q} + \mathbf{u}_{p,q+1}) - \omega^2 \frac{m(\mathbf{x}_{p,q})}{s_1(\mathbf{x}_{p,q})s_2(\mathbf{x}_{p,q})}. \tag{18}$$

In the PML, we discretize $\partial_x(s_1 \partial_x u)$ as

$$\frac{s_1(\mathbf{x}_{p+1/2,q})(\mathbf{u}_{p+1,q} - \mathbf{u}_{p,q}) - s_1(\mathbf{x}_{p-1/2,q})(\mathbf{u}_{p,q} - \mathbf{u}_{p-1,q})}{h^2}, \tag{19}$$

and analogously for $\partial_y(s_2 \partial_y u)$.

Appendix B. Proof of the annihilation condition

Let $x \in \Omega_1$, then using the definition of U and the fact that u and $G(x, \cdot)$ vanish on $\partial\Omega$, it holds

$$\begin{aligned} U(x) &= \int_{\Gamma} \Lambda(y) \nabla u(y) G(x, y) ds_y - \int_{\Gamma} u(y) [n(y) \cdot (\Lambda(y) \nabla_y G(x, y))] ds_y \\ &= \int_{\partial\Omega_2} \Lambda(y) \nabla u(y) G(x, y) ds_y - \int_{\partial\Omega_2} u(y) [n(y) \cdot (\Lambda(y) \nabla_y G(x, y))] ds_y. \end{aligned}$$

Then, employing the Divergence Theorem yields

$$\begin{aligned} U(x) &= \int_{\Omega_2} u(y) [-\text{div}_y (\Lambda(y) \nabla_y G(x, y))] dy - \int_{\Omega_2} [-\text{div} (\Lambda(y) \nabla u(y))] G(x, y) dy \\ &= \int_{\Omega_2} u(y) \underbrace{[-\text{div}_y (\Lambda(y) \nabla_y G(x, y)) - m(y)G(x, y)]}_{=0} dy - \int_{\Omega_2} \underbrace{[-\text{div} (\Lambda(y) \nabla u(y)) - m(y)u(y)]}_{=0} G(x, y) dy \end{aligned}$$

proving the Annihilation condition (6).

Appendix C. Supplementary material

Supplementary material related to this article can be found online at <https://doi.org/10.1016/j.jcp.2020.109706>.

References

- [1] P. Amestoy, R. Brossier, A. Buttari, J.-Y. L'Excellent, T. Mary, L. Métivier, A. Miniussi, S. Operto, Fast 3D frequency-domain full-waveform inversion with a parallel block low-rank multifrontal direct solver: application to OBC data from the North Sea, *Geophysics* 81 (2016) R363–R383.
- [2] P.R. Amestoy, I.S. Duff, J.-Y. L'Excellent, J. Koster, A fully asynchronous multifrontal solver using distributed dynamic scheduling, *SIAM J. Matrix Anal. Appl.* 23 (1) (2001) 15–41.
- [3] D. Aruliah, U. Ascher, Multigrid preconditioning for Krylov methods for time-harmonic Maxwell's equations in three dimensions, *SIAM J. Sci. Comput.* 24 (2) (2002) 702–718.
- [4] A.V. Astaneh, M.N. Guddati, A two-level domain decomposition method with accurate interface conditions for the Helmholtz problem, *Int. J. Numer. Methods Eng.* 107 (1) (2016) 74–90, nme.5164.
- [5] G. Ballard, J. Demmel, O. Holtz, O. Schwartz, Minimizing communication in numerical linear algebra, *SIAM J. Matrix Anal. Appl.* 32 (3) (2011) 866–901.
- [6] M. Bebendorf, *Hierarchical Matrices: A Means to Efficiently Solve Elliptic Boundary Value Problems*, Lecture Notes in Computational Science and Engineering (LNCSE), vol. 63, Springer-Verlag, ISBN 978-3-540-77146-3, 2008.
- [7] J.-P. Bérenger, A perfectly matched layer for the absorption of electromagnetic waves, *J. Comput. Phys.* 114 (2) (1994) 185–200.
- [8] F. Billelte, S. Brandsberg-Dahl, *The 2004 BP Velocity Benchmark*, EAGE, 2005.
- [9] S. Boerm, L. Grasedyck, W. Hackbusch, *Hierarchical Matrices*, Max-Planck-Institute Lecture Notes, 2006.
- [10] Y. Boubendir, An analysis of the BEM-FEM non-overlapping domain decomposition method for a scattering problem, *J. Comput. Appl. Math.* 204 (2) (2007) 282–291, Special Issue: The Seventh International Conference on Mathematical and Numerical Aspects of Waves (WAVES'05).
- [11] Y. Boubendir, X. Antoine, C. Geuzaine, A quasi-optimal non-overlapping domain decomposition algorithm for the Helmholtz equation, *J. Comput. Phys.* 231 (2) (2012) 262–280.
- [12] J. Bramble, J. Pasciak, Analysis of a finite PML approximation for the three dimensional time-harmonic Maxwell and acoustic scattering problems, *Math. Comput.* 76 (258) (2007) 597–614.
- [13] A. Brandt, I. Livshits, Wave-ray multigrid method for standing wave equations, *Electron. Trans. Numer. Anal.* 6 (1997) 162–181.
- [14] H. Calandra, S. Gratton, X. Pinel, X. Vasseur, An improved two-grid preconditioner for the solution of three-dimensional Helmholtz problems in heterogeneous media, *Numer. Linear Algebra Appl.* 20 (4) (2013) 663–688.
- [15] T.F. Chan, T.P. Mathew, Domain decomposition algorithms, *Acta Numer.* 3 (61–143) (1994) 1.
- [16] Z. Chen, X. Xiang, A source transfer domain decomposition method for Helmholtz equations in unbounded domain, *SIAM J. Numer. Anal.* 51 (4) (2013) 2331–2356.
- [17] F. Collino, S. Ghanemi, P. Joly, Domain decomposition method for harmonic wave propagation: a general presentation, *Comput. Methods Appl. Mech. Eng.* 184 (2–4) (2000) 171–211.
- [18] T.A. Davis, Algorithm 832: UMFPACK v4.3—an unsymmetric-pattern multifrontal method, *ACM Trans. Math. Softw.* 30 (2) (June 2004) 196–199.
- [19] M.V. de Hoop, S. Wang, J. Xia, On 3D modeling of seismic wave propagation via a structured parallel multifrontal direct Helmholtz solver, *Geophys. Prospect.* 59 (5) (2011) 857–873.
- [20] A. de La Bourdonnaye, C. Farhat, A. Macedo, F. Magoules, F.-X. Roux, A non-overlapping domain decomposition method for the exterior Helmholtz problem, *Contemp. Math.* 218 (1998) 42–66.
- [21] J. Demmel, L. Grigori, M. Gu, H. Xiang, Communication avoiding rank revealing qr factorization with column pivoting, Technical Report UCB/EECS-2013-46, EECS Department, University of California Berkeley, May 2013.
- [22] J.W. Demmel, S.C. Eisenstat, J.R. Gilbert, X.S. Li, J.W.H. Liu, A supernodal approach to sparse partial pivoting, *SIAM J. Matrix Anal. Appl.* 20 (3) (1999) 720–755.
- [23] B. Després, Décomposition de domaine et problème de Helmholtz, *C. R. Séances Acad. Sci., Sér. 1 Math.* 311 (1990) 313–316.
- [24] I.S. Duff, J.K. Reid, The multifrontal solution of indefinite sparse symmetric linear, *ACM Trans. Math. Softw.* 9 (3) (September 1983) 302–325.
- [25] B. Engquist, L. Ying, Sweeping preconditioner for the Helmholtz equation: hierarchical matrix representation, *Commun. Pure Appl. Math.* 64 (5) (2011) 697–735.
- [26] B. Engquist, L. Ying, Sweeping preconditioner for the Helmholtz equation: moving perfectly matched layers, *Multiscale Model. Simul.* 9 (2) (2011) 686–710.
- [27] B. Engquist, H.-K. Zhao, Absorbing boundary conditions for domain decomposition, *Appl. Numer. Math.* 27 (4) (1998) 341–365, Special Issue on Absorbing Boundary Conditions.
- [28] Y.A. Erlangga, C.W. Oosterlee, C. Vuik, A novel multigrid based preconditioner for heterogeneous Helmholtz problems, *SIAM J. Sci. Comput.* 27 (4) (2006) 1471–1492.
- [29] O.G. Ernst, M.J. Gander, Why it is difficult to solve Helmholtz problems with classical iterative methods, in: Ivan G. Graham, Thomas Y. Hou, Omar Lakkis, Robert Scheichl (Eds.), *Numerical Analysis of Multiscale Problems*, in: Lecture Notes in Computational Science and Engineering, vol. 83, Springer, Berlin, Heidelberg, 2012, pp. 325–363.
- [30] J. Fang, J. Qian, L. Zepeda-Núñez, H. Zhao, Learning dominant wave directions for plane wave methods for high-frequency Helmholtz equations, *Res. Math. Sci.* 4 (1) (May 2017) 9.
- [31] M. Gander, F. Nataf, AILU for Helmholtz problems: a new preconditioner based on an analytic factorization, *C. R. Acad. Sci., Sér. 1 Math.* 331 (3) (2000) 261–266.
- [32] M. Gander, H. Zhang, A class of iterative solvers for the Helmholtz equation: factorizations, sweeping preconditioners, source transfer, single layer potentials, polarized traces, and optimized Schwarz methods, *SIAM Rev.* 61 (1) (2019) 3–76.
- [33] M.J. Gander, Optimized Schwarz methods, *SIAM J. Numer. Anal.* 44 (2) (2006) 699–731.
- [34] M.J. Gander, F. Kwok, Optimal interface conditions for an arbitrary decomposition into subdomains, in: Yunqing Huang, Ralf Kornhuber, Olof Widlund, Jinchao Xu (Eds.), *Domain Decomposition Methods in Science and Engineering XIX*, in: Lecture Notes in Computational Science and Engineering, vol. 78, Springer, Berlin, Heidelberg, 2011, pp. 101–108.
- [35] M.J. Gander, F. Magoules, F. Nataf, Optimized Schwarz methods without overlap for the Helmholtz equation, *SIAM J. Sci. Comput.* 24 (1) (2002) 38–60.
- [36] M.J. Gander, Y. Xu, Optimized Schwarz method with two-sided transmission conditions in an unsymmetric domain decomposition, in: *Domain Decomposition Methods in Science and Engineering XXII*, Springer International Publishing, Cham, 2016, pp. 631–639.
- [37] M.J. Gander, H. Zhang, Domain decomposition methods for the Helmholtz equation: a numerical investigation, in: Randolph Bank, Michael Holst, Olof Widlund, Jinchao Xu (Eds.), *Domain Decomposition Methods in Science and Engineering XX*, in: Lecture Notes in Computational Science and Engineering, vol. 91, Springer, Berlin, Heidelberg, 2013, pp. 215–222.
- [38] M.J. Gander, H. Zhang, Optimized Schwarz methods with overlap for the Helmholtz equation, in: *Domain Decomposition Methods in Science and Engineering XXI*, Springer International Publishing, Cham, 2014, pp. 207–215.
- [39] M.J. Gander, I.G. Graham, E.A. Spence, Applying GMRES to the Helmholtz equation with shifted Laplacian preconditioning: what is the largest shift for which wavenumber-independent convergence is guaranteed?, *Numer. Math.* (2015) 1–48.
- [40] A. George, Nested dissection of a regular finite element mesh, *SIAM J. Numer. Anal.* 10 (1973) 345–363.
- [41] S. Ghanemi, A domain decomposition method for Helmholtz scattering problems, in: *Ninth International Conference on Domain Decomposition Methods*, 1998, pp. 105–112.

- [42] A. Gillman, A.H. Barnett, P.-G. Martinsson, A spectrally accurate direct solution technique for frequency-domain scattering problems with variable media, *BIT Numer. Math.* 55 (1) (Mar 2015) 141–170.
- [43] Dan Gordon, Rachel Gordon Carp-cg, A robust and efficient parallel solver for linear systems, applied to strongly convection dominated {PDEs}, *Parallel Comput.* 36 (9) (2010) 495–515.
- [44] R. Hiptmair, A. Moiola, I. Perugia, Plane wave discontinuous Galerkin methods for the 2D Helmholtz equation: analysis of the p-version, *SIAM J. Numer. Anal.* 49 (1) (2011) 264–284.
- [45] R. Hiptmair, A. Moiola, I. Perugia, A survey of Trefftz methods for the Helmholtz equation, in: G.R. Barrenea, F. Brezzi, A. Cangiani, E.H. Georgoulis (Eds.), *Building Bridges: Connections and Challenges in Modern Approaches to Numerical Partial Differential Equations*, Springer, 2016, pp. 237–278.
- [46] R.W. Hockney, A fast direct solution of Poisson's equation using Fourier analysis, *J. ACM* 12 (1) (January 1965) 95–113.
- [47] Q. Hu, H. Zhang, Substructuring preconditioners for the systems arising from plane wave discretization of Helmholtz equations, *SIAM J. Sci. Comput.* 38 (4) (2016) A2232–A2261.
- [48] S. Johnson, March 2010, Notes on perfectly matched layers (PMLs).
- [49] D. Kourounis, A. Fuchs, O. Schenk, Toward the next generation of multiperiod optimal power flow solvers, *IEEE Trans. Power Syst.* 33 (4) (July 2018) 4005–4014.
- [50] A. Laird, M. Giles, Preconditioned iterative solution of the 2D Helmholtz equation, Technical Report NA 02-12, Computing Lab, Oxford University, May 2002.
- [51] W. Leng, L. Ju, An additive overlapping domain decomposition method for the Helmholtz equation, *SIAM J. Sci. Comput.* 41 (2) (2019) A1252–A1277.
- [52] X.S. Li, J.W. Demmel, SuperLU DIST: a scalable distributed-memory sparse direct solver for unsymmetric linear systems, *ACM Trans. Math. Softw.* 29 (2) (June 2003) 110–140.
- [53] Y. Li, L. Métivier, R. Brossier, B. Han, J. Virieux, 2D and 3D frequency-domain elastic wave modeling in complex media with a parallel iterative solver, *Geophysics* 80 (2015) T101–T118.
- [54] P.-L. Lions, On the Schwarz alternating method II, in: Tony Chan, Roland Glowinski, Jacques Periaux, Olof Widlund (Eds.), *Domain Decomposition Methods*, in: *Lecture Notes in Computational Science and Engineering*, SIAM, 1989, pp. 47–70.
- [55] F. Liu, L. Ying, Recursive sweeping preconditioner for the 3D Helmholtz equation. ArXiv e-prints, 2015.
- [56] F. Magoules, K. Meerbergen, J.-P. Coyette, Application of a domain decomposition with Lagrange multipliers to acoustic problems arising from the automotive industry, *J. Comput. Acoust.* 08 (03) (2000) 503–521.
- [57] L.C. McInnes, R.F. Susan-Resiga, D.E. Keyes, H.M. Atassi, Additive Schwarz methods with nonreflecting boundary conditions for the parallel computation of Helmholtz problems, *Contemp. Math.* 218 (1998) 325–333.
- [58] J.M. Melenk, S. Sauter, Wavenumber explicit convergence analysis for Galerkin discretizations of the Helmholtz equation, *SIAM J. Numer. Anal.* 49 (3) (2011) 1210–1243.
- [59] A. Modave, X. Antoine, C. Geuzaine, an efficient domain decomposition method with cross-point treatment for Helmholtz problems, 2018.
- [60] A. Moiola, R. Hiptmair, I. Perugia, Plane wave approximation of homogeneous Helmholtz solutions, *Z. Angew. Math. Phys.* 62 (5) (2011) 809–837.
- [61] A. Moiola, E. Spence, Is the Helmholtz equation really sign-indefinite?, *SIAM Rev.* 56 (2) (2014) 274–312.
- [62] W.A. Mulder, A multigrid solver for 3D electromagnetic diffusion, *Geophys. Prospect.* 54 (5) (2006) 633–649.
- [63] Frédéric Nataf, Francis Nier, Convergence rate of some domain decomposition methods for overlapping and nonoverlapping subdomains, *Numer. Math.* 75 (3) (1997) 357–377.
- [64] R.-E. Plessix, A Helmholtz iterative solver for 3D seismic-imaging problems, *Geophysics* 72 (2007), SM185–SM194.
- [65] R.-E. Plessix, W.A. Mulder, Separation-of-variables as a preconditioner for an iterative Helmholtz solver, *Appl. Numer. Math.* 44 (3) (2003) 385–400.
- [66] J. Poulson, L. Demanet, N. Maxwell, L. Ying, A parallel butterfly algorithm, *SIAM J. Sci. Comput.* 36 (1) (2014) C49–C65.
- [67] J. Poulson, B. Engquist, S. Li, L. Ying, A parallel sweeping preconditioner for heterogeneous 3D Helmholtz equations, *SIAM J. Sci. Comput.* 35 (3) (2013) C194–C212.
- [68] R.G. Pratt, Seismic waveform inversion in the frequency domain; part 1: theory and verification in a physical scale model, *Geophysics* 64 (3) (1999) 888–901.
- [69] F.-H. Rouet, X.S. Li, P. Ghysels, A. Napov, A distributed-memory package for dense hierarchically semi-separable matrix computations using randomization, *ACM Trans. Math. Softw.* 42 (4) (June 2016) 27:1–27:35.
- [70] Y. Saad, M.H. Schultz GMRES, A generalized minimal residual algorithm for solving nonsymmetric linear systems, *SIAM J. Sci. Stat. Comput.* 7 (3) (July 1986) 856–869.
- [71] H.A. Schwarz, Über einen grenzübergang durch alternierendes verfahren, *Vierteljahrsschr. Nat.forsch. Ges. Zür.* 15 (1870) 272–286.
- [72] A.H. Sheikh, D. Lahaye, C. Vuik, On the convergence of shifted Laplace preconditioner combined with multilevel deflation, *Numer. Linear Algebra Appl.* 20 (4) (2013) 645–662.
- [73] F. Sourbier, A. Haiddar, L. Giraud, H. Ben-Hadj-Ali, S. Operto, J. Virieux, Three-dimensional parallel frequency-domain visco-acoustic wave modelling based on a hybrid direct/iterative solver, *Geophys. Prospect.* 59 (5) (2011) 834–856.
- [74] E.A. Spence, Wavenumber-explicit bounds in time-harmonic acoustic scattering, *SIAM J. Math. Anal.* 46 (4) (2014) 2987–3024.
- [75] C.C. Stolk, A rapidly converging domain decomposition method for the Helmholtz equation, *J. Comput. Phys.* 241 (0) (2013) 240–252.
- [76] C.C. Stolk, A dispersion minimizing scheme for the 3-D Helmholtz equation based on ray theory, *J. Comput. Phys.* 314 (2016) 618–646.
- [77] C.C. Stolk, An improved sweeping domain decomposition preconditioner for the Helmholtz equation, *Adv. Comput. Math.* 43 (1) (Feb 2017) 45–76.
- [78] M. Taus, L. Demanet, L. Zepeda-Núñez, A short note on a fast and high-order hybridizable discontinuous Galerkin solver for the 2D high-frequency Helmholtz equation, in: *SEG Technical Program Expanded Abstracts 2016*, 2016, pp. 3835–3840.
- [79] M. Taus, L. Demanet, L. Zepeda-Núñez, A short note on a fast and high-order hybridizable discontinuous Galerkin solver for the 2D high-frequency Helmholtz equation, in: *SEG Technical Program Expanded Abstracts 2016*, Society of Exploration Geophysicists, 2016, pp. 3835–3840.
- [80] M. Taus, L. Zepeda-Núñez, R.J. Hewett, L. Demanet, L-Sweeps/L-Sweeps-2D, Available at, <https://github.com/L-Sweeps/L-Sweeps-2D>, version 1.0.0, <https://doi.org/10.5281/zenodo.3383883>, Sep. 2019.
- [81] M. Taus, L. Zepeda-Núñez, R.J. Hewett, L. Demanet, L-Sweeps/L-Sweeps-2D-examples, Available at, <https://github.com/L-Sweeps/L-Sweeps-2D-examples>, version 1.0.0, <https://doi.org/10.5281/zenodo.3383923>, Sep. 2019.
- [82] Matthias Taus, Laurent Demanet, Leonardo Zepeda Núñez, Russell Hewett, Pollution-free and fast hybridizable discontinuous Galerkin solvers for the high-frequency Helmholtz equation, 2017, pp. 4068–4073.
- [83] R. Thakur, R. Rabenseifner, W. Gropp, Optimization of collective communication operations in mpich, *Int. J. High Perform. Comput. Appl.* 19 (1) (2005) 49–66.
- [84] A. Toselli, O.B. Widlund, *Domain Decomposition Methods – Algorithms and Theory*, Springer Series in Computational Mathematics, vol. 34, Springer, Berlin, Heidelberg, 2005.
- [85] H.A. van der Vorst Bi-CGSTAB, A fast and smoothly converging variant of Bi-CG for the solution of nonsymmetric linear systems, *SIAM J. Sci. Stat. Comput.* 13 (2) (1992) 631–644.
- [86] A. Vion, C. Geuzaine, Double sweep preconditioner for optimized Schwarz methods applied to the Helmholtz problem, *J. Comput. Phys.* 266 (0) (2014) 171–190.

- [87] J. Virieux, S. Operto, An overview of full-waveform inversion in exploration geophysics, *Geophysics* 74 (6) (2009), WCC1–WCC26.
- [88] S. Wang, M.V. de Hoop, J. Xia, X.S. Li, Massively parallel structured multifrontal solver for time-harmonic elastic waves in 3-D anisotropic media, *Geophys. J. Int.* 191 (1) (2012) 346–366.
- [89] J. Xia, S. Chandrasekaran, M. Gu, X.S. Li, Superfast multifrontal method for large structured linear systems of equations, *SIAM J. Matrix Anal. Appl.* 31 (3) (2010) 1382–1411.
- [90] L. Zepeda-Núñez, L. Demanet, A short note on the nested-sweep polarized traces method for the 2D Helmholtz equation, in: *SEG Technical Program Expanded Abstracts 2015*, 2015, pp. 3682–3687.
- [91] L. Zepeda-Núñez, L. Demanet, The method of polarized traces for the 2D Helmholtz equation, *J. Comput. Phys.* 308 (2016) 347–388.
- [92] L. Zepeda-Núñez, L. Demanet, Nested domain decomposition with polarized traces for the 2D Helmholtz equation, *SIAM J. Sci. Comput.* 40 (3) (2018) B942–B981.
- [93] L. Zepeda-Núñez, H. Zhao, Fast alternating bidirectional preconditioner for the 2D high-frequency Lippmann–Schwinger equation, *SIAM J. Sci. Comput.* 38 (5) (2016) B866–B888.
- [94] L. Zepeda-Núñez, R.J. Hewett, L. Demanet, Preconditioning the 2D Helmholtz equation with polarized traces, in: *SEG Technical Program Expanded Abstracts 2014*, 2014, pp. 3465–3470.
- [95] L. Zepeda-Núñez, A. Scheuer, R.J. Hewett, L. Demanet, The method of polarized traces for the 3D Helmholtz equation, *Geophysics* 84 (4) (2019) 1–86.