

NESTED DOMAIN DECOMPOSITION WITH POLARIZED TRACES FOR THE 2D HELMHOLTZ EQUATION*

LEONARDO ZEPEDA-NÚÑEZ[†] AND LAURENT DEMANET[‡]

Abstract. We present a solver for the two-dimensional high-frequency Helmholtz equation in heterogeneous, constant density, acoustic media, with online parallel complexity that scales empirically as $\mathcal{O}(\frac{N}{P})$, where N is the number of volume unknowns, and P is the number of processors, as long as $P = \mathcal{O}(N^{1/5})$. This sublinear scaling is achieved by domain decomposition, not distributed linear algebra, and improves on the $P = \mathcal{O}(N^{1/8})$ scaling reported earlier in [L. Zepeda-Núñez and L. Demanet, *J. Comput. Phys.*, 308 (2016), pp. 347–388]. The solver relies on a two-level nested domain decomposition: a layered partition on the outer level and a further decomposition of each layer in cells at the inner level. The Helmholtz equation is reduced to a surface integral equation (SIE) posed at the interfaces between layers, efficiently solved via a nested version of the polarized traces preconditioner [L. Zepeda-Núñez and L. Demanet, *J. Comput. Phys.*, 308 (2016), pp. 347–388]. The favorable complexity is achieved via an efficient application of the integral operators involved in the SIE.

Key words. high-frequency, wavepropagation, Helmholtz equation, fast methods

AMS subject classification. 65N22

DOI. 10.1137/15M104582X

1. Introduction. It has become clear over the past several years that the right mix of ideas to obtain an efficient Helmholtz solver, in the high-frequency regime, involves domain decomposition with accurate transmission conditions. The first empirical $\mathcal{O}(N)$ complexity algorithm, where N is the number of degrees of freedom, was the sweeping preconditioner of Engquist and Ying [33] that uses a decomposition in grid-spacing-thin layers coupled with an efficient multifrontal solver at each layer. Subsequently, Stolk proposed different instances of domain decomposition methods (DDMs) [88, 89] that restored the ability to use arbitrarily thick layers, improving the efficiency of the local solves at each layer, with the same $\mathcal{O}(N)$ claim. Liu and Ying recently presented a recursive version of the sweeping preconditioner in three dimensions (3D) that decreases the offline cost to linear complexity [68] and a variant of the sweeping preconditioner using an additive Schwarz preconditioner [67]. Many other authors have proposed algorithms with similar complexity claims that we review in section 1.2, which we briefly compare in section 1.3.

Although these algorithms are instances of efficient iterative methods, they have to revisit all the degrees of freedom inside the volume in a sequential (or sweeping) fashion at each iteration. This sequential computation thus hinders the scalability of the algorithms to large-scale parallel architectures.

*Submitted to the journal’s Computational Methods in Science and Engineering section November 13, 2015; accepted for publication (in revised form) February 1, 2018; published electronically June 20, 2018.

<http://www.siam.org/journals/sisc/40-3/M104582.html>

Funding: This project is sponsored by Total SA. The second author was funded by NSF, Air Force Office of Scientific Research and ONR.

[†]Department of Mathematics and Earth Resources Laboratory, Massachusetts Institute of Technology, Cambridge, MA 02139. Current address: Computational Research Division, Lawrence Berkeley National Laboratory, Berkeley, CA 94708 (lzepeda@lbl.gov).

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

Two solutions have been proposed so far to mitigate the lack of asymptotic scalability in the high-frequency regime:

- Poulson et al. [80] parallelized the sweeping preconditioner in 3D by using distributed linear algebra to solve the local system at each layer, obtaining a $\mathcal{O}(N)$ complexity. This approach should in principle reach sublinear complexity scalings in two dimensions (2D). The resulting codes are complex, and using distributed linear algebra libraries can be cumbersome for industrial application because of licensing issues.
- Zepeda-Núñez and Demanet [103] proposed the method of polarized traces, with an online empirical runtime $\mathcal{O}(N/P)$ in 2D, where N is the number of degrees of freedom in the volume¹ and P is the number of nodes in a distributed memory environment, provided that $P = \mathcal{O}(N^{1/8})$ and that the medium does not contain microstructures at the wavelength level with a preferred orientation aligned with the sweeping direction.² The communication cost is a negligible $\mathcal{O}(N^{1/2}P)$.

In this paper we follow and improve on the latter approach. The method of polarized traces relies on a layered domain decomposition coupled with a surface integral equation (SIE) posed at the interfaces between subdomains that is easy to precondition. The algorithm has two stages: an expensive but parallel offline stage that is performed only once and a fast online stage that is performed for each right-hand side (source). The above-mentioned online complexity is the result of, first, the efficient preconditioner, and second, the precomputation and compression of the operators involved in the SIE and its preconditioner during the offline stage.

The main improvements of the solver presented in this paper stem from a nested layered domain decomposition in two levels. In the outer level we use an equivalent matrix-free formulation of the SIE that relies on solving a local problem within each layer, with equivalent sources at the interfaces. In the inner level, we decompose each layer in cells, and we use the original method of polarized traces to solve the local problems efficiently via a local SIE posed at the interfaces between cells within a layer. As will be explained in what follows, the operators involved in the local SIE at each local problem are much smaller, making them cheaper to precompute and to apply. Finally, the number of layers and cells can be balanced to increase the parallelism and hence reduce the asymptotic runtime.

1.1. Results. We propose a variant of the method of polarized traces using a nested domain decomposition approach. This novel approach results in an algorithm with an empirical asymptotic online runtime of $\mathcal{O}(N/P)$ provided that $P = \mathcal{O}(N^{1/5})$, which results in a lower asymptotic online runtime in a distributed memory environment. Moreover, the nested polarized traces method also has lower memory footprint,³ and lower offline complexity, as shown in Table 1.1. The parameter α in Table 1.1 is the exponent of the empirical asymptotic complexity.⁴

The nested polarized traces method inherits the modularity from the original method of polarized traces; it can easily be extended to more complex physics and

¹Throughout this paper, the notation $\mathcal{O}(\cdot)$ may hide logarithmic factors.

²As will be explained in what follows, the combination of higher frequencies and thinner subdomains increases the number of iterations, thus resulting in a mild deterioration of the scalings.

³The offline complexity for the method of polarized traces in this case (Table 1.1) is higher than in [103] because we assume that we have P nodes instead of $N^{1/2}P$.

⁴We choose to express the complexity with respect to the global number of degrees of freedom, N , in order to provide a cleaner expression in Table 1.1. of the application of a compressed numerical Green's function with respect to N . The exact definition of α is given in section 4.3. In this case, the value of $\alpha = 3/4$ is estimated empirically.

TABLE 1.1

Runtime of both formulations (up to poly-logarithmic factors), supposing only P nodes, one node per layer for the method of polarized traces, and one node per cell for the nested domain decomposition with polarized traces. The value of α depends on the ability to compress local Green's functions in partitioned low-rank form; there exist theoretical and practical arguments for setting $\alpha = 3/4$. In that case, the N/P contribution dominates $P^{1-\alpha}N^\alpha$ as long as $P = O(N^{1/5})$. These scalings hold provided that the medium does not contain microstructures at the wavelength level presenting a preferred orientation aligned with the sweeping direction.

Stage	Polarized traces	Nested polarized traces
Offline	$\mathcal{O}(N^{3/2}/P)$	$\mathcal{O}((N/P)^{3/2})$
Online	$\mathcal{O}(PN^\alpha + N/P)$	$\mathcal{O}(P^{1-\alpha}N^\alpha + N/P)$

higher order discretizations, and it takes advantage of new developments in direct methods such as [52, 83, 99] and in better block-low-rank and butterfly matrix compression techniques such as [6, 12, 65].

In addition, we propose a few minor improvements to the original scheme presented in [103] to obtain better accuracy and to accelerate the convergence rate. We provide

1. an equivalent formulation of the method of polarized traces that involves a discretization using Q1 finite elements on a Cartesian grid using a suitable quadrature rule to compute the mass matrix. This formulation can be easily generalized to high order finite differences and high order finite elements, thus circumventing the labor-intensive summation by parts; and
2. a variant of the preconditioner introduced in [103], in which we used a block Gauss–Seidel iteration instead of a block Jacobi iteration, that improves the convergence rate.

Even though the nested approach presented in this paper has a lower asymptotic online runtime than the method of polarized traces, in practice, the latter is faster. This is due to the large constants resulting from iterating within each layer in the nested approach. In order to reduce the constants, we introduce a small variant of the nested approach, which relies on a compressed LU factorization to solve the local SIE within each layer. The asymptotic online runtime remains unchanged, with much lower constants, albeit with a more thorough precomputation.

1.2. Related work. Using domain decomposition to solve PDEs can be dated back to Schwarz [86] and Lions [66], in which the Laplace equation is solved iteratively. However, using such techniques to solve the Helmholtz problem was proposed for the first time by Després in [29], which led to the development of many different approaches focusing mostly on the discretization of the Helmholtz equation; in particular, the ultra weak variational formulation (UWVF) [17], which, in return, spawned plane wave methods such as the Trefftz formulation of Moiola, Hiptmair, and Perugia [76], the plane wave discontinuous Galerkin method [53, 58], the discontinuous enrichment method of Farhat et al. [39], the partition of unity method (PUM) by Babuska and Melenk [4], the least-squares method by Monk and Wang [78], and, more recently, the multitrace formulation of Claeys, Hiptmair, and Jerez-Hanckes [24, 57], among many others. A recent and thorough review can be found in [59].

Moreover, the ideas of Lions and Després led to the development of various domain decomposition algorithms, which can be classified as Schwarz algorithms (for a review on classical Schwarz methods see [18, 90]) with or without overlap for the Helmholtz equation [13, 25, 28, 51, 70, 74]. However, soon it became evident that the

convergence rate of such algorithms was spectacularly dependent on the boundary conditions prescribed at the interfaces between subdomains.

In the quest to design boundary conditions that ensured a fast convergence Gander introduced the framework of optimized Schwarz methods in [41]. Within that framework, Gander, Magoules, and Nataf provided an optimal nonlocal boundary condition, which is then approximated by an optimized Robin boundary condition [44]. How to design better approximations has been studied in [43, 46, 47, 49]. More recently Boubendir, Antoine, and Geuzaine [14] presented a quasi-optimal optimized Schwarz method using Padé approximations of the Dirichlet to Neumann (DtN) map.

The idea of mixing domain decomposition and absorbing boundary condition was first explored by Engquist and Zhao [34] for elliptic problems. The application of such ideas to the Helmholtz problem can be traced back, to great extent, to the analytic incomplete LU preconditioner of Gander and Nataf [45], in which a layered domain decomposition was used, and to Plessix and Mulder [79], in which a similar idea is used with separation of variables. However, it was Engquist and Ying who showed in [33, 32] that such ideas could yield fast methods to solve the high-frequency Helmholtz equation, by introducing the sweeping preconditioner, which was then extended by Tsuji and collaborators to different discretizations and physics [92, 93, 94]. Since then, many other papers have proposed methods with similar claims. Stolk [88] proposed a DDM using single layer potentials to transfer the information between subdomains. Geuzaine and Vion explored randomized techniques [23] to probe the DtN map [8] in order to approximate the absorbing boundary conditions within a multiplicative Schwarz iteration [96, 97]. Chen and Xiang proposed another instance of efficient domain decomposition, where the emphasis is on transferring sources from one subdomain to another [21, 22], which has inspired similar methods [64, 63, 30]. Most of the methods mentioned above can be recast as optimized Schwarz methods (cf. [20]). Closely related to the content of this paper, we find the method of polarized traces [103] and an earlier version of this work [102].

Luo et al. proposed the fast sweeping Huygens method, based on an approximate Green's function computed via geometric optics, coupled with a butterfly algorithm, which can handle transmitted waves in very favorable complexity [69, 81].

Alongside iterative methods, some advances have also been made on multigrid methods. Brandt and Livshits developed the wave-ray method [15], in which the oscillation error is eliminated by a ray-cycle using a geometric optics approximation; Haber and MacLachlan proposed an alternative formulation that can be solved by standard multigrid methods [56]; Erlangga, Oosterlee, and Vuik [37] showed how to implement a simple, although suboptimal, complex-shifted Laplace preconditioner with multigrid. Another variant of complex-shifted Laplacian method with deflation was studied by Sheikh, Lahaye, and Vuik [87]. The choice of the optimal complex-shift was studied by Cools and Vanroose [26] and by Gander, Grahan, and Spence [40]. Finally, the generalization of the complex shifted Laplacian preconditioner to the elastic wave equation has been studied by Rizzuti and Mulder [82]. We point out that most of the multigrid methods mentioned above exhibit a suboptimal dependence of the number of iterations to converge with respect to the frequency, making them ill-suited for high frequency problems. However, they are easy to parallelize, resulting in small runtimes, as shown by Calandra et al. [16].

A good early review of iterative methods for the Helmholtz equation is in [36]. Another review paper that discussed the difficulties generated by the high-frequency limit is [38].

In another exciting direction, much progress has been made on making direct methods efficient for the Helmholtz equation. Such is the case of Wang et al.'s method [99], which couples nested dissection and multifrontal elimination with \mathcal{H} -matrices (see [31] for the multifrontal method and [50] for nested dissection). Another example is the work of Gillman, Barnett, and Martinsson on computing DtN maps in a multiscale fashion [52]. Recently, Ambikasaran et al. [1] proposed a direct solver for acoustic scattering for heterogeneous media using compressed linear algebra to solve an equivalent Lippmann–Schwinger equation. These methods are extremely efficient for elliptic and low frequency problems due to the high compressibility of the Green's functions [7]. However, it is not yet clear whether offline linear complexity scalings for high-frequency problems can be achieved this way (cf. [35]), though good direct methods are often faster in practice than the iterative methods mentioned at the beginning of this section. The main issue with direct methods is the lack of scalability to very large-scale problems due to the memory requirements and prohibitively expensive communication overheads.

Finally, beautiful mathematical reviews of the Helmholtz equation are [77] and [19]. A more systematic and extended exposition of the references mentioned here can be found in [101].

1.3. Comparison to other preconditioners. There has been a recent effort to express the fast preconditioners based on domain decomposition and absorbing boundary conditions in terms of optimized Schwarz method [20, 48, 55], in particular, the method of polarized traces can be recast, algorithmically, as an optimized Schwarz method; however, the polarization of the waves using discrete integral relations, and in particular, the recasting of the volumetric problem as an extended algebraically exact integral problem posed at the interfaces between layers, seems to escape the general framework. We do not aim to compare the method of polarized traces nor its nested variant under the framework of optimized Schwarz methods, which would be out of the scope of this paper. Instead, we aim to briefly compare the method of polarized traces against some related methods in the literature.⁵

We start by comparing the method of polarized traces to the original sweeping preconditioner. The sweeping preconditioner is based on a block LDL^t factorization, coupled with the remarkable observation that the inverse of the blocks of D correspond to a half-space problem. The inverse of the blocks can be either compressed using \mathcal{H} -matrices⁶ [32] or applied using an auxiliary problem [33], in which the half-space problem is truncated using high-quality perfectly matched layers (PML). The sweeping preconditioner acts on the volume problem, and it needs to revisit all the degrees of freedom, thus hindering parallelization, which is mostly achieved by parallelized linear algebra [80]. On the other hand, the method of polarized traces preconditioning an SIE posed on the interfaces reducing the number of degrees of freedom. The application of the SIE and the preconditioner can be parallelized and accelerated using fast summation algorithms, and the number of iterations to convergence is, in general, much smaller.

Another method based on the sweeping preconditioner is the source transfer method [21], an overlapping domain decomposition. In this method, the residual in a generous overlap, times a smooth window, is used to propagate the information

⁵We refer the interested reader to [42] for a short exposition on possible shortcomings of sweeping-like preconditioners.

⁶This is possible, in homogeneous media, thanks to the compressibility of the Green's function when considering elongated structures [72].

to the neighboring domains in a fashion that resembles an iterative refinement iteration. If the window is chosen as a Heaviside function, the source transfer method has some striking similarities to the rapidly converging DDM [88] and the method of polarized traces. However, the source transfer method is a preconditioner for the volume problem, needing to revisit all the degrees of freedom each iteration, and the generous overlap between subdomain can be prohibitively expensive in 3D.

The closest algorithm to the method of polarized traces is the rapidly converging DDM of Stolk [88], in which the information is transferred between subdomains using a single layer potential in [88] and using both single and double layer potential in [89]. The latter has the same form as the sweeps in the method of polarized traces (see [104]). Stolk's method preconditions the volume problem, which requires him to move the boundaries of the subdomains and to use a step of iterative refinement between sweeps. In practice, Stolk's method has lower constants but a higher asymptotic runtime in 2D.

Finally, we compare the method of polarized traces to the double sweep preconditioner, which is, by construction, an optimized Schwarz method. In [97] the Helmholtz problem is posed as a boundary problem between interfaces using different boundary conditions. The boundary problem is then solved iteratively, using an Schwarz iteration accelerated with preconditioned GMRES. The main disadvantage of the double sweep preconditioner is its complexity. To converge in $\mathcal{O}(1)$ iterations, the double sweep preconditioner requires a good knowledge of the DtN map, which is probed using randomized methods (cf. [8]) during an offline stage. As a consequence, the linear systems at each subdomain have a nonlocal boundary condition, which results in matrices with large dense blocks, making the local linear systems expensive to solve.

At the level of the formulation there exist some differences between the two methods; in [97] the resulting system posed on the the interfaces has a block-tridiagonal structure, which is hard to solve iteratively; however, some off-diagonal blocks, which represent self-interactions within layers, are small, so they can be neglected. The resulting sparsified system has an interlaced structure, and it can be solved by back-substitution, in which each block is a local solve with a nonlocal boundary condition. In contrast, the method of polarized traces transitions from a volumetric discretization to an extended *equivalent* boundary problem, without approximation, making it better suited to cases in which the self-interactions, such as multiple scattering within a layer, cannot be neglected. By exploiting the block structure of the latter we can easily construct an efficient preconditioner based on a block Gauss–Seidel iteration within a GMRES [85] or Bi-CGstab [95] iteration. Finally, we point out that the different types of sweeps described in [98] can be seamlessly deduced using different block splittings of the extended SIE.

1.4. Organization. The present paper is organized as follows:

- we review the formulation of the Helmholtz problem and the reduction to a boundary integral equation in section 2;
- in section 3 we review the method of polarized traces;
- in section 4 we present two variants of the nested solver, and we provide the empirical complexity observed; Finally,
- in section 5 we present numerical experiments that corroborate the complexity claims.

2. Formulation. Let Ω be a rectangle in \mathbb{R}^2 , and consider a layered partition of Ω into L slabs, or layers $\{\Omega^\ell\}_{\ell=1}^L$, as shown in Figure 2.1. Define the squared slowness as $m(\mathbf{x}) = 1/c(\mathbf{x})^2$, $\mathbf{x} = (x, z)$. As in geophysics, we may refer to z as depth, and we




FIG. 2.1. Layered domain decomposition. The orange grid-points represent the PML for the original problem, the light-blue grid-points represent the artificial PML between layers, and the red grid-points correspond to \mathbf{u} in (2.11).

suppose that it points downwards. Define the global Helmholtz operator at frequency ω as

$$(2.1) \quad \mathcal{H}u = (-\Delta - m\omega^2) u \quad \text{in } \Omega,$$

with an absorbing boundary condition on $\partial\Omega$ implemented via PML [9, 10, 61].

Let us define f^ℓ as the restriction of f to Ω^ℓ , i.e., $f^\ell = f\chi_{\Omega^\ell}$, where χ_{Ω^ℓ} is the characteristic function of Ω^ℓ . Define the local Helmholtz operators as

$$(2.2) \quad \mathcal{H}^\ell u = (-\Delta - m\omega^2) u \quad \text{in } \Omega^\ell,$$

with absorbing boundary conditions on $\partial\Omega^\ell$. Finally, let u be the solution to $\mathcal{H}u = f$.

As for any DDM, we seek to find u by solving the local systems $\mathcal{H}^\ell v^\ell = f^\ell \chi_{\Omega^\ell}$. In order to compute u , classical DDMs (cf. [90]) require a coupling between the subdomains, which usually takes the form of continuity or boundary conditions. In this case, the global coupling between subdomains is realized via a reduction of the problem posed on the volume to a problem posed on the interfaces between layers, resulting in an SIE. The main tool used in this endeavor is the Green’s representation formula (GRF) in each layer.

If the global Green’s function G , given by $\mathcal{H}G(\mathbf{x}, \mathbf{y}) = \delta(\mathbf{x} - \mathbf{y})$, is known, then we can write in the interior of each layer

$$(2.3) \quad u(\mathbf{x}) = Gf^\ell(\mathbf{x}) + \int_{\partial\Omega^\ell} (G(\mathbf{x}, \mathbf{y})\partial_{\nu_y}u(\mathbf{y})dy - \partial_{\nu_y}G(\mathbf{x}, \mathbf{y})u(\mathbf{y})) dS_y$$

for $\mathbf{x} \in \Omega^\ell$, and $Gf^\ell(\mathbf{x}) = \int_{\Omega} G(\mathbf{x}, \mathbf{y})f^\ell(\mathbf{y})d\mathbf{y}$. One remarkable, and mostly undervalued, property of the GRF is that (2.3) remains true even if we change the Green’s function for a local Green’s function, which we denote G^ℓ , provided that

$$(2.4) \quad \mathcal{H}G^\ell(\mathbf{x}, \mathbf{y}) = \delta(\mathbf{x} - \mathbf{y})$$

is satisfied for $\mathbf{x}, \mathbf{y} \in \Omega^\ell$, where \mathcal{H} acts on \mathbf{x} .

Using the local GRF, the solution can be written *without approximation* in each layer as

$$(2.5) \quad u(\mathbf{x}) = G^\ell f^\ell(\mathbf{x}) + \int_{\partial\Omega^\ell} (G^\ell(\mathbf{x}, \mathbf{y})\partial_{\nu_y}u(\mathbf{y})dy - \partial_{\nu_y}G^\ell(\mathbf{x}, \mathbf{y})u(\mathbf{y})) dS_y$$

for $\mathbf{x} \in \Omega^\ell$, where $G^\ell f^\ell(\mathbf{x}) = \int_{\Omega^\ell} G^\ell(\mathbf{x}, \mathbf{y}) f^\ell(\mathbf{y}) d\mathbf{y}$ and $G^\ell(\mathbf{x}, \mathbf{y})$ is the solution of $\mathcal{H}^\ell G^\ell(\mathbf{x}, \mathbf{y}) = \delta(\mathbf{x} - \mathbf{y})$. Where \mathcal{H}^ℓ coincides with \mathcal{H} in Ω^ℓ by construction, thus G^ℓ satisfies (2.4).

Denote $\Gamma_{\ell, \ell+1} = \partial\Omega^\ell \cap \partial\Omega^{\ell+1}$. Supposing that Ω^ℓ are thin slabs either extending to infinity, or surrounded by a damping layer on the lateral sides, we can rewrite (2.5) as

$$(2.6) \quad \begin{aligned} u(\mathbf{x}) = & G^\ell f^\ell(\mathbf{x}) - \int_{\Gamma_{\ell-1, \ell}} G^\ell(\mathbf{x}, \mathbf{x}') \partial_z u(\mathbf{x}') dx' + \int_{\Gamma_{\ell, \ell+1}} G^\ell(\mathbf{x}, \mathbf{x}') \partial_z u(\mathbf{x}') dx' \\ & + \int_{\Gamma_{\ell-1, \ell}} \partial_z G^\ell(\mathbf{x}, \mathbf{x}') u(\mathbf{x}') dx' - \int_{\Gamma_{\ell, \ell+1}} \partial_z G^\ell(\mathbf{x}, \mathbf{x}') u(\mathbf{x}') dx'. \end{aligned}$$

The knowledge of u and $\partial_z u$ on the interfaces $\Gamma_{\ell, \ell+1}$ therefore suffices to recover the solution everywhere in Ω . We show in section 2.2 how to build, via an algebraic reduction, a discrete SIE (based on (2.6)) posed on the interfaces between boundaries $\{\Gamma_{\ell, \ell+1}\}_{\ell=1}^{L-2}$, and whose solution is exactly the restriction of the global solution to the interfaces between layers. Once the solution at the interfaces is known, the solution can be reconstructed at each layer exactly using (2.6). The remaining question is how to efficiently solve the SIE using an iterative method. The answer lies within the concept of polarization, which is the main topic of the next section.

2.1. Polarization. In this section we provide the rationale for the polarization of waves, a concept that plays a crucial role to solve the SIE mentioned above efficiently. We say that a wave is *polarized* at an interface when it is generated by sources supported only on one side of that interface.

In order to express the polarizing conditions in boundary integral form, we briefly recall the rationale provided in [103]. Let us consider Figure 2.2, where $\mathbf{x} = (x, z)$ with z pointing down. Consider an interface Γ partitioning \mathbb{R}^2 as $\Omega^{\text{down}} \cup \Omega^{\text{up}}$, with $f \neq 0$ in Ω^{down} . Let $\mathbf{x} \in \Omega^{\text{up}}$, and consider a contour made up of Γ and a semicircle D at infinity in the upper half-plane Ω^{up} . In case the wave speed becomes uniform past some large radius, the Sommerfeld radiation condition (SRC) puts to zero the contribution on D in GRF [62, 75], resulting in the *incomplete Green's formula*

$$(2.7) \quad u(\mathbf{x}) = \int_{\Gamma} \left(\frac{\partial G}{\partial z_y}(\mathbf{x}, \mathbf{y}) u(\mathbf{y}) - G(\mathbf{x}, \mathbf{y}) \frac{\partial u}{\partial z_y}(\mathbf{y}) \right) dS_{\mathbf{y}}, \quad \mathbf{x} \in \Omega_{\text{up}} \setminus \Gamma.$$

On the other hand, if x approaches Γ from below, then we obtain the *annihilation formula*

$$(2.8) \quad 0 = \int_{\Gamma} \left(\frac{\partial G}{\partial z_y}(\mathbf{x}, \mathbf{y}) u(\mathbf{y}) - G(\mathbf{x}, \mathbf{y}) \frac{\partial u}{\partial z_y}(\mathbf{y}) \right) dS_{\mathbf{y}}, \quad \mathbf{x} \rightarrow \Gamma \text{ from below.}$$

We can observe that (2.7) and (2.8) are equivalent; either one can be used as the definition of a polarizing boundary condition on Γ . In this case we can say that a wavefield is up-going with respect to Γ if it was generated by a source located below γ , or equivalently, if its Dirichlet and Neumann traces, denoted u^\uparrow and $\partial_z u^\uparrow$, satisfy the annihilation condition (2.8). We can easily express the same annihilation condition for a down-going wavefield by the annihilation formula

$$(2.9) \quad 0 = \int_{\Gamma} \left(-\frac{\partial G}{\partial z_y}(\mathbf{x}, \mathbf{y}) u(\mathbf{y}) + G(\mathbf{x}, \mathbf{y}) \frac{\partial u}{\partial z_y}(\mathbf{y}) \right) dS_{\mathbf{y}}, \quad \mathbf{x} \rightarrow \Gamma \text{ from above.}$$

In order to efficiently solve the globally coupling SIE, we define an extended system by introducing extra variables. We split $u = u^\uparrow + u^\downarrow$ and $\partial_z u = \partial_z u^\uparrow + \partial_z u^\downarrow$ on $\Gamma_{\ell, \ell+1}$, by letting $(u^\uparrow, \partial_z u^\uparrow)$ be polarized up in Ω^ℓ (according to (2.8)) and




FIG. 2.2. Illustration of (2.7) and (2.8).

$(u^\downarrow, \partial_z u^\downarrow)$ polarized down (according to (2.9)) in $\Omega^{\ell+1}$. Together, the interface fields $u^\uparrow, u^\downarrow, \partial_z u^\uparrow, \partial_z u^\downarrow$ are the polarized traces that serve as unknowns for the numerical method.

The discrete system is then set up from discrete algebraic reformulations⁷ of the local GRF (2.6) with the polarizing conditions (2.8) and (2.9), in a manner that will be made explicit in section 3. The resulting system has a 2×2 structure with block-triangular submatrices on the diagonal and comparably small off-diagonal submatrices. A very good preconditioner consists in inverting the block-triangular submatrices by back- and forward-substitution. One application of this preconditioner can be seen as a sweep of the domain to compute transmitted (as well as locally reflected) waves using (2.7). From extensive numerical experiments provided in [103], the structure of the system ensures that the number of GMRES iterations grows at most as $\log \omega$, provided that the m does not contain a large resonant cavity or microstructures at the wavelength level with a preferred orientation aligned with the domain decomposition.

The rationale presented above can be seamlessly translated to the algebraic level, as will be done in section 3, in which, for the sake of reproducibility, we follow a mostly algebraic description of the algorithm.

2.2. Discrete realization. In order to compress the notation, and yet provide enough details so that the reader can implement the algorithm presented in this paper, we recall the notation used in [103].

We can discretize (2.1) using second order finite differences (Appendix A) or Q1 finite elements (Appendix B), and we obtain the global linear system

$$(2.10) \quad \mathbf{H}\mathbf{u} = \mathbf{f},$$

which we aim to solve using a domain decomposition approach that relies on solving the local systems \mathbf{H}^ℓ , which are the discrete version of (2.2).

As explained previously, the coupling between the subdomains is realized via an *equivalent* discrete SIE, which relies on a discrete version of (2.6). In this section we briefly explain, at an algebraic level, the reduction of (2.10) to an equivalent discrete SIE of the form

⁷The algebraic reformulation is performed using either summation by parts, as done in [103], or the much less labor intensive approach shown in Appendix C.

$$(2.11) \quad \mathbf{M}\mathbf{u} = \mathbf{f},$$

where \mathbf{u} are the degrees of freedom of \mathbf{u} at the interfaces between layers (see Figure 2.1) and \mathbf{M} is defined below after some basic notation has been introduced.

We suppose that the full domain has $N = n_x \times n_z$ discretization points, that each layer has $n_x \times n^\ell$ discretization points, and that the number of points in the PML, n_{pml} , is the same in both dimension and in every subdomain (light-blue and orange nodes in Figure 2.1).

In both discretizations the mesh is structured so that we can define $\mathbf{x}_{p,q} = (x_p, z_q) = (ph, qh)$. We assume the same ordering as in [103], i.e.,

$$(2.12) \quad \mathbf{u} = (\mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_{n_z}),$$

and we use the notation

$$(2.13) \quad \mathbf{u}_j = (u_{1,j}, u_{2,j}, \dots, u_{n_x,j}),$$

for the entries of \mathbf{u} sampled at constant depth z_j . We write \mathbf{u}^ℓ for the wavefield defined locally at the ℓ th layer, i.e., $\mathbf{u}^\ell = \chi_{\Omega^\ell} \mathbf{u}$, and \mathbf{u}_k^ℓ for the values at the local depth⁸ z_k^ℓ of \mathbf{u}^ℓ . In particular, \mathbf{u}_1^ℓ and $\mathbf{u}_{n^\ell}^\ell$ are the top and bottom rows⁹ of \mathbf{u}^ℓ (see red grid-points in Figure 2.1). We then gather the interface traces in the vector

$$(2.14) \quad \underline{\mathbf{u}} = (\mathbf{u}_{n^1}^1, \mathbf{u}_1^2, \mathbf{u}_{n^2}^2, \dots, \mathbf{u}_1^{L-1}, \mathbf{u}_{n^{L-1}}^{L-1}, \mathbf{u}_1^L)^t.$$

Define the numerical local Green's function in layer ℓ by

$$(2.15) \quad \mathbf{H}^\ell \mathbf{G}^\ell(\mathbf{x}_{i,j}, \mathbf{x}_{i',j'}) = \mathbf{H}^\ell \mathbf{G}_{i,j,i',j'}^\ell = \delta(\mathbf{x}_{i,j} - \mathbf{x}_{i',j'}),$$

if $(i, j) \in \llbracket -n_{\text{pml}} + 1, n_x + n_{\text{pml}} \rrbracket \times \llbracket -n_{\text{pml}} + 1, n^\ell + n_{\text{pml}} \rrbracket$,¹⁰ where the discrete dirac delta¹¹ is given by

$$(2.16) \quad \delta(\mathbf{x}_{i,j} - \mathbf{x}_{i',j'}) = \begin{cases} \frac{1}{h^2} & \text{if } \mathbf{x}_{i,j} = \mathbf{x}_{i',j'}, \\ 0 & \text{if } \mathbf{x}_{i,j} \neq \mathbf{x}_{i',j'}, \end{cases}$$

and where the operator \mathbf{H}^ℓ acts on the (i, j) indices.

Furthermore, for notational convenience we consider \mathbf{G}^ℓ as an operator acting on unknowns at the interfaces, as follows.

DEFINITION 2.1. *We consider $\mathbf{G}^\ell(z_j, z_k)$ as the linear operator defined from $\llbracket -n_{\text{pml}} + 1, n_x + n_{\text{pml}} \rrbracket \times \{z_k\}$ to $\llbracket -n_{\text{pml}} + 1, n_x + n_{\text{pml}} \rrbracket \times \{z_j\}$ given by*

$$(2.17) \quad (\mathbf{G}^\ell(z_j, z_k)\mathbf{v})_i = h \sum_{i'=-n_{\text{pml}}+1}^{n_x+n_{\text{pml}}} \mathbf{G}^\ell((x_i, z_j), (x_{i'}, z_k))\mathbf{v}_{i'},$$

where \mathbf{v} is a vector in $\mathbb{C}^{n_x+2n_{\text{pml}}}$ and $\mathbf{G}^\ell(z_j, z_k)$ are matrices in $\mathbb{C}^{(n_x+2n_{\text{pml}}) \times (n_x+2n_{\text{pml}})}$. Moreover, (2.17) is the discrete counterpart of

$$(2.18) \quad \int_{\mathbb{R}} G^\ell((x, z_j), (x', z_k))v(x', z_k)dx'.$$

⁸We hope that there is little risk of confusion in overloading z_j (local indexing) for $z_{n_c^\ell+j}$ (global indexing), where $n_c^\ell = \sum_{j=1}^{\ell-1} n^j$ is the cumulative number of points in depth.

⁹We do not consider the PML points here.

¹⁰The set denoted by $\llbracket a, b \rrbracket$ is equivalent to $\{i \in \mathbb{Z}; a \leq i \leq b\}$.

¹¹The definition is for a finite difference discretization. We refer the reader to the Appendix C for a more general definition.

The interface-to-interface operator $\mathbf{G}^\ell(z_j, z_k)$ is indexed by two depths—following the jargon from fast methods; we call them source depth (z_k) and target depth (z_j). In particular, it represents the wavefield sampled at z_j produced by a source (in this case a measure) located at z_k .

DEFINITION 2.2. We consider $\mathcal{G}_j^{\uparrow, \ell}(\mathbf{v}_{n^\ell}, \mathbf{v}_{n^{\ell+1}})$ the up-going local incomplete Green’s integral and $\mathcal{G}_j^{\downarrow, \ell}(\mathbf{v}_0, \mathbf{v}_1)$ the down-going local incomplete Green’s integral, as defined by

$$(2.19) \quad \mathcal{G}_j^{\uparrow, \ell}(\mathbf{v}_{n^\ell}, \mathbf{v}_{n^{\ell+1}}) = \mathbf{G}^\ell(z_j, z_{n^{\ell+1}}) \left(\frac{\mathbf{v}_{n^{\ell+1}} - \mathbf{v}_{n^\ell}}{h} \right) - \left(\frac{\mathbf{G}^\ell(z_j, z_{n^{\ell+1}}) - \mathbf{G}^\ell(z_j, z_{n^\ell})}{h} \right) \mathbf{v}_{n^{\ell+1}},$$

$$(2.20) \quad \mathcal{G}_j^{\downarrow, \ell}(\mathbf{v}_0, \mathbf{v}_1) = -\mathbf{G}^\ell(z_j, z_0) \left(\frac{\mathbf{v}_1 - \mathbf{v}_0}{h} \right) + \left(\frac{\mathbf{G}^\ell(z_j, z_1) - \mathbf{G}^\ell(z_j, z_0)}{h} \right) \mathbf{v}_0.$$

In what follows we use the shorthand notation $\mathbf{G}^\ell(z_j, z_k) = \mathbf{G}_{j,k}^\ell$ when explicitly building the matrix form of the integral systems.

We can observe that Definition 2.2 is the discrete counterpart of (2.7), in which we used two neighboring traces to define the normal derivative. The expression above is the result of a discrete Green’s representation formula that can be deduced from a laborious summation by parts (for more details, see [103]).

Finally, we define the Newton potential as resulting from a local solve inside each layer.

DEFINITION 2.3. Consider the local Newton potential \mathcal{N}_k^ℓ applied to a local source \mathbf{f}^ℓ as

$$(2.21) \quad \mathcal{N}_k^\ell \mathbf{f}^\ell = \sum_{j=1}^{n^\ell} \mathbf{G}^\ell(z_k, z_j) \mathbf{f}_j^\ell.$$

By construction, $\mathcal{N}^\ell \mathbf{f}^\ell$ satisfies the equation $(\mathbf{H}^\ell \mathcal{N}^\ell \mathbf{f}^\ell)_{i,j} = \mathbf{f}_{i,j}$ for $-n_{pml} + 1 \leq i \leq n_x + n_{pml}$ and $1 \leq j \leq n^\ell$.

Following the notation introduced above, the discrete SIE reduction of the original discrete Helmholtz equation, issued from (2.6), takes the form

$$(2.22) \quad \mathcal{G}_1^{\downarrow, \ell}(\mathbf{u}_0^\ell, \mathbf{u}_1^\ell) + \mathcal{G}_1^{\uparrow, \ell}(\mathbf{u}_{n^\ell}^\ell, \mathbf{u}_{n^{\ell+1}}^\ell) + \mathcal{N}_1^\ell \mathbf{f}^\ell = \mathbf{u}_1^\ell,$$

$$(2.23) \quad \mathcal{G}_{n^\ell}^{\downarrow, \ell}(\mathbf{u}_0^\ell, \mathbf{u}_1^\ell) + \mathcal{G}_{n^\ell}^{\uparrow, \ell}(\mathbf{u}_{n^\ell}^\ell, \mathbf{u}_{n^{\ell+1}}^\ell) + \mathcal{N}_{n^\ell}^\ell \mathbf{f}^\ell = \mathbf{u}_{n^\ell}^\ell,$$

$$(2.24) \quad \mathbf{u}_{n^\ell}^\ell = \mathbf{u}_0^{\ell+1}, \quad \mathbf{u}_{n^{\ell+1}}^\ell = \mathbf{u}_1^{\ell+1},$$

if $1 < \ell < L$, with

$$(2.25) \quad \mathcal{G}_{n^1}^{\uparrow, 1}(\mathbf{u}_{n^1}^1, \mathbf{u}_{n^1+1}^1) + \mathcal{N}_{n^1}^1 \mathbf{f}^1 = \mathbf{u}_{n^1}^1, \quad \mathbf{u}_{n^1}^1 = \mathbf{u}_0^2, \quad \mathbf{u}_{n^1+1}^1 = \mathbf{u}_1^2,$$

and

$$(2.26) \quad \mathcal{G}_1^{\downarrow, L}(\mathbf{u}_0^L, \mathbf{u}_1^L) + \mathcal{N}_{n^L}^L \mathbf{f}^L = \mathbf{u}_1^L, \quad \mathbf{u}_{n^L-1}^L = \mathbf{u}_0^L, \quad \mathbf{u}_{n^L-1+1}^L = \mathbf{u}_1^L.$$

This was what we referred to as $\mathbf{Mu} = \mathbf{f}$ in (2.11) and whose structure is depicted in Figure 3.1 (left).

Finally, the online stage of the algorithm is summarized in Algorithm 1.

Algorithm 1. Online computation using the SIE reduction.

```

1: function  $\mathbf{u} = \text{HELMHOLTZ SOLVER}(\mathbf{f})$ 
2:   for  $\ell = 1 : L$  do
3:      $\mathbf{f}^\ell = \mathbf{f}_{\chi_{\Omega^\ell}}$  ▷ partition the source
4:   end for
5:   for  $\ell = 1 : L$  do
6:      $\mathcal{N}^\ell \mathbf{f}^\ell = (\mathbf{H}^\ell)^{-1} \mathbf{f}^\ell$  ▷ solve local problems (parallel)
7:   end for
8:    $\underline{\mathbf{f}} = (\mathcal{N}_1^1 \mathbf{f}^1, \mathcal{N}_1^2 \mathbf{f}^2, \mathcal{N}_2^2 \mathbf{f}^2, \dots, \mathcal{N}_1^L \mathbf{f}^L)^t$  ▷ form right-hand side (r.h.s.) for the
integral system
9:    $\underline{\mathbf{u}} = (\underline{\mathbf{M}})^{-1} \underline{\mathbf{f}}$  ▷ solve (2.11) for the traces
10:  for  $\ell = 1 : L$  do
11:     $\mathbf{u}_j^\ell = \mathcal{G}_j^{\uparrow, \ell}(\mathbf{u}_{n^\ell}^\ell, \mathbf{u}_{n^{\ell+1}}^\ell) + \mathcal{G}_j^{\downarrow, \ell}(\mathbf{u}_0^\ell, \mathbf{u}_1^\ell) + \mathcal{N}_j^\ell \mathbf{f}^\ell$  ▷ reconstruct local solutions
12:  end for
13:   $\mathbf{u} = (\mathbf{u}^1, \mathbf{u}^2, \dots, \mathbf{u}^{L-1}, \mathbf{u}^L)^t$  ▷ concatenate the local solutions
14: end function

```




FIG. 3.1. Sparsity pattern of the SIE matrix in (2.11) (left) and the polarized SIE matrix in (3.8) (right).

3. The method of polarized traces. In this section, we review succinctly the formulation of the method of polarized traces at an algebraic level; for further details see [101]. From Algorithm 1 we can observe that the local solves (lines 5-7 in Algorithm 1) and the reconstruction (lines 10-12 in Algorithm 1) can be performed concurrently; the only sequential bottleneck is the solution of (2.11) (line 9 in Algorithm 1). For the sake of clarity we use matrices to explain the preconditioner; however, all the operations can be performed in a matrix-free fashion, as will be explained in section 4.1.

The method of polarized traces was developed to solve (2.11) efficiently. The method utilizes an extended equivalent SIE formulation which relies on

- a decomposition of the wavefield at the interfaces in two components, up-going and down-going;
- integral relations to close the new extended system;
- a permutation of the unknowns to obtain an easily preconditionable system via classical matrix splitting (see section 4.2.2 of [84]).

Following the notation introduced in section 2.2, the resulting extended system (see section 3.5 in [103]) can be written as

$$(3.1) \quad \mathcal{G}_1^{\uparrow,\ell}(\mathbf{u}_{n^\ell}^{\ell,\uparrow}, \mathbf{u}_{n^{\ell+1}}^{\ell,\uparrow}) + \mathcal{G}_1^{\downarrow,\ell}(\mathbf{u}_0^{\ell,\downarrow}, \mathbf{u}_1^{\ell,\downarrow}) + \mathcal{G}_1^{\uparrow,\ell}(\mathbf{u}_{n^\ell}^{\ell,\downarrow}, \mathbf{u}_{n^{\ell+1}}^{\ell,\downarrow}) + \mathcal{N}_1^\ell \mathbf{f}^\ell = \mathbf{u}_1^{\ell,\uparrow} + \mathbf{u}_1^{\ell,\downarrow},$$

$$(3.2) \quad \mathcal{G}_{n^\ell}^{\downarrow,\ell}(\mathbf{u}_0^{\ell,\uparrow}, \mathbf{u}_1^{\ell,\uparrow}) + \mathcal{G}_{n^\ell}^{\uparrow,\ell}(\mathbf{u}_{n^\ell}^{\ell,\uparrow}, \mathbf{u}_{n^{\ell+1}}^{\ell,\uparrow}) + \mathcal{G}_{n^\ell}^{\downarrow,\ell}(\mathbf{u}_0^{\ell,\downarrow}, \mathbf{u}_1^{\ell,\downarrow}) + \mathcal{N}_{n^\ell}^\ell \mathbf{f}^\ell = \mathbf{u}_{n^\ell}^{\ell,\uparrow} + \mathbf{u}_{n^\ell}^{\ell,\downarrow},$$

$$(3.3) \quad \mathcal{G}_0^{\uparrow,\ell}(\mathbf{u}_{n^\ell}^{\ell,\uparrow}, \mathbf{u}_{n^{\ell+1}}^{\ell,\uparrow}) + \mathcal{G}_0^{\downarrow,\ell}(\mathbf{u}_0^{\ell,\downarrow}, \mathbf{u}_1^{\ell,\downarrow}) + \mathcal{G}_0^{\uparrow,\ell}(\mathbf{u}_{n^\ell}^{\ell,\downarrow}, \mathbf{u}_{n^{\ell+1}}^{\ell,\downarrow}) + \mathcal{N}_0^\ell \mathbf{f}^\ell = \mathbf{u}_0^{\ell,\uparrow},$$

$$(3.4) \quad \mathcal{G}_{n^{\ell+1}}^{\downarrow,\ell}(\mathbf{u}_0^{\ell,\uparrow}, \mathbf{u}_1^{\ell,\uparrow}) + \mathcal{G}_{n^{\ell+1}}^{\uparrow,\ell}(\mathbf{u}_{n^\ell}^{\ell,\uparrow}, \mathbf{u}_{n^{\ell+1}}^{\ell,\uparrow}) + \mathcal{G}_{n^{\ell+1}}^{\downarrow,\ell}(\mathbf{u}_0^{\ell,\downarrow}, \mathbf{u}_1^{\ell,\downarrow}) + \mathcal{N}_{n^{\ell+1}}^\ell \mathbf{f}^\ell = \mathbf{u}_{n^{\ell+1}}^{\ell,\downarrow},$$

for $\ell = 1, \dots, L$ or equivalently

$$(3.5) \quad \underline{\underline{\mathbf{M}}} \underline{\underline{\mathbf{u}}} = \underline{\underline{\mathbf{f}}}, \quad \underline{\underline{\mathbf{u}}} = \begin{pmatrix} \underline{\mathbf{u}}^\downarrow \\ \underline{\mathbf{u}}^\uparrow \end{pmatrix},$$

where we write

$$(3.6) \quad \underline{\mathbf{u}}^\downarrow = \left(\mathbf{u}_{n^1}^{\downarrow,1}, \mathbf{u}_{n^{1+1}}^{\downarrow,1}, \mathbf{u}_{n^2}^{\downarrow,2}, \dots, \mathbf{u}_{n^{L-1}}^{\downarrow,L-1}, \mathbf{u}_{n^{L-1+1}}^{\downarrow,L-1} \right)^t,$$

$$(3.7) \quad \underline{\mathbf{u}}^\uparrow = \left(\mathbf{u}_0^{\uparrow,2}, \mathbf{u}_1^{\uparrow,2}, \mathbf{u}_0^{\uparrow,3}, \dots, \mathbf{u}_0^{\uparrow,L}, \mathbf{u}_1^{\uparrow,L} \right)^t,$$

to define the components of the polarized wavefields, and $\underline{\mathbf{u}}^\downarrow + \underline{\mathbf{u}}^\uparrow = \underline{\mathbf{u}}$. The indices and the arrows are chosen such that they reflect the propagation direction. For example, $\mathbf{u}_{n^1}^{\downarrow,\ell}$ represents the wavefield leaving the layer ℓ at its bottom, i.e., propagating downwards and sampled at the bottom of the layer.

After a permutation of the entries (see section 3.5 in [103], in particular Figure 5), and some basic algebraic operations, the matrix in (3.5) takes the form

$$(3.8) \quad \underline{\underline{\mathbf{M}}} = \begin{bmatrix} \underline{\mathbf{D}}^\downarrow & \underline{\mathbf{U}} \\ \underline{\mathbf{L}} & \underline{\mathbf{D}}^\uparrow \end{bmatrix},$$

where $\underline{\mathbf{D}}^\downarrow$ and $\underline{\mathbf{D}}^\uparrow$ are, respectively, block-lower triangular and block-upper triangular matrices with identity diagonal blocks, thus easily invertible using a block back-substitution (see Figure 3.1 (right)).

Finally, the method of polarized traces seeks to solve the system in (3.5) using an iterative method, such as GMRES, coupled with an efficient preconditioner issued from a matrix splitting, which relies on the application of $(\underline{\mathbf{D}}^\downarrow)^{-1}$ and $(\underline{\mathbf{D}}^\uparrow)^{-1}$.

We point out that the blocks of $\underline{\underline{\mathbf{M}}}$ have a physical meaning. $\underline{\mathbf{D}}^\downarrow$ takes into account the waves propagating downwards, $\underline{\mathbf{D}}^\uparrow$ considers the waves propagating upwards, $\underline{\mathbf{U}}$ takes into account the reflections of waves propagating upwards and being reflected downwards, and $\underline{\mathbf{L}}$ takes into account the down-going waves reflected upwards.

3.1. Gauss–Seidel preconditioner. In this paper, we use a block Gauss–Seidel iteration as a preconditioner to solve the polarized system in (3.5) instead of the block Jacobi iteration used in [103]. The Gauss–Seidel preconditioner is given by

$$(3.9) \quad P^{\text{GS}} \begin{pmatrix} \underline{\mathbf{v}}^\downarrow \\ \underline{\mathbf{v}}^\uparrow \end{pmatrix} = \begin{pmatrix} (\underline{\mathbf{D}}^\downarrow)^{-1} \underline{\mathbf{v}}^\downarrow \\ (\underline{\mathbf{D}}^\uparrow)^{-1} \left(\underline{\mathbf{v}}^\uparrow - \underline{\mathbf{L}}(\underline{\mathbf{D}}^\downarrow)^{-1} \underline{\mathbf{v}}^\downarrow \right) \end{pmatrix},$$

and the Jacobi preconditioner is given by

$$(3.10) \quad P^{\text{Jac}} \begin{pmatrix} \underline{\mathbf{v}}^\downarrow \\ \underline{\mathbf{v}}^\uparrow \end{pmatrix} = \begin{pmatrix} (\underline{\mathbf{D}}^\downarrow)^{-1} \underline{\mathbf{v}}^\downarrow \\ (\underline{\mathbf{D}}^\uparrow)^{-1} \underline{\mathbf{v}}^\uparrow \end{pmatrix}.$$

In our experiments, solving (3.8) using GMRES, or alternatively Bi-CGstab, preconditioning with P^{GS} converges twice as fast as using P^{Jac} as a preconditioner, and

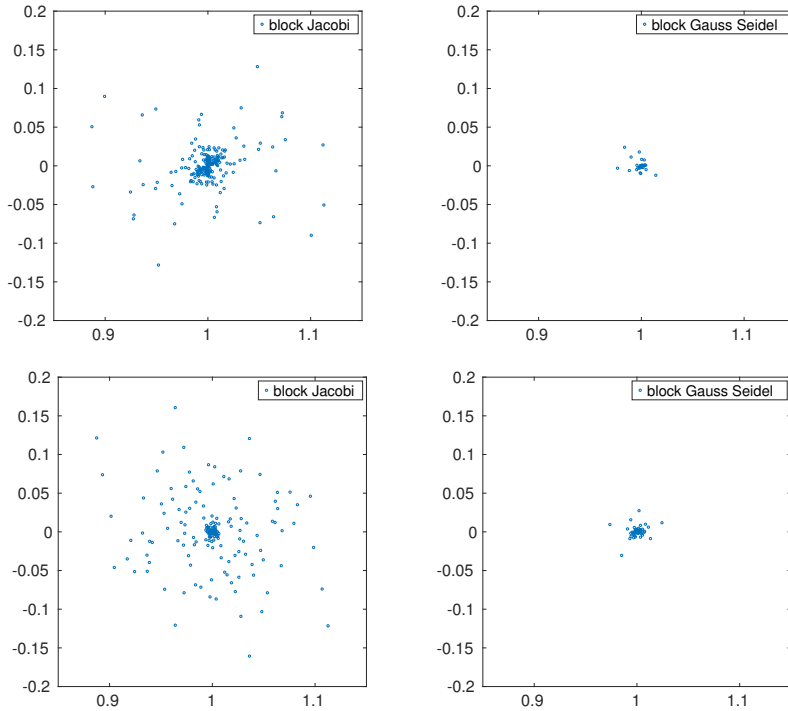


FIG. 3.2. Eigenvalues for the preconditioned polarized systems using the block Jacobi (left) and the block Gauss–Seidel (right) preconditioner, using the BP 2004 model [11] with $L = 5$, $npml = 10$, and $\omega = 34\pi$ (top row) and $\omega = 70\pi$ (bottom row).

the former exhibits a weaker dependence of the number of iterations for convergence with respect to the frequency. We considered other standard preconditioners, such as symmetric successive overrelaxation (SSOR) (see section 10.2 in [84]), but they failed to yield faster convergence, while being more computationally expensive to apply.

We point out that $\underline{\underline{\mathbf{M}}}$ can be partitioned in smaller blocks. In that case, we can recover the X and NX sweeps (see Table 3 of [89]) when we precondition the system with a Jacobi or Gauss–Seidel iteration, respectively.

Figure 3.2 depicts the eigenvalues for $\underline{\underline{\mathbf{M}}}$ preconditioned with P^{GS} and P^{Jac} . We can observe that for P^{GS} the eigenvalues are more clustered and there exist fewer outliers. There exists extensive numerical evidence that indicates that a tighter clustering of the spectrum away from zero can be related to a smaller number of iterations needed for convergence¹² thus explaining the faster convergence of the GMRES iterations preconditioned with P^{Jac} .

The system in (3.5) is solved using GMRES preconditioned with P^{GS} . Moreover, as in [103], one can use an adaptive \mathcal{H} -matrix fast algorithm for the application of integral kernels, which in this case are used for the solution of the local problems defined in each layer.

¹²If the preconditioned systems were represented by normal matrices, then from Theorem 35.2 of [91] the clustering of the eigenvalues would explain the smaller number of iterations needed for convergence. For a more extensive treatment of nonnormal matrices, see [54] and references therein.




FIG. 4.1. *Nested decomposition in cells. The orange grid-points represent the PML for the original problem, the light-blue grid-points represent the artificial PML between layers, and the pink grid-points represent the artificial PML between cells in the same layer. Compare to Figure 2.1.*

4. Nested solver. The main drawback of the method of polarized traces is its offline precomputation that involves computing, storing, and compressing the interface-to-interface Green's functions needed to assemble $\underline{\underline{\mathbf{M}}}$. In 3D this approach would become impractical given the sheer size of the resulting matrices. To alleviate this issue, we present an equivalent matrix-free approach that relies on local solves with sources at the interfaces between layers.

As will be explained in what follows, the matrix-free approach relies on the fact that the blocks of $\underline{\underline{\mathbf{M}}}$ (as well as the blocks of $\underline{\mathbf{D}}^\downarrow$ and $\underline{\mathbf{D}}^\uparrow$) are the restrictions of local Green's functions. Thus they can be applied via a local solve (using, for example, a multifrontal sparse direct solver such as [2, 60, 27] among many others) with sources at the interfaces. This same observation was used in [103] to reconstruct the solution in the volume (see section 2.3, in particular (28), in [103]). However, given the iterative nature of P^{GS} (that relies on inverting $\underline{\mathbf{D}}^\downarrow$ and $\underline{\mathbf{D}}^\uparrow$ by block-back-substitution), solving the local problems naively would incur a deterioration of the online complexity, in particular, the parallelization. This deterioration can be circumvented if we solve the local problems inside the layer via the same boundary integral strategy as in the method of polarized traces, in a nested fashion. This procedure can be written as a factorization of the Green's integral in block-sparse factors, as will be explained in section 4.2.

The nested domain decomposition approach involves a layered decomposition in $L \sim \sqrt{P}$ layers, such that each layer is further decomposed in $L_c \sim \sqrt{P}$ cells, as shown in Figure 4.1.

In addition to the lower online complexity achieved by the nested approach, the offline complexity is much reduced; instead of computing large Green's functions for each layer, we compute much smaller interface-to-interface operators between the interfaces of adjacent cells within each layer, resulting in a lower memory footprint.

The nested approach consists of two levels:

- the *outer solver*, which solves the global Helmholtz problem, (2.10), using the matrix-free version of the method of polarized traces to solve (2.11) at the interfaces between layers, and

- the *inner solver*, which solves the local Helmholtz problems at each layer, using an integral boundary equation to solve for the degrees of freedom at the interfaces between cells within a layer.

4.1. Matrix-free approach. We proceed to explain how to implement the method of polarized traces using the matrix-free approach. We point out that the matrix-free approach is only used in the outer SIE; we still need to assemble and factorize the local systems (or local SIEs); however, the outer SIE is never assembled.

As stated before, the backbone of the method of polarized traces is to solve (3.8) iteratively with GMRES using (3.9) as a preconditioner. In this section we explain how to apply $\underline{\mathbf{M}}$ and P^{GS} in a matrix-free fashion, and we provide the pseudocode for the method of polarized traces using the matrix-free approach with the local solves explicitly identified.

From (3.1), (3.2), (3.3), and (3.4), each block of $\underline{\mathbf{M}}$ is a Green’s integral, and its application to a vector is equivalent to sampling a wavefield generated by suitable sources at the boundaries. The application of the Green’s integral to a vector $\underline{\mathbf{v}}$, in a matrix-free approach, consists of three steps:

- from $\underline{\mathbf{v}}$ we form the sources at the interfaces,
- we perform a local direct solve inside the layer,
- and we sample the solution at the interfaces.

The precise algorithm to apply $\underline{\mathbf{M}}$ in a matrix-free fashion is provided in Algorithm 2.

Algorithm 2. Application of the boundary integral matrix $\underline{\mathbf{M}}$.

```

1: function  $\underline{\mathbf{u}} = \text{BOUNDARY INTEGRAL}(\underline{\mathbf{v}})$ 
2:    $\tilde{\mathbf{f}}^1 = -\delta(z_{n^1+1} - z)\mathbf{v}_{n^1}^1 + \delta(z_{n^1} - z)\mathbf{v}_{n^1}^2$            ▷ forming equivalent sources
3:    $\mathbf{w}^1 = (\mathbf{H}^1)^{-1}\tilde{\mathbf{f}}^1$ 
4:    $\mathbf{u}_{n^1}^\ell = \mathbf{w}_{n^1}^\ell - \mathbf{v}_{n^1}^\ell$                                        ▷ sampling
5:   for  $\ell = 2 : L - 1$  do
6:      $\tilde{\mathbf{f}}^\ell = \delta(z_1 - z)\mathbf{v}_{n^{\ell-1}}^{\ell-1} - \delta(z_0 - z)\mathbf{v}_1^\ell$ 
            $-\delta(z_{n^{\ell+1}} - z)\mathbf{v}_{n^{\ell}}^\ell + \delta(z_{n^{\ell}} - z)\mathbf{v}_1^{\ell+1}$            ▷ forming equivalent sources
7:      $\mathbf{w}^\ell = (\mathbf{H}^\ell)^{-1}\tilde{\mathbf{f}}^\ell$                                        ▷ inner solve
8:      $\mathbf{u}_1^\ell = \mathbf{w}_1^\ell - \mathbf{v}_1^\ell$ ;    $\mathbf{u}_{n^{\ell}}^\ell = \mathbf{w}_{n^{\ell}}^\ell - \mathbf{v}_{n^{\ell}}^\ell$            ▷ sampling
9:   end for
10:   $\tilde{\mathbf{f}}^L = \delta(z_1 - z)\mathbf{v}_{n^{L-1}}^{L-1} - \delta(z_0 - z)\mathbf{v}_1^L$            ▷ forming equivalent sources
11:   $\mathbf{w}^L = (\mathbf{H}^L)^{-1}\tilde{\mathbf{f}}^L$ 
12:   $\mathbf{u}_1^L = \mathbf{w}_1^L - \mathbf{v}_1^L$                                        ▷ sampling
13: end function

```

Algorithm 2 can be easily generalized for $\underline{\mathbf{M}}$. We observe that there is no data dependency within the **for** loop, which yields an embarrassingly parallel algorithm.

Matrix-free preconditioner. For the sake of clarity we present a high level description of the implementation of (3.9) using the matrix-free version.

We use the notation introduced in section 3 (in particular, (3.6) and (3.7)) to write explicitly the matrix-free operations for the block Gauss–Seidel preconditioner in (3.9). Algorithm 3 and 4 have the physical interpretation of propagating the waves across the domains, and Algorithm 5 can be seen as the up-going reflections generated by a down-going wavefield. The following algorithms can be easily derived from section 3.5 in [103].

Algorithm 3. Downward sweep, application of $(\underline{\mathbf{D}}^\downarrow)^{-1}$.

```

1: function  $\underline{\mathbf{u}}^\downarrow = \text{DOWNWARD SWEEP}(\underline{\mathbf{v}}^\downarrow)$ 
2:    $\mathbf{u}_{n^1}^{\downarrow,1} = -\mathbf{v}_{n^1}^{\downarrow,1}$  ▷ invert the diagonal block
3:    $\mathbf{u}_{n^1+1}^{\downarrow,1} = -\mathbf{v}_{n^1+1}^{\downarrow,1}$ 
4:   for  $\ell = 2 : L - 1$  do
5:      $\mathbf{w}^\ell = (\mathbf{H}^\ell)^{-1} \left[ \delta(z_0 - z) \mathbf{u}_{n^{\ell-1}+1}^{\downarrow,\ell-1} - \delta(z_1 - z) \mathbf{u}_{n^{\ell-1}}^{\downarrow,\ell-1} \right]$  ▷ inner solve
6:      $\mathbf{u}_{n^\ell}^{\downarrow,\ell} = \mathbf{w}_{n^\ell}^\ell - \mathbf{v}_{n^\ell}^{\downarrow,\ell}$  ▷ sample the wavefield and subtract the r.h.s.
7:      $\mathbf{u}_{n^{\ell+1}}^{\downarrow,\ell} = \mathbf{w}_{n^{\ell+1}}^\ell - \mathbf{v}_{n^{\ell+1}}^{\downarrow,\ell}$  ▷ sample the wavefield and subtract the r.h.s.
8:   end for
9:    $\underline{\mathbf{u}}^\downarrow = \left( \mathbf{u}_{n^1}^{\downarrow,1}, \mathbf{u}_{n^1+1}^{\downarrow,1}, \mathbf{u}_0^{\downarrow,2}, \dots, \mathbf{u}_0^{\downarrow,L-1}, \mathbf{u}_1^{\downarrow,L-1} \right)^t$ 
10: end function

```

Algorithm 4. Upward sweep, application of $(\underline{\mathbf{D}}^\uparrow)^{-1}$.

```

1: function  $\underline{\mathbf{u}}^\uparrow = \text{UPWARD SWEEP}(\underline{\mathbf{v}}^\uparrow)$ 
2:    $\mathbf{u}_0^{\uparrow,L} = -\mathbf{v}_0^{\uparrow,L}$  ▷ invert the diagonal block
3:    $\mathbf{u}_1^{\uparrow,L} = -\mathbf{v}_1^{\uparrow,L}$ 
4:   for  $\ell = L - 1 : 2$  do
5:      $\mathbf{w}^\ell = (\mathbf{H}^\ell)^{-1} \left[ -\delta(z_{n^{\ell+1}} - z) \mathbf{u}_1^{\uparrow,\ell-1} + \delta(z_{n^\ell} - z) \mathbf{u}_0^{\uparrow,\ell-1} \right]$  ▷ inner solve
6:      $\mathbf{u}_1^{\uparrow,\ell} = \mathbf{w}_1^\ell - \mathbf{v}_1^{\uparrow,\ell}$  ▷ sample the wavefield and subtract the r.h.s.
7:      $\mathbf{u}_0^{\uparrow,\ell} = \mathbf{w}_0^\ell - \mathbf{v}_0^{\uparrow,\ell}$  ▷ sample the wavefield and subtract the r.h.s.
8:   end for
9:    $\underline{\mathbf{u}}^\uparrow = \left( \mathbf{u}_0^{\uparrow,1}, \mathbf{u}_1^{\uparrow,1}, \mathbf{u}_{n^2}^{\uparrow,2}, \dots, \mathbf{u}_{n^{L-1}}^{\uparrow,L-1}, \mathbf{u}_{n^{L-1}+1}^{\uparrow,L-1} \right)^t$ 
10: end function

```

Algorithm 5. Upward reflections, application of $\underline{\mathbf{L}}$.

```

1: function  $\underline{\mathbf{u}}^\uparrow = \text{UPWARD REFLECTIONS}(\underline{\mathbf{v}}^\downarrow)$ 
2:   for  $\ell = 2 : L - 1$  do
3:      $\mathbf{f}^\ell = \delta(z_1 - z) \mathbf{v}_0^{\downarrow,\ell} - \delta(z_0 - z) \mathbf{v}_1^{\downarrow,\ell} - \delta(z_{n^{\ell+1}} - z) \mathbf{v}_1^{\downarrow,\ell+1} + \delta(z_{n^\ell} - z) \mathbf{v}_0^{\downarrow,\ell+1}$ 
4:      $\mathbf{w}^\ell = (\mathbf{H}^\ell)^{-1} \mathbf{f}^\ell$  ▷ inner solve
5:      $\mathbf{u}_1^{\uparrow,\ell} = \mathbf{w}_1^\ell - \mathbf{v}_1^{\downarrow,\ell}$  ▷ sample the wavefield and subtract the identity
6:      $\mathbf{u}_0^{\uparrow,\ell} = \mathbf{w}_0^\ell$  ▷ sample the wavefield
7:   end for
8:    $\mathbf{f}^L = \delta(z_1 - z) \mathbf{v}_0^{\uparrow,L} - \delta(z_0 - z) \mathbf{v}_1^{\uparrow,L}$ 
9:    $\mathbf{w}^L = (\mathbf{H}^L)^{-1} \mathbf{f}^L$  ▷ local solve
10:   $\mathbf{u}_1^{\uparrow,L} = \mathbf{w}_1^L - \mathbf{v}_1^{\downarrow,L}$  ▷ sample the wavefield and subtract the identity
11:   $\mathbf{u}_0^{\uparrow,L} = \mathbf{w}_0^L$  ▷ sample the wavefield
12:   $\underline{\mathbf{u}}^\uparrow = \left( \mathbf{u}_0^{\uparrow,2}, \mathbf{u}_1^{\uparrow,2}, \mathbf{u}_{n^2}^{\uparrow,3}, \dots, \mathbf{u}_{n^{L-1}}^{\uparrow,L-1}, \mathbf{u}_{n^L+1}^{\uparrow,L} \right)^t$ 
13: end function

```

We observe that the **for** loop in lines 2–7 in Algorithm 5 is completely parallel. On the other hand, in Algorithms 3 and 4, the data dependency within the for loop forces the algorithm to run sequentially. The most expensive operation is the inner solve

Algorithm 6. Matrix-free polarized traces solver.

```

1: function  $\mathbf{u} = \text{MATRIX-FREE SOLVER}(\mathbf{f})$ 
2:   for  $\ell = 1 : L$  do
3:      $\mathbf{f}^\ell = \mathbf{f}\chi_{\Omega^\ell}$  ▷ partition the source
4:      $\mathbf{w}^\ell = (\mathbf{H}^\ell)^{-1}(\mathbf{f}^\ell)$  ▷ solve local problems
5:   end for
6:    $\underline{\mathbf{f}} = (\mathbf{w}_{n_1}^1, \mathbf{w}_1^2, \dots, \mathbf{w}_1^L)^t$ ;    $\underline{\mathbf{f}}_0 = (\mathbf{w}_{n_1+1}^1, \mathbf{w}_0^2, \dots, \mathbf{w}_0^L)^t$ 
7:    $\underline{\underline{\mathbf{f}}} = \begin{pmatrix} \underline{\mathbf{f}} \\ \underline{\mathbf{f}}_0 \end{pmatrix}$  ▷ form the r.h.s. for the polarized integral system
8:    $\begin{pmatrix} \underline{\mathbf{u}}^\downarrow \\ \underline{\mathbf{u}}^\uparrow \end{pmatrix} = \underline{\underline{\mathbf{u}}} = (P^{\text{GS}}\underline{\underline{\mathbf{M}}})^{-1} P^{\text{GS}}\underline{\underline{\mathbf{f}}}$  ▷ solve using GMRES
9:    $\underline{\mathbf{u}} = \underline{\mathbf{u}}^\uparrow + \underline{\mathbf{u}}^\downarrow$  ▷ add the polarized components
10:   $\tilde{\mathbf{f}}^1 = \mathbf{f}^1 - \delta(z_{n_1+1} - z)\mathbf{u}_{n_1}^1 + \delta(z_{n_1} - z)\mathbf{u}_1^2$  ▷ reconstruct local solutions
11:   $\mathbf{u}^1 = (\mathbf{H}^1)^{-1}(\tilde{\mathbf{f}}^1)$ 
12:  for  $\ell = 2 : L - 1$  do
13:     $\tilde{\mathbf{f}}^\ell = \mathbf{f}^\ell + \delta(z_1 - z)\mathbf{u}_{n_{\ell-1}}^{\ell-1} - \delta(z_0 - z)\mathbf{u}_1^\ell$ 
14:     $\quad - \delta(z_{n_\ell+1} - z)\mathbf{u}_{n_\ell}^\ell + \delta(z_{n_\ell} - z)\mathbf{u}_1^{\ell+1}$ 
15:     $\mathbf{u}^\ell = (\mathbf{H}^\ell)^{-1}(\tilde{\mathbf{f}}^\ell)$ 
16:  end for
17:   $\tilde{\mathbf{f}}^L = \mathbf{f}^L + \delta(z_1 - z)\mathbf{u}_{n_{L-1}}^{L-1} - \delta(z_0 - z)\mathbf{u}_1^L$ 
18:   $\mathbf{u}^L = (\mathbf{H}^L)^{-1}(\tilde{\mathbf{f}}^L)$ 
19:   $\mathbf{u} = (\mathbf{u}^1, \mathbf{u}^2, \dots, \mathbf{u}^{L-1}, \mathbf{u}^L)^t$  ▷ concatenate the local solutions
20: end function

```

performed locally at each layer. We will argue in section 4.2 that by using a nested approach, with an appropriate reduction of the degrees of freedom, we can obtain a highly efficient inner solve, which yields a fast application of the preconditioner.

Matrix-free solver. We provide the full algorithm of the matrix-free solver using the method of polarized traces coupled with the Gauss–Seidel preconditioner. The main difference with the original method of polarized traces in [103] is that we use Algorithms 2, 3, 4, and 5 within the GMRES iteration (line 8 of Algorithm 6) instead of compressed matrix-vector multiplications.

The matrix-free solver in Algorithm 6 has three stages:

- lines 2–7: preparation of the r.h.s. for the outer polarized integral system, which can be done concurrently within each layer;
- lines 8–9: solve for the traces at the interfaces between layers, using preconditioned GMRES and applying $\underline{\underline{\mathbf{M}}}$ and the preconditioner via the matrix-free approach with Algorithms 2, 3, 4, and 5;
- lines 10–18: reconstruction of the solution inside the volume at each layer, which can be performed independently of the other layers.

4.2. Nested inner and outer solver. In the presentation of the matrix-free solver (Algorithm 6), we have extensively relied on the assumption that the inner systems \mathbf{H}^ℓ can be solved efficiently in order to apply the Green’s integrals fast. In this section we describe the algorithms to compute the solutions to the inner, or local, systems efficiently, and then we describe how the outer solver calls the inner solver.

From the analysis of the rank of the off-diagonal blocks of the Green’s functions, we know that the Green’s integrals can be compressed in a way that results in a

fast application in $\mathcal{O}(n^{3/2})$ time (see section 5 in [103]), but this approach requires precomputation and storage of the Green's functions. The matrix-free approach in Algorithm 6 does not need expensive precomputations, but it would naively perform a direct solve in the volume (inverting \mathbf{H}^ℓ), resulting in an application of the Green's integral in $\mathcal{O}(N/L)$ complexity up to logarithmic factors (assuming that a good direct method is used at each layer). This becomes problematic when applying the preconditioner, which involves $\mathcal{O}(L)$ sequential applications of the Green's integrals as Algorithms 3 and 4 show. This means that the unaided application of the preconditioner using the matrix-free approach would result in an algorithm with linear online complexity; in particular, the method would behave similarly to a sweeping-like preconditioner. The nested strategy (which we present below) mitigates this issue resulting in a lower $\mathcal{O}(L_c(n/L)^{3/2})$ (up to logarithmic factors) complexity for the application of each Green's integral, where L_c is the number of cells per layer.

We follow the matrix-free approach of Algorithm 6, but instead of a direct solver to invert \mathbf{H}^ℓ , we use a nested solver, i.e., we use the same reduction used in Ω to each layer Ω^ℓ . We reduce the local problem at each layer to solving a discrete integral system analog to (2.11) with a layered decomposition in the transverse direction given by

$$(4.1) \quad \mathbf{M}^\ell \mathbf{u}^\ell = \mathbf{f}^\ell \quad \text{for } \ell = 1, \dots, L_c.$$

We suppose that we have $L_c \sim L \sim \sqrt{P}$ cells in each layer.

The nested solver uses the inner boundaries, or interfaces between cells, as proxies to perform the local solve inside the layer efficiently. The efficiency can be improved when the inner solver is used in the applications of the Green's integral within the preconditioner. The improved efficiency stems from the localization of the sources, and the sampling of the solution on the interfaces, which allows us to precompute and compress some for the operations. In that case, the application of the Green's integral can be decomposed into three steps:

- using precomputed Green's functions at each cell, we evaluate the wavefield generated from the sources to form \mathbf{f}^ℓ (from red to pink in Figure 4.2 (left)); this operation can be represented by a sparse block matrix \mathbf{M}_f^ℓ ;
- we solve (4.1) to obtain \mathbf{u}^ℓ (from pink to blue in Figure 4.2 (right)); finally,
- we use the Green's representation formula to sample the wavefield at the interfaces (from blue to green in Figure 4.2); this operation is represented by another sparse-block matrix \mathbf{M}_u^ℓ .




FIG. 4.2. Sketch of the application of the Green's functions using a nested approach. The sources are in red (left) and the sampled field in green (right). The application uses the inner boundaries as proxies to perform the solve.

Using the definition of the incomplete integrals in section 2.2, the algorithm described above leads to the factorization

$$(4.2) \quad \begin{bmatrix} \mathcal{G}_0^{\downarrow,\ell}(\mathbf{v}_0, \mathbf{v}_1) + \mathcal{G}_0^{\uparrow,\ell}(\mathbf{v}_{n^\ell}, \mathbf{v}_{n^\ell+1}) \\ \mathcal{G}_1^{\downarrow,\ell}(\mathbf{v}_0, \mathbf{v}_1) + \mathcal{G}_1^{\uparrow,\ell}(\mathbf{v}_{n^\ell}, \mathbf{v}_{n^\ell+1}) \\ \mathcal{G}_2^{\downarrow,\ell}(\mathbf{v}_0, \mathbf{v}_1) + \mathcal{G}_2^{\uparrow,\ell}(\mathbf{v}_{n^\ell}, \mathbf{v}_{n^\ell+1}) \\ \mathcal{G}_{n^\ell+1}^{\downarrow,\ell}(\mathbf{v}_0, \mathbf{v}_1) + \mathcal{G}_{n^\ell+1}^{\uparrow,\ell}(\mathbf{v}_{n^\ell}, \mathbf{v}_{n^\ell+1}) \end{bmatrix} = \underline{\mathbf{M}}_u^\ell \left(\underline{\mathbf{M}}^\ell \right)^{-1} \underline{\mathbf{M}}_f^\ell \cdot \begin{bmatrix} \mathbf{v}_0 \\ \mathbf{v}_1 \\ \mathbf{v}_{n^\ell} \\ \mathbf{v}_{n^\ell+1} \end{bmatrix},$$

in which the blocks of $\underline{\mathbf{M}}_f^\ell$ and $\underline{\mathbf{M}}_u^\ell$ are dense but compressible in partitioned low rank (PLR)¹³ form.

Algorithms. We now provide the algorithms in pseudocode for the the inner solver, Algorithm 7, and we provide the necessary modifications to Algorithm 6 for the outer solve. In addition, we provide a variant of the inner solver that is crucial to obtain the online complexity mentioned at the beginning of the paper (i.e., $\mathcal{O}(N/P)$ provided that $P = \mathcal{O}(N^{1/5})$).

In order to reduce the notational burden, we define the inner solve using the same notation as before. We suppose that each layer Ω^ℓ is decomposed in L_c cells, noted $\{\Omega^{\ell,c}\}_{c=1}^{L_c}$. We extend all the definitions from the matrix-free solver to the inner solver by indexing the operations by ℓ and c , in which ℓ stands for the layer-index and c for the cell-index within the layer.

For each Ω^ℓ , we apply the variable swap $\tilde{\mathbf{x}} = (z, x)$, which is noted by \mathcal{R} such that $\mathcal{R}^2 = I$. Under the variable swap, we can decompose Ω^ℓ in L_c layers, or cells, $\{\Omega^{\ell,c}\}_{c=1}^{L_c}$, to which we can apply the machinery of the boundary integral reduction at the interfaces between cells. The resulting algorithm has the same structure as before. The variable swap is a suitable tool that allows us to reuse to great extent the notation introduced in [103]. Numerically, the variable swap just introduced is implemented by transposing the matrices that represent the different wavefields.

To obtain the nested solver, we modify Algorithm 6 and the algorithm it calls by replacing $(\mathbf{H}^\ell)^{-1}$ by the inner solver.

- If the support of the source is the whole layer and the wavefield is required in the volume, we use the inner solve as prescribed in Algorithm 7 without modifications. In Algorithm 6 we modify lines 4, 11, 14, and 17, in which $(\mathbf{H}^\ell)^{-1}$ is replaced by Algorithm 7.
- If the source term is concentrated at the interfaces between layers, and the wavefield is needed only at the interfaces, we reduce the computational cost by using a slight modification of Algorithm 7, according to (4.2). In this variant, the local solves in line 7 of Algorithm 7 (which is performed via a LU back-substitution) and the reconstruction (lines 11 to 15 in Algorithm 7) are replaced by precomputed operators as explained above. Within the GMRES loop (line 10 in Algorithm 6), we replace $(\mathbf{H}^\ell)^{-1}$ with the variant of Algorithm 7 following (4.2) in line 5 of Algorithm 3, line 4 of Algorithm 4, and lines 4 and 9 of Algorithm 5.

The choice of algorithm to solve (4.1) and to apply the Green’s integrals dictates the scaling of the offline complexity and the constant of the online complexity. We can either use the method of polarized traces or the compressed-block LU solver, which are explained below.

¹³A PLR matrix is a \mathcal{H} -matrix obtained by an adaptive dyadic partitioning and multilevel compression; see section 3 in [103].

Algorithm 7. Inner solve, for applying $(\mathbf{H}^\ell)^{-1}$ in Algorithms 3, 4, and 5.

```

1: function  $\mathbf{w} = \text{INNER SOLVER}^\ell(\mathbf{f}^\ell)$ 
2:    $\mathbf{g}^\ell = \mathcal{R} \circ \mathbf{f}^\ell$  ▷ variable swap
3:   for  $c = 1 : L_c$  do
4:      $\mathbf{g}^{\ell,c} = \mathbf{g}\chi_{\Omega^{\ell,c}}$  ▷ partition the source
5:   end for
6:   for  $c = 1 : L_c$  do
7:      $\mathcal{N}^{\ell,c} \mathbf{g}^{\ell,c} = (\mathbf{H}^{\ell,c})^{-1} \mathbf{g}^{\ell,c}$  ▷ solve local problems
8:   end for
9:    $\underline{\mathbf{g}}^\ell = \left( \mathcal{N}_{n^1}^{\ell,1} \mathbf{g}^{\ell,1}, \mathcal{N}_1^{\ell,2} \mathbf{g}^{\ell,2}, \mathcal{N}_{n^2}^{\ell,2} \mathbf{g}^{\ell,2}, \dots, \mathcal{N}_1^{\ell,L_c} \mathbf{g}^{\ell,L_c} \right)^t$  ▷ form r.h.s.
10:   $\underline{\mathbf{v}}^\ell = \left( \underline{\mathbf{M}}^\ell \right)^{-1} \underline{\mathbf{g}}^\ell$  ▷ solve for the traces (4.1)
11:  for  $c = 1 : L_c$  do ▷ local reconstruction
12:     $\mathbf{v}_j^{\ell,c} = \mathcal{G}_j^{\uparrow,\ell,c}(\mathbf{v}_{n^\ell}^{\ell,c}, \mathbf{v}_{n^{\ell+1}}^{\ell,c}) + \mathcal{G}_j^{\downarrow,\ell,c}(\mathbf{v}_0^{\ell,c}, \mathbf{v}_1^{\ell,c}) + \mathcal{N}_j^{\ell,c} \mathbf{g}^{\ell,c}$ 
13:  end for
14:   $\mathbf{v}^\ell = (\mathbf{v}^{\ell,1}, \mathbf{v}^{\ell,2}, \dots, \mathbf{v}^{\ell,L_c-1}, \mathbf{v}^{\ell,L_c})^t$  ▷ concatenate the local solutions
15:   $\mathbf{w} = \mathcal{R} \circ \mathbf{v}^\ell$  ▷ variable swap
16: end function

```

4.2.1. Nested polarized traces. To efficiently apply the Green's integrals using Algorithm 7, we need to solve (4.1) efficiently. One alternative is to use the method of polarized traces in a recursive fashion to solve the system at each layer. We call this approach the method of nested polarized traces. Following [103] this approach has the same empirical scalings, at the inner level, as those found in [103] when the blocks are compressed in PLR form.

Although, as will be explained in section 4.3, in this case the complexity is lower, we have to iterate inside each layer to solve each system, which produces large constants for the application of the Green's integrals in the online stage.

4.2.2. Inner compressed-block LU. An alternative to efficiently apply the Green's integrals via Algorithm 7 is to use the compressed-block LU solver (see Chapter 3 in [101]) to solve (4.1). Given the banded structure of $\underline{\mathbf{M}}^\ell$ (see Figure 3.1 (left)), we perform a block LU decomposition without pivoting. The resulting LU factors are block sparse and tightly banded. We have the factorization

$$(4.3) \quad \underline{\mathbf{M}}^\ell = \underline{\mathbf{L}}^\ell \underline{\mathbf{U}}^\ell,$$

which leads to

$$(4.4) \quad \mathcal{G}^\ell = \underline{\mathbf{M}}_u^\ell (\underline{\mathbf{U}}^\ell)^{-1} (\underline{\mathbf{L}}^\ell)^{-1} \underline{\mathbf{M}}_f^\ell,$$

in which \mathcal{G}^ℓ represents the linear operator at the left-hand side of (4.2). Following section 3 in [101], $(\underline{\mathbf{U}}^\ell)^{-1} (\underline{\mathbf{L}}^\ell)^{-1}$ can be done in the same complexity as the nested polarized traces, at the price of a more thorough precomputation. The improved complexity is achieved by inverting the diagonal blocks of the LU factors, thus reducing $(\underline{\mathbf{U}}^\ell)^{-1}$ and $(\underline{\mathbf{L}}^\ell)^{-1}$ to a sequence of matrix-vector multiplications that are further accelerated by compressing the matrices in PLR form.

The main advantage of the inner compressed solver with respect to using the method of polarized traces in the layer solve is that we do not need to iterate, and the system to solve is half the size. Therefore, the online constants are much lower than using the method of polarized traces for the inner solve.

TABLE 4.1

Number of nodes active, complexity, and communication cost of the different steps of the preconditioner, in which α depends on the compression of the local matrices, thus on the scaling of the frequency with respect to the number of unknowns. The value of α depends on the compression rate of the discrete integral operator, which depends on the scaling of ω , N , and the maximum rank of the blocks (see section 5.3, in particular, Table 4, of [103]). Typically $\alpha = 3/4$.

Step	N_{nodes} active	Complexity per node	Communication
LU factorizations	$\mathcal{O}(P)$	$\mathcal{O}((N/P + \log(N))^{3/2})$	$\mathcal{O}(1)$
Green's functions	$\mathcal{O}(P)$	$\mathcal{O}((N/P + \log(N))^{3/2})$	$\mathcal{O}(1)$
Local solves	$\mathcal{O}(P)$	$\mathcal{O}(N/P + \log(N)^2)$	$\mathcal{O}(1)$
Sweeps	1	$\mathcal{O}(P(N/P + \log(N)^2)^\alpha)$	$\mathcal{O}(PN^{1/2})$
Reconstruction	$\mathcal{O}(P)$	$\mathcal{O}(N/P + \log(N)^2)$	$\mathcal{O}(1)$

4.3. Complexity. Table 4.1 summarizes the complexities and number of processors at each stage of the nested polarized traces method in section 4.2.1.

The runtimes and complexities presented in Table 4.1 use the following setup: We suppose that we have available $\mathcal{O}(P)$ nodes, and they are organized following the domain decomposition in Figure 4.1. In particular, we have $L \sim \sqrt{P}$ layers, $\mathcal{O}(\sqrt{P})$ nodes per layer, and $\mathcal{O}(1)$ nodes per cell, i.e., $L_c \sim \sqrt{P}$ and $L \cdot L_c \sim P$. We suppose that the factorized matrices and compressed Green's functions are local to the nodes associated to a cell. The cells communicate only with their neighbors induced by the topology of the decomposition. For simplicity we do not count the logarithmic factors from the nested dissection; however, we consider the logarithmic factors coming from the extra degrees of freedom in the PML.¹⁴

4.3.1. Offline complexity. The offline stage is composed of the LU factorizations at each cell containing $\mathcal{O}(N/P + \log(N))$ points, the computation of the Green's function that involves solving $\mathcal{O}(n/\sqrt{P})$ local systems in each cell, and the compression of the resulting Green's functions in PLR form. The complexities for each node are presented in Table 4.1; furthermore, the compression of the Green's functions takes a negligible $\mathcal{O}(n/\sqrt{P} \log(n/\sqrt{P}))$ time per node (see Table 2 and section 5.1 in [103]) using randomized methods (cf. [73]) to accelerate the compression step.

From Table 4.1, the offline stage is embarrassingly parallel at the cell level, and it has an overall runtime $\mathcal{O}((N/P)^{3/2})$, up to logarithmic factors, as stated in Table 1.1.

When using the inner compressed-block LU method, we have the same complexity but with an extra $\mathcal{O}(N^{3/2}/P)$ cost in the offline stage, making it comparable to the method of polarized traces in [103], although with a lower online complexity.

4.3.2. Online complexity. For the online stage, the runtime of the local solves and the reconstruction in each cell is independent of the frequency and embarrassingly parallel; the runtime is then dominated by the complexity of multifrontal methods, as stated in Table 4.1.

The sweeps are, however, fully sequential, as shown in Table 4.1. Moreover, given the nested nature of the preconditioner, we have two kinds of sweeps:

- the inner sweeps within a layer that are used to apply the Green's integral in the outer SIE,
- and the outer sweeps, which sweep from layer to layer applying the Green's integrals in a matrix-free fashion, therefore relying on the inner sweeps.

The runtime of the inner sweeps depends on the compression rate of the integral operators involved in the local SIE, which is given by α that is exponent of the

¹⁴They are more visible in the runtime scalings presented below.

empirical asymptotic complexity of the application of integral operators with respect to N . The value of α depends on the scaling between the frequency and the number of degrees of freedom per dimension. If the frequency scales as $\omega \sim \sqrt{n}$, the regime in which second order finite-differences and Q1 finite elements are expected to be quasi-optimal, we empirically obtain $\alpha = 5/8$ (see Figure 13 in [103]); however, we assume the more conservative value $\alpha = 3/4$. The latter is in better agreement with a theoretical analysis of the ranks of the off-diagonal blocks of the Green's functions under a geometrical optics approximation (see section 5, in particular, Table 4 and Lemma 7, in [103]; for further details on the compressibility of Green's functions for smooth heterogeneous media see [35]).

In such a scenario, the complexity of the application of the Green's integral at the layer level depends on the complexity of the application of each of the factors in the right-hand side of (4.2), i.e., $\underline{\mathbf{M}}_u^\ell$, $(\underline{\mathbf{M}}^\ell)^{-1}$, and $\underline{\mathbf{M}}_f^\ell$.

For the operators $\underline{\mathbf{M}}_u^\ell$ and $\underline{\mathbf{M}}_f^\ell$, we can compress their blocks in PLR format (see section 5 of [103]). From (4.2) and Figure 4.2, we can clearly see that $\underline{\mathbf{M}}_u^\ell$ and $\underline{\mathbf{M}}_f^\ell$ are integral operators with blocks involving the numerical Green's function local to each cell sampled at the interfaces. Each block can be represented by a matrix of size $(n/L + \log(n)) \times (n/L + \log(n))$,¹⁵ thus after compression it is possible to apply each block in $\mathcal{O}((n/L + \log(n))^{2\alpha})$ time. This last statement is backed from the extensive numerical experiments whose results are summarized in Table 4 of [103], in which the complexity of the application of a compressed $n \times n$ matrix issued from the discrete Green's function is $\mathcal{O}(n^{2\alpha})$ or, following the fact that $N = n^2$, $\mathcal{O}(N^\alpha)$. Finally, given that the operators $\underline{\mathbf{M}}_u^\ell$ and $\underline{\mathbf{M}}_f^\ell$ have L_c blocks, they can be applied in $\mathcal{O}(L_c(n/L + \log(n))^{2\alpha})$ time.

For the application of $(\underline{\mathbf{M}}^\ell)^{-1}$, the remaining term in (4.2), we can solve (4.1) using either the compressed-block LU or the nested polarized traces in $\mathcal{O}(L_c(n/L + \log(n))^{2\alpha})$ time. We follow the same procedure as in the method of polarized traces: we build an extended local SIE, and we solve the local SIE iteratively using preconditioned GMRES. The solve is accelerated by compressing the blocks of the local SIE. Each block of the local SIE is represented by an $(n/L + \log(n)) \times (n/L + \log(n))$ matrix; thus, following section 4 in [103], the application of each block can be performed in $\mathcal{O}((n/L + \log(n))^{2\alpha})$ time. Therefore, given that each layer has L_c cells, the application of the local SIE and the local preconditioner after compression can be performed in $\mathcal{O}(L_c(n/L + \log(n))^{2\alpha})$ time. For the case of the compressed-block LU variant, a similar argument provides the same asymptotic complexity.

This yields a runtime of $\mathcal{O}(L_c(n/L + \log(n))^{2\alpha})$ for each application of the Green's integral, at the layer level, using the factorization in (4.2).

Finally, to apply the Gauss–Seidel preconditioner on the outer SIE, we need to perform outer sweeps, each requiring $\mathcal{O}(L)$ applications of the Green's integrals, resulting in a runtime of $\mathcal{O}(L \cdot L_c(n/L + \log(n))^{2\alpha})$ to solve (3.5). Using the fact that $L \sim \sqrt{P}$, $L_c \sim \sqrt{P}$, $N = n^2$ and adding the contribution of the other steps of the online stage, we have that the overall online runtime is given by $\mathcal{O}(P^{1-\alpha}N^\alpha + P \log(N)^\alpha + N/P + \log(N)^2)$. Supposing that $P = \mathcal{O}(N)$ and neglecting the logarithmic factors, we have that the overall runtime is given by $\mathcal{O}(P^{1-\alpha}N^\alpha + N/P)$ as stated in Table 1.1.

Moreover, if $\alpha = 3/4$, then we have that the online complexity is $\mathcal{O}(N/P)$ (up to logarithmic factors) provided that $P = \mathcal{O}(N^{1/5})$. The communication cost for

¹⁵The logarithmic factor comes from the fact that the traces are taken transversally to each layer; then we need to consider the extra PML nodes, which grow as $\mathcal{O}(\log n)$.

the online part is $\mathcal{O}(nP)$, and the memory footprint is $\mathcal{O}(P^{1/4}N^{3/4} + P \log(N)^{3/4} + N/P + \log(N)^2)$, which represents an asymptotic improvement with respect to [103], in which the memory footprint is $\mathcal{O}(PN^{3/4} + N/P + \log(N)^2)$.

It was already explained in [103] why $\alpha = 3/4$ is a reasonable assumption. Empirically, α is often closer to $5/8$, but this seems to be an overly optimistic preasymptotic scaling. Theoretically, the case can be made for $\alpha = 3/4$ in the continuous geometrical optics scenario when $G(x, y) = A_\omega(x, y)e^{i\omega\Phi(x, y)}$ for smooth $A_\omega(x, y)$ and $\Phi(x, y)$ except when $x \neq y$, and $A_\omega(x, y)$ further depends on ω in a harmless polyhomogeneous way.¹⁶ When $\omega \sim n$, then it is easy to show (by factoring out the leading plane wave) that the largest constant-rank blocks have size $O(\sqrt{\omega}) = O(\sqrt{n})$. Further bookkeeping on the partitioned low-rank structure induced by these blocks shows that the compressed matrix-vector multiplication can be realized in $O(n^{3/2}) = O(N^{3/4})$ operations, yielding $\alpha = 3/4$. This argument is not rigorous for two reasons: (i) geometrical optics may not be a good approximation, and (ii) it does not take into account the fact that the Green's function for the discretized problem may be far from that of the PDE. If the Green's functions are not compressed, it is clear that $\alpha = 1$.

5. Numerical results. The code used for the numerical experiments was written in MATLAB, and the experiments were performed in a dual socket server with two Xeon E5-2670 and 384 GB of RAM. Given the lack of parallelism of the MATLAB implementation, we only benchmark the sequential bottleneck of the online computation, which is the only nonembarrassingly parallel operation. Following Table 4.1 the communication cost is asymptotically negligible, so we focus the benchmarks on the number of iterations needed for convergence and on the compressibility of the Green's functions at the inner level.

In this section we provide numerical evidence of the following claims:

- the fast convergence of the method of polarized traces for a typical geophysical benchmark model and its agnosticity to the type of sources;
- the behavior of the preconditioner using two different matrix splittings and different Krylov iterative methods to solve the outer problem;
- the dependence of the number of iteration to convergence with respect to the frequency, the contrast of the medium, the number of subdomains, and the number of layers in layered media;
- the effectiveness of the compression of the local matrices to reduce the runtime; and,
- the complexity reduction due to the compression of the local matrices coupled with a proper scaling of the number of subdomains.




FIG. 5.1. *BP 2004 geophysical benchmark model* [11].

¹⁶i.e., $A_\omega(x, y)$ has no highly oscillatory components depending on ω .

TABLE 5.1

Maximum number of iterations required to reduce the relative residual to 10^{-5} , using GMRES with the Gauss–Seidel preconditioner, for 50 realizations of a randomly located point source and a randomly generated global source. In this case, $f = \omega/2\pi \sim n$, the number of layers and the number of cells inside each layer grow as n , the number of points in the PML scales as $\log(N)$, and the sound speed is given by the BP 2004 model (see [11]).

N	$\omega/2\pi$ [Hz]	$L \times L_c$	$s = \delta(\mathbf{x})$	$s = \text{rand}$
60×169	1.25	6×3	3	3
120×338	2.5	12×6	4	4
239×675	5.0	24×12	4	5
478×1349	10.0	48×24	5	5
955×2697	20.0	96×48	5	6




FIG. 5.2. Two iterations of the preconditioner, from left to right and top to bottom: initial guess with only local solves, first iteration, second iteration, solution. The background model is given by the BP 2004 model [11] shown in Figure 5.1.

5.1. Fast convergence. Figure 5.2 depicts the fast convergence of the method when using the BP 2004 model¹⁷ (see Figure 5.1), point sources, and a layered domain decomposition with vertical layers. After a couple of iterations, the exact and approximated solutions are indistinguishable to the naked eye. Figure 5.3 depicts the fast convergence of the method of polarized traces when using the BP 2004 model as wavespeed, but with a global random source. We can observe that the method converges extremely fast; after two iterations the exact and approximated solutions are indistinguishable to the naked eye. Even though most of the experiments shown in this manuscript were performed using randomly located point sources, which are of paramount importance in geophysics to approximate seismic sources, the behavior of the solver is agnostic to the form and location of the source as shown in Table 5.1, which presents the number of iterations to convergence when using a randomly located point source or a global source randomly generated by the `rand` function in MATLAB.

5.2. Convergence using different matrix splittings and Krylov solvers. Previously we defined two different preconditioners: the block Jacobi (3.10), and block Gauss–Seidel preconditioner (3.9), which are essentially different block matrix splittings of the polarized system in (3.8). We compare the performance of both preconditioners for two different Krylov subspace methods used to solve the polarized system. In particular, the polarized system is solved iteratively using either GMRES

¹⁷The BP 2004 model is a typical benchmark model in the geophysical community, and it includes simplified geologic features usually found in the Eastern/Central Gulf of Mexico and off-shore Angola.




FIG. 5.3. Two iterations of the preconditioner using a random source supported in the physical domain, from left to right and top to bottom: initial guess with only local solves, first iteration, second iteration, solution. The background model is given by the BP 2004 model [11] shown in Figure 5.1.

TABLE 5.2

Number of iterations required to reduce the relative residual to 10^{-7} , using GMRES and Bi-CGstab preconditioned with the Gauss–Seidel and Jacobi preconditioners. $f = \omega/2\pi \sim n$; the number of points in the PML scales as $\log(N)$, and the sound speed is given by the BP 2004 model (see [11]).

N	$\omega/2\pi$ [Hz]	$L \times L_c$	GS_{gmres}	J_{gmres}	$GS_{bicgstab}$	$J_{bicgstab}$
120×338	2.50	9×3	5	10	2.5	5
239×675	5.0	18×6	6	11	3	6
478×1349	10.0	27×9	6	12	3.5	6.5
955×2697	20.0	36×12	8	15	4	7
1910×5394	40.0	45×15	8	17	4.5	9.5

or Bi-CGstab [95] preconditioned with either the Gauss–Seidel preconditioner (3.9) or the Jacobi preconditioner (3.10). Table 5.2 summarizes the typical behavior of the different combinations of preconditioners and Krylov solvers; it shows the number of iterations needed to converge when the polarized system (3.5) is solved for different frequencies and number of subdomains using the BP 2004 model.

We can observe that for this particular model, the number of iterations needed to converge depends weakly on the frequency and the number of subdomains when a horizontally layered decomposition is used. Moreover, when the system is preconditioned with P^{GS} , the iterative solver converges twice as fast as when using P^{Jac} as a preconditioner, which can be loosely explained by the clustering of the eigenvalues of the preconditioned systems, as shown in Figure 3.2. Table 5.2 shows that the number of iterations for Bi-CGstab is half the number of iterations for GMRES; however, Bi-CGstab needs two applications of the matrix and preconditioner per iteration, resulting in a comparable computational cost, albeit with a smaller memory footprint for Bi-CGstab.

We point out that we used a zero initial guess for both GMRES and Bi-CGstab iterative solvers. Even though it is possible to use a random initial guess, it is not recommended given that the number of iterations will increase compared to a zero initial guess. Table 5.3 shows that the number of iterations to convergence using a random initial guess (generated using the `rand` command in MATLAB) retains the same asymptotic behavior as before, albeit with larger constants.

TABLE 5.3

Maximum number of iterations required to reduce the relative residual to 10^{-5} , using GMRES with the Gauss–Seidel preconditioner, for 50 realizations of a random source, using a zero initial guess and a randomly generated initial guess. In this case, $f = \omega/2\pi \sim n$, the number of layers and the number of cells inside each layer grow as n , the number of points in the PML scales as $\log(N)$, and the sound speed is given by the BP 2004 model (see [11]).

N	$\omega/2\pi$ [Hz]	$L \times L_c$	$u_0 = 0$	$u_0 = rand$
60×169	1.25	6×3	3	5
120×338	2.5	12×6	4	6
239×675	5.0	24×12	5	8
478×1349	10.0	48×24	6	10
955×2697	20.0	96×48	6	12

TABLE 5.4

Number of iterations required to reduce the relative residual to 10^{-5} , using GMRES preconditioned with the Gauss–Seidel preconditioner, in two different patches of the Marmousi2 model [71] using either a horizontal or vertical layered decomposition. The subscripts indicate in which medium the Helmholtz equation was solved: *lay* stands for the layered patch in Figure 5.4 (right); whereas *cent* stands for the centered patch in Figure 5.4 (left). $f = \omega/2\pi \sim n$, the number of layers scales as $L \sim n$; and the number of points in the PML scales as $\log(N)$.

N	$\omega/2\pi$	L	$Iter_{cent,vert}$	$Iter_{cent,hor}$	$Iter_{lay,vert}$	$Iter_{lay,hor}$
88×88	8	6	4	4	3	4
175×175	17	12	4	5	3	6
350×350	35	24	5	6	4	7
700×700	70	48	6	7	4	7
1400×1400	140	96	7	8	5	12
2800×2800	280	192	8	9	7	18

5.3. Layered media. Although the asymptotic behavior of the sweeping type solvers seems to be weakly dependent on the frequency, the number of subdomains, the source, and the initial guess, this kind of method can be very sensitive to the medium and to the orientation of the layered decomposition. One typical example of the former is a medium containing large resonant cavities, for which the number of iterations becomes highly frequency dependent. One example of the latter is layered media with multiscale structures at the wavelength level: if the domain decomposition is performed perpendicularly to the layers in the media (i.e., the layers in the decomposition are perpendicular to the layers present in the medium), the empirical behavior of the method remains the same; however, if the domain decomposition is performed parallel to the layers, there is a detriment on the convergence rate.

We showcase this behavior using two different parts of the Marmousi2 [71] model. We use the central part of the model, as in Figure 5.4 (left), in which we observe several layers that span different lengthscales, but without a preferred orientation, and the left-most part of the model, as in Figure 5.4 (right), in which most of the layers are oriented horizontally. Figure 5.5 depicts the solutions to the Helmholtz equation using both patches of the Marmousi2 model. Table 5.4 shows the number of iterations to convergence for each portion of the model, using a decomposition with either horizontal or vertical layers. When the medium does not possess a preferred orientation, both decompositions perform roughly the same; however, when there is a preferred orientation, the decomposition with layers perpendicular to the preferred orientation performs better. This difference in the performance is due to the increasing number of layers and frequency *simultaneously*; in fact, the number of iterations depends weakly on the frequency when the number of layers is constant, as shown in Table 5.5.




FIG. 5.4. *Truncated Marmousi models: left: the central part without a preferred layered structure, right: left-most part, with a horizontally layered structure.*




FIG. 5.5. *Solutions to the Marmousi2 patches for $\omega/(2\pi) = 35$: left: solution to the Helmholtz equation for the central patch of the Marmousi2 model, right: solution to the Helmholtz equation for the left-most patch of the Marmousi2 model.*

We interpret the detriment of the convergence rate as follows: from [103] it is known that the performance of the method of polarized traces is roughly proportional to energy scattered back at each sweep compared to the energy transmitted, which is closely linked to the number of reflections of the wavefield when interacting with the medium. In particular, the interactions within each subdomain are treated seamlessly using a direct solver, and interactions across subdomains are treated iteratively. The method handles the propagating waves in a sequential fashion by sweeping the waves traveling across the subdomains and treating the reflections in an ulterior sweep in the opposite direction. For example, if the medium is constant, then at each sweep most of the energy of the waves will be transmitted across subdomains, up to a small amount that is reflected given the imperfection of the absorbing boundary conditions. However, if the medium exhibits discontinuities in the wave speed, which will produce new reflections, then energy will be scattered back. These new reflections need to be propagated in an ulterior sweep in the opposite direction, which will provoke new reflections that need to be propagated in a subsequent sweep, and so on.

If the medium presents a handful of layers as in Figure 5.6, and the frequency and the number of subdomains increase, we have that most of the subdomains will not

TABLE 5.5

Number of iterations required to reduce the relative residual to 10^{-5} , using GMRES preconditioned with the Gauss–Seidel preconditioner, using the layered patch of the Marmousi2 model [71] as depicted in Figure 5.4 (right), and using either a horizontal or vertical layered decomposition. $f = \omega/2\pi \sim n$, the number of layers is fixed and equal to 6, and the number of points in the PML scales as $\log(N)$.

N	$\omega/2\pi$	L	$Iter_{\text{lay,vert}}$	$Iter_{\text{lay,hor}}$
88×88	8	6	3	4
175×175	17	6	3	5
350×350	35	6	4	6
700×700	70	6	4	6
1400×1400	140	6	4	6
2800×2800	280	6	5	6




FIG. 5.6. Layered media used for the numerical experiments with different contrasts; the left medium has a maximum contrast of 2.0 in the wave speed, and the right medium has a maximum contrast of 3.0.

present any interior reflections. Thus, during the sweeps the waves are transmitted seamlessly across most subdomains until they hit a subdomain containing a discontinuity, in which a reflection is created, which will need to be transmitted in an ulterior sweep. In this case, the number of reflections in the domain is mostly frequency independent. Therefore, provided that enough absorption is used for the transmission conditions between subdomains, the performance of the method is preserved as shown in Table 5.6.

It is well known that waves interact with structures that have a characteristic length of the same order than the wavelength. In the case of a layered medium with multiscale structures, as the one presented in Figure 5.4 (right), as the frequency increases, or the wavelength decreases, the waves will increasingly interact with increasingly thinner layers, thus increasing the amount of energy scattered back.

If the frequency increases but the number of layers is constant, then the increased number of interactions due to the increment on the frequency is handled with a direct method within each subdomain; the only interactions handled iteratively are the transmission and reflection between subdomains, which remain almost constant as the frequency increases, as depicted by Table 5.5.

However, if the number of domains increases as the frequency increases, then the orientation of the domain decomposition is crucial. If the orientation is perpendicular

TABLE 5.6

Number of iterations required to reduce the relative residual to 10^{-5} , using GMRES preconditioned with the Gauss–Seidel preconditioner, for the layered models in Figure 5.6, using a vertical domain decomposition, for two different contrasts. $f = \omega/2\pi \sim n$; the number of layers scales as $L \sim n$, and the number of points in the PML scales as $\log(N)$.

N	$\omega/2\pi$	L	$Iter_{2.0}$	$Iter_{3.0}$
60×60	6	4	2	2
120×120	12	8	2	2
240×240	24	16	2	3
480×480	48	32	3	4
960×960	96	64	3	5
1920×1920	192	128	5	7

to the layered media, then, as the frequency increases, the energy scattered from the micro structures is efficiently handled via a direct method within each subdomain. The interaction between subdomains will be composed of mostly transmitted waves propagating within each of the layers present in the medium. These interactions are efficiently handled by the sweeps. In such cases the method retains its performance as shown in Table 5.4. However, if the domain decomposition aligned to the layered media, then the method needs to properly handle all the internal reflection between the increasing numbers of subdomains in a sequential fashion, thus resulting in an increasing number of iterations to convergence as shown in Table 5.4.

Finally, the performance of the method is not only a function of the number of layers, it is a function of the strength of the reflections, which is directly linked to the contrast between layers. For example, in a layered medium with strong contrast, the amount of energy scattered back will increase, meaning that the sweep needs to take into account more reflections, as shown in Table 5.6.

We point out that in these numerical experiments we used $L \sim n$, which is a much more stringent condition than the one used in the nested version of the algorithm, in which $L \sim L_c \sim \sqrt{P}$, where $P = \mathcal{O}(N^{1/5})$. In order to keep an online complexity sublinear, the number of layers and cells in the decomposition grows slowly with respect to the frequency. This fact, coupled with the very efficient transmission conditions between cells, greatly attenuates the growth in the number of iterations for the nested solver within a layer, even in the case of multiscale layers with a preferred orientation.

5.4. Fast methods. Table 5.2 depicts the efficiency of the preconditioner measured in the number of iterations for convergence; however, in order to obtain sub-linear runtimes, we need to compress the integral kernels. As noted in [103] the scaling of the number of degrees of freedom with the frequency is critical to obtain the correct asymptotic compression rate. If the scaling is too aggressive, as in Table 5.2, the pollution error will be overwhelming and the compression rate of the integral operator will suffer. In order to account for the pollution error, in Table 5.7 we use the scaling $n \sim \omega^2$, which is known to be quasi-optimal for finite elements and finite differences (even though it is widely believed that $n \sim \omega^{3/2}$ is enough, cf. [5])

Table 5.7 shows the sublinear¹⁸ $\mathcal{O}(P^{1-\alpha}N^\alpha + P \log(N)^\alpha)$ scaling of the runtime of one GMRES iteration for $\alpha = 5/8$, as shown by Figure 5.7. Once again we can observe that the number of iterations to converge depends weakly on the frequency and the number of subdomains. Figure 5.7 depicts the efficient compression of the

¹⁸The same scaling holds for other typical geophysical benchmarks such as Marmousi2 [71], in which convergence is achieved in 4–6 iterations.

TABLE 5.7

Number of GMRES iterations (bold) required to reduce the relative residual to 10^{-5} , along with average execution time (in seconds) of one GMRES iteration using the compressed direct method, for different N and $P = L \times L_c$. The frequency is scaled such that $f = \omega/2\pi \sim \sqrt{n}$, the number of points in the PML scales as $\log(N)$, and the sound speed is given by the BP 2004 model (see [11]).

N	$\omega/2\pi$ [Hz]	6×2	24×8	42×14	60×20
120×338	2.50	(4) 0.42	(4) 8.30	(4) 24.8	(4) 51.7
239×675	3.56	(4) 0.74	(5) 9.15	(5) 26.1	(5) 52.8
478×1349	5.11	(4) 1.52	(5) 11.6	(5) 30.8	(5) 59.9
955×2697	7.25	(5) 3.32	(5) 17.9	(6) 38.5	(6) 68.8
1910×5394	10.3	(5) 6.79	(6) 29.6	(6) 58.7	(6) 98.3




FIG. 5.7. Runtime for one GMRES iteration using the two different nested solves, for $L = 9$ and $L_c = 3$, $\omega \sim \sqrt{n}$, in which the maximum ϵ -rank in the adaptive PLR compression (see section 5.1 in [103] for further details) scales as $\max_{\text{rank}} \sim \sqrt{\omega}$ and $\epsilon = 10^{-8}$. Moreover, for the nested polarized traces, the accuracy for the GMRES inner solve is fixed to 10^{-6} .

discrete integral operators; we can observe that both methods (nested polarized traces and compressed-block LU) have the same asymptotic runtime, but with different constants;¹⁹ the same scaling holds for different numbers of cells and layers.

6. Conclusion. We presented an extension to the method of polarized traces introduced in [103], with improved asymptotic runtimes in a distributed memory environment. The method has sublinear runtime even in the presence of rough media of geophysical interest. Moreover, its performance is completely agnostic to the source.

The method can be embedded efficiently within algorithms that require solving systems in which the medium is locally updated in an inversion loop. The method only needs to be locally modified in order to solve the updated system, thus reducing the overall computational effort. This algorithm is of special interest in the context of time-lapse full-waveform inversion, continuum reservoir monitoring, and local inversion.

We point out that this approach can be further parallelized using distributed linear algebra libraries. Moreover, it is possible to solve multiple right-hand sides simultaneously without an asymptotic penalty. This can be achieved by pipelining the sweeps, i.e., performing additional sweeps before the first one has finished, in order to maintain a constant load among all the nodes.

¹⁹We point out that some gains can be made by using different compressed operators. One can use one compressed operator with high accuracy to apply $\underline{\mathbf{M}}$, an operation that is easily parallelizable, and another with low accuracy to apply the preconditioner that represents the sequential bottleneck.

Appendix A. Discretization using finite differences.

In order to impose the absorbing boundary conditions, we extend the rectangular domain $\Omega = (0, L_x) \times (0, L_z)$ to $\Omega^{\text{ext}} = (-\delta_{\text{pml}}, L_x + \delta_{\text{pml}}) \times (-\delta_{\text{pml}}, L_z + \delta_{\text{pml}})$. The Helmholtz operator in (2.1) then takes the form

$$(A.1) \quad \mathcal{H} = -\partial_{xx} - \partial_{zz} - m\omega^2 \quad \text{in } \Omega^{\text{ext}},$$

where m is an extension²⁰ of the squared slowness. The differential operators are redefined following

$$(A.2) \quad \partial_x \rightarrow \alpha_x(\mathbf{x})\partial_x, \quad \partial_z \rightarrow \alpha_z(\mathbf{x})\partial_z,$$

where

$$(A.3) \quad \alpha_x(\mathbf{x}) = \frac{1}{1 + i \frac{\sigma_x(\mathbf{x})}{\omega}}, \quad \alpha_z(\mathbf{x}) = \frac{1}{1 + i \frac{\sigma_z(\mathbf{x})}{\omega}}.$$

Moreover, $\sigma_x(\mathbf{x})$ is defined as

$$(A.4) \quad \sigma_x(\mathbf{x}) = \begin{cases} \frac{C}{\delta_{\text{pml}}} \left(\frac{x}{\delta_{\text{pml}}} \right)^2 & \text{if } x \in (-\delta_{\text{pml}}, 0), \\ 0 & \text{if } x \in [0, L_x], \\ \frac{C}{\delta_{\text{pml}}} \left(\frac{x-L_x}{\delta_{\text{pml}}} \right)^2 & \text{if } x \in (L_x, L_x + \delta_{\text{pml}}), \end{cases}$$

and similarly for $\sigma_z(\mathbf{x})$.²¹ In general, δ_{pml} goes from a couple of wavelengths in a uniform medium to a large number independent of ω in a highly heterogeneous medium, and C is chosen to provide enough absorption.

With this notation we rewrite (2.1) as

$$(A.5) \quad \mathcal{H}u = f \quad \text{in } \Omega^{\text{ext}},$$

with homogeneous Dirichlet boundary conditions (f is the zero extended version of f to Ω^{ext}).

We discretize Ω as an equispaced regular grid of stepsize h and of dimensions $n_x \times n_z$. For the extended domain Ω^{ext} , we extend this grid by $n_{\text{pml}} = \delta_{\text{pml}}/h$ points in each direction, obtaining a grid of size $(2n_{\text{pml}} + n_x) \times (2n_{\text{pml}} + n_z)$. Define $\mathbf{x}_{p,q} = (x_p, z_q) = (ph, qh)$.

We use the 5-point stencil Laplacian to discretize (A.5). For the interior points $\mathbf{x}_{i,j} \in \Omega$, we have

$$(A.6) \quad (\mathbf{H}\mathbf{u})_{p,q} = -\frac{1}{h^2} (\mathbf{u}_{p-1,q} - 2\mathbf{u}_{p,q} + \mathbf{u}_{p+1,q}) - \frac{1}{h^2} (\mathbf{u}_{p,q-1} - 2\mathbf{u}_{p,q} + \mathbf{u}_{p,q+1}) - \omega^2 m(\mathbf{x}_{p,q}).$$

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

$$(A.7) \quad \alpha_x(\mathbf{x}_{p,q}) \frac{\alpha_x(\mathbf{x}_{p+1/2,q})(\mathbf{u}_{p+1,q} - \mathbf{u}_{p,q}) - \alpha_x(\mathbf{x}_{p-1/2,q})(\mathbf{u}_{p,q} - \mathbf{u}_{p-1,q})}{h^2},$$

²⁰We assume that $m(x)$ is given to us in Ω^{ext} .

²¹In practice, δ_{pml} and C can be seen as parameters to be tuned for accuracy versus efficiency.

and analogously for $\alpha_z \partial_z (\alpha_z \partial_z u)$. Although we use a second order finite difference stencil, the method can be easily extended to higher order finite differences. In such cases, the number of traces needed in the SIE reduction will increase.

Appendix B. Discretization using Q1 finite elements. The symmetric formulation of the Helmholtz equation takes the form

$$(B.1) \quad - \left(\nabla \cdot \Lambda \nabla + \frac{\omega^2 m(\mathbf{x})}{\alpha_x(\mathbf{x}) \alpha_z(\mathbf{x})} \right) u(\mathbf{x}) = \frac{f(\mathbf{x})}{\alpha_x(\mathbf{x}) \alpha_z(\mathbf{x})},$$

where

$$(B.2) \quad \Lambda(\mathbf{x}) = \begin{bmatrix} s_x(\mathbf{x}) & 0 \\ 0 & s_z(\mathbf{x}) \end{bmatrix},$$

$s_x = \alpha_x / \alpha_z$, $s_z = \alpha_z / \alpha_x$, where α_x and α_z are defined in (A.3).

In the case of a medium with sharp interfaces, finite difference approximations give inaccurate results due to the lack of differentiability of the velocity profile. In such cases, sophisticated quadratures and adaptive meshes have to be implemented to properly approximate the finite difference operator [100, 3]. We opted for a low order Q1 finite element discretization, with an adaptive quadrature rule at the discontinuities.

(B.1) is discretized using Q1 elements, leading to a discretized matrix

$$(B.3) \quad \mathbf{H} = \mathbf{S} - \mathbf{M},$$

where the stiffness matrix \mathbf{S} is computed using a Gauss quadrature. On the other hand, the mass matrix, \mathbf{M} , is computed using a quadrature adapted to each element depending on the local smoothness of the velocity profile:

- if the medium is locally smooth, a fixed Gauss quadrature is used to approximate the integral over the square;
- if the medium is discontinuous, an adaptive trapezoidal rule is used, until a preset accuracy is achieved.

To discriminate whether the medium is discontinuous, the velocity is sampled at the Gauss-points, and the ratio between maximum and minimum velocity is computed. If the ratio is smaller than a fixed threshold, the medium is considered smooth; otherwise it is considered discontinuous.

Using a nodal basis we can write the system to solve as

$$(B.4) \quad \mathbf{H}\mathbf{u} = \mathbf{f},$$

where \mathbf{u} is the pointwise value of the solution at the corners of the mesh, and \mathbf{f} is the projection of f onto the Q1 elements, using a high order quadrature rule.

Appendix C. Green's representation formula.

We present a generalization of the domain decomposition framework developed in [103] to Q1 regular finite elements.

We start by providing the algebraic formula for the discrete Green's representation formula. We propose a technique to derive such formulas without the time consuming computations performed in the Appendix of [103]. We point out that there are clear parallels between this derivation and the reduction to an interface problem using interior Schur complements. However, for the interface system based on Schur complements, we could not define the polarizing conditions that would allow

$$(C.6) \quad \mathcal{F}_1^\ell(\mathbf{u}) = -\mathbf{H}_{1,0}\mathbf{u}_0 = -\mathbf{H}_{1,0}^\ell\mathbf{u}_0.$$

We can observe that the role of \mathcal{F}^ℓ is to complete (C.3) at the boundary with exterior data, such that \mathbf{v}^ℓ satisfies the same equation as \mathbf{u} inside the whole layer and not only in the interior.

Analogously (C.3) can be evaluated at $k = 0$, obtaining

$$(C.7) \quad \mathbf{H}_{0,1}^\ell\mathbf{v}_1^\ell = \mathcal{F}_0^\ell(\mathbf{u}),$$

and imposing that $\mathbf{v}^\ell = \mathbf{u}\chi_{\Omega^\ell}$, we obtain that

$$(C.8) \quad \mathcal{F}_0^\ell(\mathbf{u}) = \mathbf{H}_{0,1}\mathbf{u}_1 = \mathbf{H}_{0,1}^\ell\mathbf{u}_1.$$

Finally, for $k < 0$, the same argument leads to

$$(C.9) \quad \mathcal{F}_k^\ell(\mathbf{u}) = 0.$$

We can easily generalize this argument for the other side of a layer, obtaining a generic formula for \mathcal{F}^ℓ

$$\mathcal{F}^\ell(\mathbf{u}) = -\delta_{n^\ell}\mathbf{H}_{n^\ell,n^\ell+1}^\ell\mathbf{u}_{n^\ell+1} + \delta_{n^\ell+1}\mathbf{H}_{n^\ell+1,n^\ell}^\ell\mathbf{u}_{n^\ell} - \delta_1\mathbf{H}_{1,0}^\ell\mathbf{u}_0 + \delta_0\mathbf{H}_{0,1}^\ell\mathbf{u}_1,$$

which can be substituted in (C.3), leading to

$$(C.10) \quad \begin{aligned} \mathbf{H}^\ell\mathbf{v}^\ell = & -\delta_{n^\ell}\mathbf{H}_{n^\ell,n^\ell+1}^\ell\mathbf{u}_{n^\ell+1} + \delta_{n^\ell+1}\mathbf{H}_{n^\ell+1,n^\ell}^\ell\mathbf{u}_{n^\ell} \\ & - \delta_1\mathbf{H}_{1,0}^\ell\mathbf{u}_0 + \delta_0\mathbf{H}_{0,1}^\ell\mathbf{u}_1 + \mathbf{f}^\ell. \end{aligned}$$

In addition, (C.10) can be transformed into the discrete expression of the Green’s representation formula by applying the inverse of \mathbf{H}^ℓ , \mathbf{G}^ℓ . We can then reformulate the Green’s integral in Definition 2.2 in the form

$$(C.11) \quad \mathcal{G}_j^{\downarrow,\ell}(\mathbf{v}_0, \mathbf{v}_1) = h [\mathbf{G}^\ell(z_j, z_1) \quad \mathbf{G}^\ell(z_j, z_0)] \begin{pmatrix} -\mathbf{H}_{1,0}^\ell\mathbf{v}_0 \\ \mathbf{H}_{0,1}^\ell\mathbf{v}_1 \end{pmatrix},$$

$$(C.12) \quad \mathcal{G}_j^{\uparrow,\ell}(\mathbf{v}_{n^\ell}, \mathbf{v}_{n^\ell+1}) = h [\mathbf{G}^\ell(z_j, z_{n^\ell+1}) \quad \mathbf{G}^\ell(z_j, z_{n^\ell})] \begin{pmatrix} \mathbf{H}_{n^\ell+1,n^\ell}^\ell\mathbf{v}_{n^\ell} \\ -\mathbf{H}_{n^\ell,n^\ell+1}^\ell\mathbf{v}_{n^\ell+1} \end{pmatrix}.$$

Finally, we can redefine $\mathbf{G}^\ell(z_j, z_k)$ for $k = 0, 1, n^\ell, n^\ell + 1$, such that they absorb all the extra factors. In particular, we redefine

$$(C.13) \quad \mathbf{G}_{1,0}^\ell = -\delta_1((\mathbf{H}^\ell)^{-1}\delta_0\mathbf{H}_{1,0}^\ell), \quad \mathbf{G}_{1,1}^\ell = -\delta_1((\mathbf{H}^\ell)^{-1}\delta_1\mathbf{H}_{0,1}^\ell),$$

$$(C.14) \quad \mathbf{G}_{n,n}^\ell = -\delta_n((\mathbf{H}^\ell)^{-1}\delta_{n+1}\mathbf{H}_{n+1,n}^\ell), \quad \mathbf{G}_{n,n+1}^\ell = -\delta_n((\mathbf{H}^\ell)^{-1}\delta_n\mathbf{H}_{n,n+1}^\ell).$$

This redefinition allows us to seamlessly use all the machinery introduced in [103] to define the SIE and its preconditioner.

Remark 1. As an example, in the case of the unsymmetric finite difference discretization, the upper and lower diagonal blocks of \mathbf{H} are diagonal matrices rescaled by $-1/h^2$. Then the formula presented here reduces exactly to the formulas computed by summation by parts in Appendix of [103].

Acknowledgments. Both authors thank Russell Hewett for interesting discussions.

REFERENCES

- [1] S. AMBIKASARAN, C. BORGES, L.-M. IMBERT-GERARD, AND L. GREENGARD, *Fast, Adaptive, High Order Accurate Discretization of the Lippmann-Schwinger Equation in Two Dimension*, preprint, arXiv:1505.07157 [math.NA], 2015, <https://arxiv.org/abs/1505.07157>.
- [2] P. R. AMESTOY, I. S. DUFF, J.-Y. L'EXCELLENT, AND J. KOSTER, *A fully asynchronous multifrontal solver using distributed dynamic scheduling*, SIAM J. Matrix Anal. Appl., 23 (2001), pp. 15–41, <https://doi.org/10.1137/S0895479899358194>.
- [3] I. BABUSKA, F. IHLENBURG, E. T. PAIK, AND S. A. SAUTER, *A generalized finite element method for solving the Helmholtz equation in two dimensions with minimal pollution*, Comput. Methods Appl. Mech. Engrg., 128 (1995), pp. 325–359, [https://doi.org/10.1016/0045-7825\(95\)00890-X](https://doi.org/10.1016/0045-7825(95)00890-X).
- [4] I. BABUSKA AND J. M. MELENK, *The partition of unity method*, Internat. J. Numer. Methods Engrg., 40 (1997), pp. 727–758, [https://doi.org/10.1002/\(SICI\)1097-0207\(19970228\)40:4\(727::AID-NME86\)3.0.CO;2-N](https://doi.org/10.1002/(SICI)1097-0207(19970228)40:4<727::AID-NME86>3.0.CO;2-N).
- [5] A. BAYLISS, C. I. GOLDSTEIN, AND T. E., *On accuracy conditions for the numerical computation of waves*, J. Comput. Phys., 59 (1985), pp. 396–404, [https://doi.org/10.1016/0021-9991\(85\)90119-6](https://doi.org/10.1016/0021-9991(85)90119-6).
- [6] M. BEBENDORF, *Hierarchical matrices: A means to efficiently solve elliptic boundary value problems*, Lect. Notes Comput. Sci. Eng. 63, Springer-Verlag, Berlin, 2008.
- [7] M. BEBENDORF AND W. HACKBUSCH, *Existence of \mathcal{H} -matrix approximants to the inverse fe-matrix of elliptic operators with l^∞ -coefficients*, Numer. Math., 95 (2002), pp. 1–28, <https://doi.org/10.1007/s00211-002-0445-6>.
- [8] R. BÉLANGER-RIOUX AND L. DEMANET, *Compressed Absorbing Boundary Conditions via Matrix Probing*, preprint, arXiv:1401.4421 [math.NA], 2014, <https://arxiv.org/abs/1401.4421>.
- [9] J.-P. BÉRENGER, *A perfectly matched layer for the absorption of electromagnetic waves*, J. Comput. Phys., 114 (1994), pp. 185–200.
- [10] A. BERMÚDEZ, L. HERVELLA-NIETO, A. PRIETO, AND R. RODRÍGUEZ, *Perfectly matched layers for time-harmonic second order elliptic problems*, Arch. Comput. Methods Eng., 17 (2010), pp. 77–107, <https://doi.org/10.1007/s11831-010-9041-6>.
- [11] F. BILLETTE AND S. BRANDSBERG-DAHL, *The 2004 BP velocity benchmark*, presented at the 67th annual European Association of Geoscientists and Engineers Conference and Exhibition, Madrid, 2005.
- [12] S. BORM, *Directional H2-Matrix compression for High-Frequency Problems*, preprint, arXiv:1510.07087 [math.NA], 2015, <https://arxiv.org/abs/1510.07087>.
- [13] Y. BOUBENDIR, *An analysis of the BEM-FEM non-overlapping domain decomposition method for a scattering problem*, J. Comput. Appl. Math., 204 (2007), pp. 282–291, <https://doi.org/10.1016/j.cam.2006.02.044>.
- [14] Y. BOUBENDIR, X. ANTOINE, AND C. GEUZAINÉ, *A quasi-optimal non-overlapping domain decomposition algorithm for the Helmholtz equation*, J. Comput Phys., 231 (2012), pp. 262–280, <https://doi.org/10.1016/j.jcp.2011.08.007>.
- [15] A. BRANDT AND I. LIVSHITS, *Wave-ray multigrid method for standing wave equations*, Electron. Trans. Numer. Anal., 6 (1997), pp. 162–181.
- [16] H. CALANDRA, S. GRATTON, X. PINEL, AND X. VASSEUR, *An improved two-grid preconditioner for the solution of three-dimensional Helmholtz problems in heterogeneous media*, Numer. Linear Algebra Appl., 20 (2013), pp. 663–688, <https://doi.org/10.1002/nla.1860>.
- [17] O. CESSENAT AND B. DESPRÉS, *Using plane waves as base functions for solving time harmonic equations with the ultra weak variational formulation*, J. Comput. Acoust., 11 (2003), pp. 227–238, <https://doi.org/10.1142/S0218396X03001912>.
- [18] T. F. CHAN AND T. P. MATHEW, *Domain decomposition algorithms*, Acta Numer., 3 (1994), pp. 61–143, <https://doi.org/10.1017/S0962492900002427>.
- [19] S. N. CHANDLER-WILDE, I. G. GRAHAM, S. LANGDON, AND E. A. SPENCE, *Numerical-asymptotic boundary integral methods in high-frequency acoustic scattering*, Acta Numer., 21 (2012), pp. 89–305, <https://doi.org/10.1017/S0962492912000037>.
- [20] Z. CHEN, M. J. GANDER, AND H. ZHANG, *On the relation between optimized Schwarz methods and source transfer*, in Domain Decomposition Methods in Science and Engineering XXII, T. Dickopf, M. J. Gander, L. Halpern, R. Krause, and L. F. Pavarino, eds., Springer, Cham, (2016), pp. 217–225, https://doi.org/10.1007/978-3-319-18827-0_20.

- [21] Z. CHEN AND X. XIANG, *A source transfer domain decomposition method for Helmholtz equations in unbounded domain*, SIAM J. Numer. Anal., 51 (2013), pp. 2331–2356, <https://doi.org/10.1137/130917144>.
- [22] Z. CHEN AND X. XIANG, *A source transfer domain decomposition method for Helmholtz equations in unbounded domain part II: Extensions*, Numer. Math. Theory Methods Appl., 6 (2013), pp. 538–555.
- [23] J. CHIU AND L. DEMANET, *Matrix probing and its conditioning*, SIAM J. Numer. Anal., 50 (2012), pp. 171–193, <https://doi.org/10.1137/110825972>.
- [24] X. CLAEYS AND R. HIPTMAIR, *Multi-trace boundary integral formulation for acoustic scattering by composite structures*, Comm. Pure Appl. Math., 66 (2013), pp. 1163–1201, <https://doi.org/10.1002/cpa.21462>.
- [25] F. COLLINO, S. GHANEMI, AND P. JOLY, *Domain decomposition method for harmonic wave propagation: A general presentation*, Comput. Methods Appl. Mech. Engrg., 184 (2000), pp. 171–211, [https://doi.org/10.1016/S0045-7825\(99\)00228-5](https://doi.org/10.1016/S0045-7825(99)00228-5).
- [26] S. COOLS AND W. VANROOSE, *Local Fourier analysis of the complex shifted Laplacian preconditioner for Helmholtz problems*, Numer. Linear Algebra Appl., 20 (2013), pp. 575–597, <https://doi.org/10.1002/nla.1881>.
- [27] T. A. DAVIS, *Algorithm 832: UMFPACK v4.3—an unsymmetric-pattern multifrontal method*, ACM Trans. Math. Software, 30 (2004), pp. 196–199, <https://doi.org/10.1145/992200.992206>.
- [28] A. DE LA BOURDONNAYE, C. FARHAT, A. MACEDO, F. MAGOULES, AND F.-X. ROUX, *A non-overlapping domain decomposition method for the exterior Helmholtz problem*, Contemp. Math., 218 (1998), pp. 42–66.
- [29] B. DESPRÉS, *Décomposition de domaine et problème de Helmholtz*, C.R. Math. Acad. Sci. Paris, 311 (1990), pp. 313–316.
- [30] Y. DU AND H. WU, *An Improved Pure Source Transfer Domain Decomposition Method for Helmholtz Equations in Unbounded Domain*, preprint, arXiv:1505.06052 [math.NA], 2015, <https://arxiv.org/abs/1505.06052>.
- [31] I. S. DUFF AND J. K. REID, *The multifrontal solution of indefinite sparse symmetric linear*, ACM Trans. Math. Software, 9 (1983), pp. 302–325, <https://doi.org/10.1145/356044.356047>.
- [32] B. ENGQUIST AND L. YING, *Sweeping preconditioner for the Helmholtz equation: Hierarchical matrix representation*, Comm. Pure Appl. Math., 64 (2011), pp. 697–735, <https://doi.org/10.1002/cpa.20358>.
- [33] B. ENGQUIST AND L. YING, *Sweeping preconditioner for the Helmholtz equation: moving perfectly matched layers*, Multiscale Model. Simul., 9 (2011), pp. 686–710.
- [34] B. ENGQUIST AND H.-K. ZHAO, *Absorbing boundary conditions for domain decomposition*, Appl. Numer. Math., 27 (1998), pp. 341–365, [https://doi.org/10.1016/S0168-9274\(98\)00019-1](https://doi.org/10.1016/S0168-9274(98)00019-1).
- [35] B. ENGQUIST AND H.-K. ZHAO, *Approximate Separability of Green’s Function for High Frequency Helmholtz Equations*, March 2014, <ftp://ftp.math.ucla.edu/pub/camreport/cam14-71.pdf>.
- [36] Y. A. ERLANGGA, *Advances in iterative methods and preconditioners for the Helmholtz equation*, Arch. Comput. Methods Eng., 15 (2008), pp. 37–66, <https://doi.org/10.1007/s11831-007-9013-7>.
- [37] Y. A. ERLANGGA, C. W. OOSTERLEE, AND C. VUIK, *A novel multigrid based preconditioner for heterogeneous Helmholtz problems*, SIAM J. Sci. Comput., 27 (2006), pp. 1471–1492, <https://doi.org/10.1137/040615195>.
- [38] O. G. ERNST AND M. J. GANDER, *Why it is difficult to solve Helmholtz problems with classical iterative methods*, in Numerical Analysis of Multiscale Problems, I. G. Graham, T. Y. Hou, O. Lakkis, and R. Scheichl, eds., Lect. Notes Comput. Sci. Eng. 83, Springer Berlin Heidelberg, 2012, pp. 325–363, https://doi.org/10.1007/978-3-642-22061-6_10.
- [39] C. FARHAT, I. HARARI, AND L. P. FRANCA, *The discontinuous enrichment method*, Comput. Methods Appl. Mech. Engrg., 190 (2001), pp. 6455–6479, [https://doi.org/10.1016/S0045-7825\(01\)00232-8](https://doi.org/10.1016/S0045-7825(01)00232-8).
- [40] M. GANDER, I. GRAHAM, AND E. 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), pp. 1–48, <https://doi.org/10.1007/s00211-015-0700-2>.
- [41] M. J. GANDER, *Optimized schwarz methods*, SIAM J. Numer. Anal., 44 (2006), pp. 699–731, <https://doi.org/10.1137/S0036142903425409>.
- [42] M. J. GANDER, *On the limitations of sweeping type preconditioners for wave propagation*, Oberwolfach Rep., 50 (2016), pp. 2899–2902.

- [43] M. J. GANDER AND F. KWOK, *Optimal interface conditions for an arbitrary decomposition into subdomains*, in Domain Decomposition Methods in Science and Engineering XIX, Y. Huang, R. Kornhuber, O. Widlund, and J. Xu, eds., Lect. Notes Comput. Sci. Eng. 78, Springer Berlin Heidelberg, 2011, pp. 101–108, https://doi.org/10.1007/978-3-642-11304-8_9.
- [44] M. J. GANDER, F. MAGOULÈS, AND F. NATAF, *Optimized Schwarz methods without overlap for the Helmholtz equation*, SIAM J. Sci. Comput., 24 (2002), pp. 38–60, <https://doi.org/10.1137/S1064827501387012>.
- [45] M. J. GANDER AND F. NATAF, *AILU for Helmholtz problems: A new preconditioner based on the analytic parabolic factorization*, J. Comput. Acoust., 9 (2001), pp. 1499–1506, <https://doi.org/10.1142/S0218396X01001510>.
- [46] M. J. GANDER AND Y. XU, *Optimized Schwarz method with two-sided transmission conditions in an unsymmetric domain decomposition*, in Domain Decomposition Methods in Science and Engineering XXII, T. Dickopf, M. J. Gander, L. Halpern, R. Krause, and L. F. Pavarino, eds., Springer, Cham, 2016, pp. 631–639, https://doi.org/10.1007/978-3-319-18827-0_65.
- [47] M. J. GANDER AND H. ZHANG, *Domain decomposition methods for the Helmholtz equation: A numerical investigation*, in Domain Decomposition Methods in Science and Engineering XX, R. Bank, M. Holst, O. Widlund, and J. Xu, eds., Lect. Notes Comput. Sci. Eng. 91, Springer Berlin Heidelberg, 2013, pp. 215–222, https://doi.org/10.1007/978-3-642-35275-1_24.
- [48] M. J. GANDER AND H. ZHANG, *Algorithmic perspective of PML transmission conditions for domain decomposition methods*, in 2014 IEEE Conference on Antenna Measurements Applications (CAMA), November 2014, pp. 1–4, <https://doi.org/10.1109/CAMA.2014.7003360>.
- [49] M. J. GANDER AND H. ZHANG, *Optimized Schwarz methods with overlap for the Helmholtz equation*, in Domain Decomposition Methods in Science and Engineering XXI, J. Erhel, M. J. Gander, L. Halpern, G. Pichot, T. Sassi, O. Widlund, eds., Springer, Cham, 2014, pp. 207–215, https://doi.org/10.1007/978-3-319-05789-7_17.
- [50] A. GEORGE, *Nested dissection of a regular finite element mesh*, SIAM J. Numer. Anal., 10 (1973), pp. 345–363.
- [51] S. GHANEMI, *A domain decomposition method for Helmholtz scattering problems*, in Proceedings of the Ninth International Conference on Domain Decomposition Methods, 1998, pp. 105–112.
- [52] A. GILLMAN, A. BARNETT, AND P. MARTINSSON, *A spectrally accurate direct solution technique for frequency-domain scattering problems with variable media*, BIT Mathematics, (2014), pp. 1–30, <https://doi.org/10.1007/s10543-014-0499-8>.
- [53] C. J. GITTELSON, R. HIPTMAIR, AND I. PERUGIA, *Plane wave discontinuous Galerkin methods: Analysis of the h-version*, ESAIM Math. Model. Numer. Anal., 43 (2009), pp. 297–331, <https://doi.org/10.1051/m2an/2009002>.
- [54] N. GMATI AND B. PHILIPPE, *Comments on the GMRES convergence for preconditioned systems*, in Large-Scale Scientific Computing: 6th International Conference, I. Lirkov, S. Margenov, J. Waśniewski, eds., Springer-Verlag, Berlin, 2008, pp. 40–51, https://doi.org/10.1007/978-3-540-78827-0_4.
- [55] I. G. GRAHAM, E. A. SPENCE, AND E. VAINIKKO, *Domain decomposition preconditioning for high-frequency Helmholtz problems using absorption*, preprint, arXiv:1507.02097 [math.NA], 2015, <https://arxiv.org/abs/1507.02097>.
- [56] E. HABER AND S. MACLACHLAN, *A fast method for the solution of the Helmholtz equation*, J. Comput. Phys., 230 (2011), pp. 4403–4418, <https://doi.org/10.1016/j.jcp.2011.01.015>.
- [57] R. HIPTMAIR AND C. JEREZ-HANCKES, *Multiple traces boundary integral formulation for Helmholtz transmission problems*, Adv. Comput. Math., 37 (2012), pp. 39–91, <https://doi.org/10.1007/s10444-011-9194-3>.
- [58] R. HIPTMAIR, A. MOIOLA, AND I. PERUGIA, *Plane wave discontinuous Galerkin methods for the 2d Helmholtz equation: Analysis of the p-version*, SIAM J. Numer. Anal., 49 (2011), pp. 264–284, <https://doi.org/10.1137/090761057>.
- [59] R. HIPTMAIR, A. MOIOLA, AND I. PERUGIA, *A survey of Trefftz methods for the Helmholtz equation*, preprint, arXiv:1506.04521 [math.NA], 2015, <https://arxiv.org/abs/1506.04521>.
- [60] J. XIA, S. CHANDRASEKARANM, M. GU, AND X. S. LI, *Superfast multifrontal method for large structured linear systems of equations*, SIAM J. Matrix Anal. Appl., 31 (2009), pp. 1382–1411, <https://doi.org/10.1137/09074543X>.
- [61] S. JOHNSON, *Notes on Perfectly Matched Layers (PMLs)*, manuscript, March 2010, <http://www-math.mit.edu/~stevenj/18.369/pml.pdf>.

- [62] R. KRESS, *Linear integral equations*, Appl. Math. Sci., Springer, Cham, 1999.
- [63] W. LENG, *A fast propagation method for the Helmholtz equation*, preprint, arXiv:1507.02467 [math.NA], 2015, <https://arxiv.org/abs/1507.02467>.
- [64] W. LENG AND L. JU, *An overlapping domain decomposition preconditioner for the Helmholtz equation*, preprint, arXiv:1508.02897 [math.NA], 2015, <https://arxiv.org/abs/1508.02897>.
- [65] Y. LI, H. YANG, E. MARTIN, K. HO, AND L. YING, *Butterfly factorization*, preprint, arXiv:1502.01379 [math.NA], 2015, <https://arxiv.org/abs/1502.01379>.
- [66] P.-L. LIONS, *On the Schwarz alternating method II*, in Domain Decomposition Methods, T. Chan, R. Glowinski, J. Periaux, and O. Widlund, eds., Lect. Notes Comput. Sci. Eng., 1989, pp. 47–70, https://doi.org/10.1007/978-3-642-35275-1_24.
- [67] F. LIU AND L. YING, *Additive sweeping preconditioner for the Helmholtz equation*, preprint, arXiv:1504.04058 [math.NA], 2015, <https://arxiv.org/abs/1504.04058>.
- [68] F. LIU AND L. YING, *Recursive sweeping preconditioner for the 3D Helmholtz equation*, preprint, arXiv:1502.07266 [math.NA], 2015, <https://arxiv.org/abs/1502.07266>.
- [69] S. LUO, J. QIAN, AND R. BURRIDGE, *Fast Huygens sweeping methods for Helmholtz equations in inhomogeneous media in the high frequency regime*, J. Comput. Phys., 270 (2014), pp. 378–401, <https://doi.org/10.1016/j.jcp.2014.03.066>.
- [70] F. MAGOULES, K. MEERBERGEN, AND J.-P. COYETTE, *Application of a domain decomposition with Lagrange multipliers to acoustic problems arising from the automotive industry*, J. Comput. Acoust., 8 (2000), pp. 503–521, <https://doi.org/10.1142/S0218396X00000297>.
- [71] G. MARTIN, R. WILEY, AND K. MARFURT, *Marmousi2: An elastic upgrade for Marmousi*, Leading Edge, 25 (2006), pp. 156–166.
- [72] P. MARTINSSON AND V. ROKHLIN, *A fast direct solver for scattering problems involving elongated structures*, J. Comput. Phys., 221 (2007), pp. 288–302, <https://doi.org/10.1016/j.jcp.2006.06.037>.
- [73] P. MARTINSSON, V. ROKHLIN, AND M. TYGERT, *A randomized algorithm for the decomposition of matrices*, Appl. Comput. Harmon. Anal., 30 (2011), pp. 47–68, <https://doi.org/10.1016/j.acha.2010.02.003>.
- [74] L. C. MCINNES, R. F. SUSAN-RESIGA, D. E. KEYES, AND H. M. ATASSI, *Additive schwarz methods with nonreflecting boundary conditions for the parallel computation of Helmholtz problems*, Contemp. Math., 218 (1998), pp. 325–333.
- [75] W. MCLEAN, *Strongly Elliptic Systems and Boundary Integral Equations*, Cambridge University Press, London, 2000.
- [76] A. MOIOLA, R. HIPTMAIR, AND I. PERUGIA, *Plane wave approximation of homogeneous Helmholtz solutions*, Z. Angew. Math. Phys., 62 (2011), pp. 809–837, <https://doi.org/10.1007/s00033-011-0147-y>.
- [77] A. MOIOLA AND E. A. SPENCE, *Is the Helmholtz equation really sign-indefinite?*, SIAM Rev., 56 (2014), pp. 274–312, <https://doi.org/10.1137/120901301>.
- [78] P. MONK AND D.-Q. WANG, *A least-squares method for the Helmholtz equation*, Methods Appl. Mech. Engrg., 175 (1999), pp. 121–136, [https://doi.org/10.1016/S0045-7825\(98\)00326-0](https://doi.org/10.1016/S0045-7825(98)00326-0).
- [79] R.-E. PLESSIX AND W. A. MULDER, *Separation-of-variables as a preconditioner for an iterative Helmholtz solver*, Appl. Numer. Math., 44 (2003), pp. 385–400, [https://doi.org/10.1016/S0168-9274\(02\)00165-4](https://doi.org/10.1016/S0168-9274(02)00165-4).
- [80] J. POULSON, B. ENGQUIST, S. LI, AND L. YING, *A parallel sweeping preconditioner for heterogeneous 3D Helmholtz equations*, SIAM J. Sci. Comput., 35 (2013), pp. C194–C212, <https://doi.org/10.1137/120871985>.
- [81] J. QIAN, S. LUO, AND R. BURRIDGE, *Fast Huygens’ sweeping methods for multiarrival Green’s functions of Helmholtz equations in the high-frequency regime*, Geophysics, 80 (2015), pp. T91–T100, <https://doi.org/10.1190/geo2014-0420.1>.
- [82] G. RIZZUTI AND W. MULDER, *Multigrid-based ‘shifted-laplacian’ preconditioning for the time-harmonic elastic wave equation*, J. Comput. Phys., 317 (2016), pp. 47–65, <https://doi.org/10.1016/j.jcp.2016.04.049>.
- [83] F.-H. ROUET, X. S. LI, P. GHYSELS, AND A. NAPOV, *A distributed-memory package for dense Hierarchically Semi-Separable matrix computations using randomization*, preprint, arXiv:1503.05464 [cs.MS], 2015, <https://arxiv.org/abs/1503.05464>.
- [84] Y. SAAD, *Iterative Methods for Sparse Linear Systems*, second ed., SIAM, 2003, <https://doi.org/10.1137/1.9780898718003>.
- [85] Y. SAAD AND M. H. SCHULTZ, *GMRES: A generalized minimal residual algorithm for solving nonsymmetric linear systems*, SIAM J. Sci. Stat. Comput., 7 (1986), pp. 856–869, <https://doi.org/10.1137/0907058>.
- [86] H. A. SCHWARZ, *Über einen grenzübergang durch alternierendes verfahren*, Vierteljahrsschrift der Naturforschenden Gesellschaft in Zurich, 15 (1870), pp. 272–286.

- [87] A. H. SHEIKH, D. LAHAYE, AND C. VUIK, *On the convergence of shifted Laplace preconditioner combined with multilevel deflation*, Numer. Linear Algebra Appl., 20 (2013), pp. 645–662, <https://doi.org/10.1002/nla.1882>.
- [88] C. STOLK, *A rapidly converging domain decomposition method for the Helmholtz equation*, J. Comput. Phys., 241 (2013), pp. 240–252.
- [89] C. C. STOLK, *An improved sweeping domain decomposition preconditioner for the Helmholtz equation*, Adv. Comput. Math., (2016), pp. 1–32, <https://doi.org/10.1007/s10444-016-9475-y>.
- [90] A. TOSELLI AND O. B. WINDLUND, *Domain Decomposition Methods – Algorithms and Theory*, Springer Ser. Comput. Math., Springer, Berlin, Heidelberg, 2005.
- [91] L. N. TREFETHEN AND D. BAU, *Numerical Linear Algebra*, SIAM, Philadelphia, 1997.
- [92] P. TSUJI, B. ENGQUIST, AND L. YING, *A sweeping preconditioner for time-harmonic Maxwell’s equations with finite elements*, J. Comput. Phys., 231 (2012), pp. 3770–3783, <https://doi.org/10.1016/j.jcp.2012.01.025>.
- [93] P. TSUJI, J. POULSON, B. ENGQUIST, AND L. YING, *Sweeping preconditioners for elastic wave propagation with spectral element methods*, ESAIM Math. Model. Numer. Anal., 48 (2014), pp. 433–447, <https://doi.org/10.1051/m2an/2013114>.
- [94] P. TSUJI AND L. YING, *A sweeping preconditioner for Yee’s finite difference approximation of time-harmonic Maxwell’s equations*, Front. Math. China, 7 (2012), pp. 347–363, <https://doi.org/10.1007/s11464-012-0191-8>.
- [95] 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 (1992), pp. 631–644, <https://doi.org/10.1137/0913035>.
- [96] A. VION, R. BÉLANGER-RIOUX, L. DEMANET, AND C. GEUZAINÉ, *A DDM double sweep preconditioner for the Helmholtz equation with matrix probing of the DtN map*, in Proceedings of the 11th International Conference on Mathematical and Numerical Aspects of Waves (WAVES 2013), Gammarth, Tunisia, 2013, pp. 95–96, <http://enitlr.tn/waves2013/proceedings.pdf>.
- [97] A. VION AND C. GEUZAINÉ, *Double sweep preconditioner for optimized Schwarz methods applied to the Helmholtz problem*, J. Comput. Phys., 266 (2014), pp. 171–190, <https://doi.org/10.1016/j.jcp.2014.02.015>.
- [98] A. VION AND C. GEUZAINÉ, *Parallel Double Sweep Preconditioner for the Optimized Schwarz Algorithm Applied to High Frequency Helmholtz and Maxwell Equations*, Springer International Publishing, Cham, 2016, pp. 239–247, https://doi.org/10.1007/978-3-319-18827-0_22.
- [99] S. WANG, X. S. LI, J. XIA, Y. SITU, AND M. V. DE HOOP, *Efficient scalable algorithms for solving dense linear systems with hierarchically semiseparable structures*, SIAM J. Sci. Comput., 35 (2013), pp. C519–C544, <https://doi.org/10.1137/110848062>.
- [100] X. WANG AND W. W. SYMES, *Harmonic coordinate finite element method for acoustic waves*, in SEG Technical Program Expanded Abstracts 2012, pp. 1–5, <https://doi.org/10.1190/segam2012-1151.1>.
- [101] L. ZEPEDA-NÚÑEZ, *Fast and Scalable Solvers for the Helmholtz Equation*, Ph.D. thesis, Massachusetts Institute of Technology, Cambridge MA, 2015.
- [102] L. ZEPEDA-NÚÑEZ AND L. DEMANET, *A short note on the nested-sweep polarized traces method for the 2D Helmholtz equation*, in SEG Technical Program Expanded Abstracts 2015, pp. 3682–3687, <https://doi.org/10.1190/segam2015-5838886.1>.
- [103] L. ZEPEDA-NÚÑEZ AND L. DEMANET, *The method of polarized traces for the 2D Helmholtz equation*, J. Comput. Phys., 308 (2016), pp. 347–388, <https://doi.org/10.1016/j.jcp.2015.11.040>.
- [104] L. ZEPEDA-NÚÑEZ, R. J. HEWETT, AND L. DEMANET, *Preconditioning the 2D Helmholtz equation with polarized traces*, in SEG Technical Program Expanded Abstracts 2014, 2014, pp. 3465–3470, <https://doi.org/10.1190/segam2014-1275.1>.