

# FAST ALTERNATING BI-DIRECTIONAL PRECONDITIONER FOR THE 2D HIGH-FREQUENCY LIPPMANN-SCHWINGER EQUATION

LEONARDO ZEPEDA-NÚÑEZ<sup>†</sup> AND HONGKAI ZHAO<sup>†</sup>

**Abstract.** This paper presents a fast iterative solver for Lippmann-Schwinger equation for high-frequency waves scattered by a smooth medium with a compactly supported inhomogeneity. The solver is based on the sparsifying preconditioner [63] and a domain decomposition approach similar to the method of polarized traces [64]. The iterative solver has two levels, the outer level in which a sparsifying preconditioner for the Lippmann-Schwinger equation is constructed, and the inner level, in which the resulting sparsified system is solved fast using an iterative solver preconditioned with a bi-directional matrix-free variant of the method of polarized traces. The complexity of the construction and application of the preconditioner is  $\mathcal{O}(N)$  and  $\mathcal{O}(N \log N)$  respectively, where  $N$  is the number of degrees of freedom. Numerical experiments in 2D indicate that the number of iterations in both levels depends weakly on the frequency resulting in a method with an overall  $\mathcal{O}(N \log N)$  complexity.

**1. Introduction.** In this paper we present a fast solver for the high-frequency Lippmann-Schwinger equation in 2D. Let  $\omega$  be the frequency of a time-harmonic wave propagating in a medium with wave-speed denoted by  $c(\mathbf{x})$ , where  $\mathbf{x} = (x, z)$ . Let  $\Omega$  be a rectangular domain and assume that the refraction index is constant up to a compactly supported perturbation. The refraction index is denoted by  $1/c^2(\mathbf{x}) = 1 + m(\mathbf{x})$ , where  $\text{supp}(m) = D \subset \Omega \subset \mathbb{R}^2$ . In this case,  $u(\mathbf{x})$ , the wave scattered by the perturbation  $m$  from an incident wave  $u_I(\mathbf{x})$ , is the solution to the Helmholtz equation

$$(1.1) \quad (\Delta + \omega^2(1 + m(\mathbf{x}))) (u(\mathbf{x}) + u_I(\mathbf{x})) = 0, \mathbf{x} \in \mathbb{R}^d;$$

where  $u(\mathbf{x})$  satisfies the Sommerfeld radiation condition

$$(1.2) \quad \lim_{r \rightarrow \infty} r^{1/2} \left( \frac{\partial u}{\partial r} - i\omega u \right) = 0,$$

and the incoming field  $u_I(\mathbf{x})$  satisfies the homogeneous Helmholtz equation

$$(1.3) \quad (\Delta + \omega^2) u_I(\mathbf{x}) = 0,$$

in some neighborhood of  $\Omega$  containing  $D$ .

From eqs. (1.1) and (1.3) it follows that the unknown field  $u$  satisfies

$$(1.4) \quad (\Delta + \omega^2(1 + m(\mathbf{x}))) u(\mathbf{x}) = -\omega^2 m(\mathbf{x}) u_I(\mathbf{x}), \quad \mathbf{x} \in \mathbb{R}^d.$$

Convolving (1.4) with

$$(1.5) \quad G(\mathbf{x}) = \frac{i}{4} H_0^1(\omega|\mathbf{x}|),$$

the Green's function for the 2D free space Helmholtz operator, results in the Lippmann-Schwinger equation:

$$(1.6) \quad u + \omega^2 G * (mu) = -\omega^2 G * (mu_I).$$

We chose (1.6) among different formulations of the Lippmann-Schwinger equation (for other formulations see [2]) by the following reasons :

---

<sup>†</sup>Department of Mathematics, University of California, Irvine, CA 92612, USA.

- the unknown scattered field,  $u(\mathbf{x})$ , satisfies the Sommerfeld radiation conditions automatically;
- the convolution kernel is translation invariant which allows for a rapid application via fast Fourier transforms (FFT) [21].

Even though  $u$  can be computed by discretizing (1.4) using finite differences, finite elements or composite spectral discretizations [38, 39]; and then solving the resulting linear system using state-of-the art preconditioners [18, 29, 30, 56, 61, 64], there are several advantages on computing  $u$  by solving (1.6) :

1. (1.6) arises from a second kind Fredholm formulation, implying that the conditioning number of the discretized system does not deteriorate as the discretization is refined;
2. (1.6) is defined on a bounded set  $\Omega$  and the solution satisfies the Sommerfeld radiation condition automatically, whereas (1.4) is defined in the whole  $\mathbb{R}^2$  thus; and hence one needs to truncate the computational domain, and to impose absorbing boundary conditions (ABC) at the boundary of the truncated computational domain, using perfectly matched layers (PML) [12, 13, 44] or other techniques [28] (see [14] for a review);
3. most discretizations of (1.4) via finite differences or finite elements suffer from dispersion error and the pollution effect [7, 8, 9] (see [57, 59] for recent references to methods designed to alleviate this issue); whereas the introduction of the analytical Green's function in (1.6) and the use of an accurate quadrature rule automatically alleviates these issues.

However, solving (1.6) raises other numerical issues. The discretization of (1.6) produces large dense linear systems, which becomes impractical, both memory- and complexity-wise, for typical direct solvers. Although a new generation of direct methods based on hierarchical compression, such as [20, 22, 42, 2], have been able to considerably reduce the complexity of solving integral equations coming from elliptic partial differential equations (PDE), when applying these methods to integral equations coming from wave equations such as (1.6), they manage to exhibit linear or quasi-linear complexity for low frequency problems, but they fall back to  $\mathcal{O}(N^{3/2})$  in high-frequency regimes. This phenomenon can be clearly explained by recent studies of the approximate separability of the Green's functions. It was shown in [11] that the Green's functions for elliptic operators are highly separable, implying that the off-diagonal blocks of the discrete Green's function are extremely low-rank. In contrast, it was shown in [32] that the Green's functions for the Helmholtz operator in high-frequency regime are highly non-separable in general. This last results provides an asymptotic frequency-dependent lower bound on the ranks of the off-diagonal blocks of the discrete Green's function, which hinders the compressibility of the blocks, thus resulting in algorithms with higher complexity.

In addition, the Lippmann-Schwinger equation becomes numerically-ill conditioned as the frequency and contrast increases, due to multiple scattering, hindering the convergence rate of iterative methods. This difficulty has led to the development of efficient preconditioners, such as [6, 16, 17, 55, 60]; however, they tend to require significantly more iterations to converge as the frequency increases. Some of them are based on semi-analytical techniques, which require fairly strict assumptions on the perturbation (for example [17]); they tend to work well under the assumptions, but they naturally deteriorate if not.

There is an extensive literature on volume integral methods, including fast methods [10, 41, 50] and, in particular, methods solving the Lippmann-Schwinger equation

[45, 20, 26, 63, 6, 17, 22, 60, 16, 1], and we do not review it here, but to the authors' knowledge there is no method that can handle general smooth perturbations in quasi-linear time in the high-frequency regime.

In the present paper we introduce a new preconditioner for the 2D high-frequency Lippmann-Schwinger equation, with a total asymptotic cost  $\mathcal{O}(N \log N)$ , which compares favorably with respect to most of the methods (both iterative and direct) in the literature. The scaling holds for general smooth perturbations without large resonant cavities. The solver is based on the sparsifying preconditioner and a fast iterative solver in a domain decomposition setting. In particular, the sparsifying preconditioner is used to approximate the Lippmann-Schwinger equation with a sparse system, which is solved fast using a bi-directional matrix-free variant of the method of polarized traces [66].

The method has two stages, the off-line, or setup, stage that is only performed once for each linear system to solve with an asymptotic cost of  $\mathcal{O}(N)$ , and the on-line, or solve, stage that is performed for each right-hand-side, with an asymptotic cost  $\mathcal{O}(N \log N)$ . Furthermore, the on-line stage has two levels: the inner level in which the approximate sparse system is solved iteratively using a bi-directional preconditioner, and the outer level, in which the resulting solution from the inner level is used to precondition the Lippmann-Schwinger equation.

Under the assumptions that the perturbation is smooth and it does not represent a resonant cavity the number of outer iterations is essentially independent of the frequency [63], and numerical experiments presented in this paper, indicate that the number of inner iterations depends, at worst, logarithmically with the frequency, resulting in the aforementioned complexity.

**1.1. Results.** In this paper we propose a new two-level preconditioner for the 2D high-frequency Lippmann-Schwinger equation that results in a method with  $\mathcal{O}(N \log N)$  complexity. We follow the approach in [63] to construct a sparsifying preconditioner, and we develop a fast iterative solver to solve the sparsified system, in order to achieve an overall optimal complexity. The main improvement with respect to the sparsifying preconditioner is how the resulting sparse linear system is solved. In [63] the system is solved via multi-frontal methods [37, 27] incurring an asymptotic cost  $\mathcal{O}(N^{3/2})$  for the setup and  $\mathcal{O}(N \log N)$  for the application. Exchanging the multifrontal solver by an iterative one, preconditioned using a bi-directional variant of the method of polarized traces, results in a drastic reduction of the asymptotic cost.

The inner-level preconditioner is a matrix-free variant of the method of polarized traces [66], which has been adapted to solve the sparsified Lippmann-Schwinger equation efficiently. As it will be explained in the sequel, preconditioning the sparsified system using the original method of polarized traces results in an algorithm that is sensitive to the sign of the perturbation and to the direction of the incident waves. Then a new bi-directional variant is introduced to eliminate these adverse effects.

The method of polarized traces is a domain decomposition method [64], which relies on a layered domain decomposition coupled with an equivalent surface integral equation (SIE) that is easy to precondition. The matrix-free variant introduced in [66] uses local solves on extended subdomains in order to apply the SIE and the preconditioner. The number of iterations to convergence depends on the underlying medium and the quality of the absorbing boundary conditions (ABC) imposed at the interfaces between subdomains. For the Helmholtz problem, there exist a myriad of different formulations of ABC [14, 44, 12, 28]; unfortunately, for the sparse system resulting from sparsifying the Lippmann-Schwinger equation such ABC are not readily avail-

able. We give details on building such ABC for the sparsified Lippmann-Schwinger system in Section 3.3.

Alas, the formulation of the ABC in this paper is not perfect, which is evidenced by numerical tests showing that the number of iteration needed to convergence using the matrix-free version of the method of polarized traces grows sub-optimally as  $\mathcal{O}(\sqrt{\omega})$ . Optimality can be restored by introducing an alternating bi-directional sweep, and by preconditioning the volume problem directly, which reduces the number of iterations to  $\mathcal{O}(\log \omega)$ . To show the effectiveness of our proposed method, we also include an example arising from plasma physics, in particular, an electromagnetic wave impinging on a confined plasma in a fusion reactor.

In the present work we mainly focus on explaining the key ideas and verification of the method by 2D tests. The algorithm presented in this paper can be easily extended to 3D, with an offline complexity of  $\mathcal{O}(N^{4/3})$  and an on-line complexity of  $\mathcal{O}(N \log^2 N)$ . The complexities can be further reduced using a nested approach such as the one presented in [49, 66] resulting in  $\mathcal{O}(N)$  and  $\mathcal{O}(N \log N)$  complexity respectively.

**1.2. Outline of the paper.** The present paper is organized as follows: we present the discretization of the Lippmann-Schwinger equation and provide a brief review of the sparsifying preconditioner in Section 2. In Section 3, we review the method of polarized traces and discuss the modifications needed to adapt it to the sparsified Lippmann-Schwinger equation, in particular we explain the formulation for the ABC in Section 3.3 and the bi-directional preconditioner in Section 3.5. In Section 4 we present the two-level preconditioner and discuss its complexity. Finally, in Section 5, we present various numerical experiments to support the complexity claims.

**2. Discretization and sparsifying preconditioner.** In this section we discuss the discretization used and provide a brief review of the sparsifying preconditioner [63]. We refer the interested reader to Section 2 of [63] for further details.

We discretize the domain  $\Omega$  with a uniform mesh<sup>1</sup> of a step size  $h$ , such that we have 10 points per wavelength in the background wavespeed,  $h\omega = \mathcal{O}(1)$ . We denote the collocation points by  $\mathbf{x}_{i,j} = (x_i, z_j) = (hi, hj)$ , and we define a vector index  $\mathbf{k} = (i, j)$ , notation that will be useful to explain the sparsifying preconditioner. We use a simple modification of the trapezoidal rule [26] to discretize the integral within the convolution operation in (1.6). In particular, we use a quadrature that only modifies the weight in the diagonal, obtaining a quadrature rule with order  $\mathcal{O}(h^4 \log h)$ .

We denote the unknown scattered field in  $\Omega$  as  $\mathbf{u}_{\mathbf{k}} = \mathbf{u}_{i,j} = u(\mathbf{x}_{i,j})$ , and we write the resulting discretized linear system as

$$(2.1) \quad \mathbf{H}\mathbf{u} = \mathbf{f},$$

in which  $\mathbf{H}$  can be written as

$$(2.2) \quad \mathbf{H} = \mathbf{I} + \omega^2 \mathbf{G} \cdot \mathbf{D},$$

---

<sup>1</sup>If the perturbation has sharp transitions the sparsifying preconditioner can be modified to handle adaptive quadratures without asymptotic penalty. However, the convolution product will have to be handled differently [15] and the construction of the preconditioner will require randomized techniques in order to keep the same asymptotic complexity.

where  $\mathbf{I}$  is the identity matrix, and  $\mathbf{D}$  is a diagonal matrix that encodes the multiplication of  $\mathbf{u}$  times  $m(\mathbf{x}_{i,j})$ . Finally,  $\mathbf{f}$  is given by

$$(2.3) \quad \mathbf{f} = -\omega^2 \mathbf{G} \cdot \mathbf{D} \mathbf{u}_I,$$

with  $(\mathbf{u}_I)_{i,j} = u_I(\mathbf{x}_{i,j})$ .

Given that the convolution kernel is translation invariant and that the mesh is uniform we can easily embed  $\mathbf{G}$  in a Toeplitz matrix and use the FFT to apply it in  $\mathcal{O}(N \log N)$  operations (or we can use the method in [15] if the mesh is non-uniform).

The sparsifying preconditioner is based on designing a sparse matrix  $\mathbf{A}$  such that the product  $\mathbf{A}\mathbf{H}$  is approximately sparse in the sense that a large amount of its entries are extremely small. The main idea is to multiply  $\mathbf{A}$  by  $\mathbf{H}$ , and threshold all the small entries of the product to zero, obtaining a sparse matrix  $\mathbf{C}$  that approximates  $\mathbf{A}\mathbf{H}$ . As it will be explained in the sequel, the sparsity pattern of  $\mathbf{A}$  is the same as a second-order compact finite difference discretization of the Laplacian, and its entries are specially constructed to annihilate the free space Green's function away from the source. In return, the matrix  $\mathbf{C}$ , which shares the same sparsity pattern as  $\mathbf{A}$ , is similar to the matrix arising from a finite differences discretization of the Helmholtz equation (1.1).

Since  $\mathbf{A}\mathbf{H} = \mathbf{A} + \omega^2 \mathbf{A}\mathbf{G}\mathbf{D}$ , it follows that to make  $\mathbf{A}\mathbf{H}$  approximately sparse, we only need to find  $\mathbf{A}$  such that the product  $\mathbf{A}\mathbf{G}$  is approximately sparse. The last requirement is equivalent to enforce that the entries of the matrix  $\mathbf{A}$  approximately annihilate the off-diagonal blocks of  $\mathbf{G}$ . This requirement can be recast as a minimization problem, which is solved via a singular value decomposition (SVD).

Define the set of immediate neighbors (stencils) of  $\mathbf{k}$  as

$$(2.4) \quad \mu(\mathbf{k}) = \{\mathbf{p} : \|\mathbf{k} - \mathbf{p}\|_{\ell_\infty(\mathbb{N}^2)} < 1, \mathbf{x}_{\mathbf{p}} \in \Omega\}.$$

We seek  $\mathbf{A}$  such that for each  $\mathbf{k}$ :

- only  $\mathbf{A}(\mathbf{k}, \mu(\mathbf{k}))$  are non-zeros (sparsity of  $\mathbf{A}$ );
- $\mathbf{A}(\mathbf{k}, :) \mathbf{G}(:, \mu(\mathbf{k})^c) \approx 0$  (approximate sparsity of the  $\mathbf{A}\mathbf{G}$ ).

Given that  $\mathbf{G}$  is translation invariant we impose  $\mathbf{A}$  to be translation invariant. In such case, it is sufficient to consider only the case when  $\mathbf{k} = 0$ . Then the problem can be recast as

$$(2.5) \quad \min_{\alpha: \|\alpha\|=1} \|\alpha \cdot \mathbf{G}(\mu(0), \mu(0)^c)\|,$$

which can be efficiently solved via a singular value decomposition of  $\mathbf{G}(\mu(0), \mu(0)^c) = U\Sigma V^*$  such that

$$(2.6) \quad \alpha = U(:, \mu(0))^*.$$

Finally,  $\mathbf{A}$  is built such that  $\mathbf{A}(\mathbf{k}, \mu(\mathbf{k})) = \alpha$ , for  $\mathbf{k}$  in the interior of  $\Omega$ . For  $\mathbf{k}$  on  $\partial\Omega$  a special treatment is needed, which depends on the geometry of  $\Omega$ . For the sake of simplicity we omit such treatment but we refer the interested reader to Section 4 of [63].

By multiplying (2.1) by  $\mathbf{A}$  and sparsifying the product, we obtain the approximated *sparse* system

$$(2.7) \quad \mathbf{C}\mathbf{u} \approx \mathbf{A}\mathbf{f},$$

for which the preconditioner is defined as

$$(2.8) \quad \mathbf{u}_{approx} = \mathbf{C}^{-1}\mathbf{g}, \quad \text{where } \mathbf{g} = \mathbf{A}\mathbf{f}.$$

It has been shown in [63] that when the preconditioner mentioned above is used within GMRES iterations, the number of iterations to convergence is almost independent of the frequency.

One alternative to this preconditioner would be to discretize (1.4) via finite differences or finite elements, and then use this approximate solution to precondition (2.1). Unfortunately, such preconditioners tend to behave poorly, due to the pollution error that is normally introduced with the discretizations. The pollution error creates serious phase errors making the preconditioner incompatible with the linear system being preconditioned (see [57] for a similar example for multigrid). The optimization procedure used in the construction of the sparsifying preconditioner takes into account the exact free space Green's function and hence minimizes the dispersion error. The size of the stencils used in  $\mathbf{A}$  and  $\mathbf{C}$  should strike a balance between sparsity and dispersion error. Since the matrix  $\mathbf{A}$  mimics the discretization of the Helmholtz differential operator, it is expected that the ideas for fast solver for the Helmholtz equation can be applied to the resulting sparse system. This is what we propose in the next Section.

**3. Method of polarized traces.** As seen from Section 2, the sparsifying preconditioner for the Lippmann-Schwinger equation relies on solving the sparsified system

$$(3.1) \quad \mathbf{C}\mathbf{v} = \mathbf{g},$$

which, in the present paper, is solved fast using a bi-directional variant of the method of polarized traces. In this section, we first briefly introduce the matrix-free version of the method of polarized traces [66] in one direction, and then we discuss the modifications needed to adapt it to the sparsified Lippmann-Schwinger equation. We provide a simple description of the three key elements in the method of polarized traces for the Helmholtz equation as follows

1. the layered domain decomposition with the extended subdomains;
2. how to define the transmission conditions between subdomains using a discrete Green's representation formula issued from an algebraic decomposition;
3. how to implement an absorbing boundary conditions for the subdomains to obtain a fast converging iterative method.

We then provide the full algorithm for the one-directional preconditioner; we provide the physical intuition and discuss its short-comings, which are rectified by the bi-directional preconditioner that is introduced right after.

**3.1. Domain decomposition.** The domain  $\Omega$  is partitioned in  $L$  horizontal layers  $\{\Omega^\ell\}_{\ell=1}^L$ . Following [64] we extend each subdomain, which we denote by  $\{\Omega_{ext}^\ell\}_{\ell=1}^L$ . Within each extended subdomain we assemble the discrete local problem

$$(3.2) \quad \mathbf{C}^\ell \mathbf{w}^\ell = \mathbf{g}^\ell = \chi_{\Omega^\ell} \mathbf{g}, \quad \text{for } \ell = 1, \dots, L.$$

We impose  $\mathbf{C}^\ell$  to satisfy the *compatibility condition*, i.e.,  $\mathbf{C}^\ell$  coincides exactly with  $\mathbf{C}$  in the interior of  $\Omega^\ell$ . Within  $\Omega_{ext}^\ell - \Omega^\ell$ , i.e., the extension of the subdomain, we do not require  $\mathbf{C}^\ell$  to coincide with  $\mathbf{C}$ , instead we use the extra degrees of freedom to impose an analogue to an absorbing boundary condition. Given that  $\mathbf{C}$  in the



in which all the sub-matrices are sparse. Given that the mesh is uniform we use the following ordering of the unknowns

$$(3.4) \quad \mathbf{v} = (\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_{n_z}),$$

where  $\mathbf{v}_j \in \mathbb{C}^{n_x}$  is given by

$$(3.5) \quad \mathbf{v}_j = (v_{1,j}, v_{2,j}, \dots, v_{n_x,j}),$$

which corresponds to  $\mathbf{v}$  sampled at a constant  $z$ .

We write  $\mathbf{v}^\ell$  for the wavefield defined locally at the  $\ell$ -th layer, i.e.,  $\mathbf{v}^\ell = \chi_{\Omega^\ell} \mathbf{v}$ , and  $\mathbf{v}_k^\ell$  for the values at the local depth<sup>2</sup>  $z_k^\ell$  of  $\mathbf{v}^\ell$ .

In particular,  $\mathbf{v}_1^\ell$  represents the wavefield sampled at the bottom of the layer, and  $\mathbf{v}_{n^\ell}^\ell$  is the wavefield at the top of the layer. Moreover,  $\mathbf{v}_0^\ell$  is the local wavefield sampled immediately below the bottom of the layer, in the extended region within the PML;  $\mathbf{v}_{n^\ell+1}^\ell$  is the local wavefield sampled immediately above the top of the layer.

Finally, we use the notation  $\delta_k$  to denote a discrete measure. For a generic vector  $\mathbf{p} \in \mathbb{C}^{n_x}$ ,  $\delta_k \mathbf{p} \in \mathbb{C}^{n_x \times n_z}$  is defined as follows:

$$(3.6) \quad (\delta_k \mathbf{p})_{i,j} = \begin{cases} \mathbf{p}_i & \text{if } j = k, \\ 0 & \text{if } j \neq k. \end{cases}$$

This notation will be used extensively to characterize the forcing terms necessary to transfer the information from one subdomain to its neighbors.

The correct Green's representation formula is given by the equivalent and more manageable expression<sup>3</sup>

$$(3.7) \quad \begin{aligned} \mathbf{C}^\ell \mathbf{w}^\ell &= -\delta_{n^\ell} \mathbf{C}_{n^\ell, n^\ell+1}^\ell \mathbf{v}_{n^\ell+1} + \delta_{n^\ell+1} \mathbf{C}_{n^\ell+1, n^\ell}^\ell \mathbf{v}_{n^\ell} \\ &\quad - \delta_1 \mathbf{C}_{1,0}^\ell \mathbf{v}_0 + \delta_0 \mathbf{C}_{0,1}^\ell \mathbf{v}_1 + \mathbf{g}^\ell. \end{aligned}$$

It can be shown that if the traces of the global  $\mathbf{v}$ , i.e.  $\mathbf{C}\mathbf{v} = \mathbf{f}$ , are known and used in (3.7), then  $\mathbf{w}^\ell$  coincides exactly with  $\mathbf{v}$  inside the layer, i.e.,  $\mathbf{w} = \chi_{\Omega^\ell} \mathbf{v}$  (see Appendix C of [64] for a rigorous proof for matrices arising from finite differences). In [66] the above formula is used to build an equivalent surface integral formulation in a matrix-free fashion, which is then solved iteratively. The convergence is fast provided that high-quality ABC are implemented, which is the topic of the next subsection.

**3.3. Absorbing boundary conditions.** One of most the important elements of the method of polarized traces is the implementation of ABC at the interfaces between subdomains. ABC have successfully been used in domain decomposition for elliptic equations [31], and a simpler version of ABC were extensively used in the first domain decomposition method for the Helmholtz equation to ensure convergence [25]. High-quality absorbing conditions between subdomains are key to the fast convergence of iterative solvers for the high-frequency Helmholtz equation [56, 61, 30, 48, 18, 19, 46, 49, 58, 47].

The high-quality ABC are needed in order to minimize the artificial reverberations at the interfaces between subdomains that arise from the truncation introduced

<sup>2</sup>We hope that there is little risk of confusion in overloading  $z_j$  (local indexing) for  $z_{n^\ell+j}$  (global indexing), where  $n_c^\ell = \sum_{j=1}^{\ell-1} n^j$  is the cumulative number of points in depth.

<sup>3</sup>It is possible to recover the Green's representation formula by multiplying (3.7) by the inverse of  $\mathbf{C}^\ell$



by the domain decomposition. One of the objectives of the ABC, is to make the local Green's functions, which are defined by the local problem (3.2), as close as possible to the global Green's function restricted to the subdomain without scattering outside the subdomain. In other words, the local Green's functions should capture the local interactions within the subdomain, but should not contain pollution effects due to artificial reflections at the truncated boundary. Given that medium outside the subdomain is not involved in the local system, the local Green's functions do not capture those long range interactions among subdomains; these interactions will be captured in an iterative manner during the sweeps.

In practice, without careful treatment at the truncated boundaries, the artificial reverberations can heavily pollute the local Green functions and hinder the convergence of the iterative method. For example, the adverse effect of the artificial reverberations can be overwhelming for the cases when Dirichlet or Neumann boundary conditions are imposed at the boundaries between subdomains [61, 33], causing the method to converge in  $\mathcal{O}(\omega)$  iterations, thus resulting in a super-linear complexity.

Given that the matrix  $\mathbf{C}$  in our problem does not arise from a direct discretization of the Helmholtz equation, absorbing boundary conditions are not readily available. A natural way is to localize the perturbation,  $m(\mathbf{x})$ , to the subdomain via a smooth cut-off. The localization produces a local discrete system in the form of (3.2) that satisfies the compatibility conditions.

Now, we proceed to give an explanation of how to define the local problems (3.2) using the localization to obtain an approximated ABC. Intuitively, the solution to (3.2) should provide an approximation of the wave scattered by the perturbation supported in  $\Omega^\ell$ . Thus, one naive approach would be to define the local problem as

$$(3.8) \quad \mathcal{H}^\ell u^\ell := u^\ell + \omega^2 G * (\chi_{\Omega^\ell} m u^\ell),$$

and then use the sparsifying preconditioner to obtain a sparse local problem in the form of (3.2).

However, the local perturbation  $\chi_{\Omega^\ell} m$  is discontinuous at the interfaces between subdomains, which will produce strong artificial reflections for waves impinging on the artificial boundary of  $\Omega^\ell$ . These strong reflections will pollute the local Green's function and result in a slow convergence. Fortunately, this adverse effect can be easily rectified by using a window, or cut-off, function that transitions the perturbation smoothly to zero.

Define a set of smooth cut-off functions  $\{\xi^\ell(\mathbf{x})\}_{\ell=1}^L$ , such that

$$(3.9) \quad \xi^\ell(\mathbf{x}) = \begin{cases} 1 & , \text{if } \mathbf{x} \in \Omega^\ell; \\ 0 & , \text{if } \mathbf{x} \notin \Omega_{ext}^\ell. \end{cases}$$

We can define the continuous counterpart of (3.2) as the problem

$$(3.10) \quad \mathcal{H}^\ell u^\ell := u^\ell + \omega^2 G * (\xi^\ell m u^\ell).$$

which is posed for  $\mathbf{x} \in \Omega$ . We recall that the matrix  $\mathbf{D}$ , is a diagonal matrix that encodes the multiplication by  $m(\mathbf{x}_{i,j})$ ; we define  $\mathbf{D}^\ell$  as the diagonal matrix that represents the multiplication by  $\xi^\ell(\mathbf{x})m(\mathbf{x})$ . We can define the discrete local Lippmann-Schwinger equation as,

$$(3.11) \quad \mathbf{H}_{global}^\ell = \mathbf{I} + \omega^2 \mathbf{G} \cdot \mathbf{D}^\ell.$$

The global subscript indicates that  $\mathbf{H}_{global}^\ell$  is a  $N \times N$  matrix. In this case we can multiply (3.11) by  $\mathbf{A}$ , and we would obtain a local system that satisfies the

compatibility condition of (3.2). However, this would be prohibitively expensive given that would need to solve a  $N \times N$  system for each subdomain.

The key observation is that only a small portion of  $\mathbf{H}_{global}^\ell$  in (3.11) deviates from an identity matrix. This is caused by the small support of  $\xi^\ell m$ , in which  $\text{supp}(\xi^\ell m) \subset \Omega_{ext}^\ell$ . Thus, we can re-write (3.10) in  $\Omega_{ext}^\ell$  instead of  $\Omega$ , and discretize it resulting in

$$(3.12) \quad \mathbf{H}^\ell = \mathbf{I} + \omega^2 \mathbf{G} \cdot \mathbf{D}^\ell,$$

which is a quasi-1D (see [30])  $qn \times qn$  system, where  $q$  is the thickness, in grid points, of each subdomain. In the implementation of the method, the cut-off is constant in the direction tangential to the boundary,  $x$  direction in our layered setting, and it is a third order spline in the direction normal to the boundary,  $z$  direction in our setting.

In theory, we could apply the sparsifying preconditioner to (3.12); however, given that the mesh is also truncated, the optimization via a SVD can provide different results, hindering the compatibility condition. Instead, we multiply (3.12) by the corresponding block of  $\mathbf{A}$ , but with a special care at the boundaries of  $\Omega_{ext}^\ell$ , where we use the stencil generated at the boundary of  $\Omega$ , which provides a satisfying ABC [63] for the background wave-speed at the outer boundary of the extended domain.

Unfortunately, the smooth cut-off used to localize the perturbation is not equivalent to an ABC, which would approximate the Dirichlet to Neumann (DtN) map at the boundary; such that it can absorb and dampen waves propagating out of the boundary in all directions with little reflections. The smooth cut-off in the normal direction to the boundary minimally affects waves propagating out of the domain close to the normal direction. However, for waves propagating near and almost tangential to the boundary, a smooth cut-off may not be good enough depending on the sign of the perturbation as illustrated in Figures 2 and 3. If the perturbation is negative, the waves propagate faster inside the subdomain than in the cut-off region. If a wave propagates almost tangentially to the boundary to the cut-off region it will bend outwards, and it will be ultimately absorbed by the numerical ABC situated at the artificial boundary of the extended subdomain. On the other hand, if the perturbation is positive, then the waves propagate slower inside the subdomain than in the background media. Under these circumstances, the subdomain with a smooth cut-off to the background media creates an artificial wave-guide effect at the subdomain boundary. Waves propagating almost normal to the boundaries will still cross the smooth cut-off and get absorbed by the numerical ABC at the artificial boundary of the extended subdomain. However, waves propagating almost transversally along the subdomain boundary will bend inwards due to the artificial cut-off, producing spurious grazing waves that pollute the local solution within the subdomain. This last problem can be attenuated to some extent by introducing a  $\mathcal{O}(\omega)$  complex shift within the cut-off region, thus dampening the grazing waves. We point out that we are only introducing a complex shift in the extended domain; it would be possible to introduce a complex shift in  $\Omega$  as in [40], however, that would violate the compatibility condition.

Although the complex shift helps to reduce the number of iterations for the original method of polarized traces from  $\mathcal{O}(\omega)$  to  $\mathcal{O}(\sqrt{\omega})$ , which is better than standard iterative methods, it is not optimal. In order to obtain the lower rate of  $\mathcal{O}(\log \omega)$  iterations to converge, we introduce in the next section an effective bi-directional alternating sweeping preconditioner to handle those grazing waves due to either the incoming waves or the smooth cut-off at the artificial boundaries between two subdomains in one directional sweeping.




FIG. 2. Example of a global (top) and local (bottom) Green's function for a positive perturbation, we can observe that there is some interference in the local Green's function given by the spurious grazing wave that is concentrated at the top of the slab.




FIG. 3. Example of a global (top) and local (bottom) Green's function for a negative perturbation, we can observe some differences between the Green's functions, but we do not observe the kind of interference, nor the concentration at the boundaries, as observed in Fig. 2.

**3.4. One-directional algorithm.** In this section we discuss the preconditioner for the sparse problem based on the method of polarized traces for one-directional sweep and provide the algorithm in pseudo-code with a physical interpretation.

The main intuition behind the method of polarized traces is the directionality of the waves, which are imposed by the polarizing conditions (see Section 3 of [64]) and the absorbing boundary conditions. For example, the wavefield generated by a local source,  $\mathbf{g}^\ell$

$$(3.13) \quad \mathbf{C}^\ell \mathbf{v}^\ell = \mathbf{g}^\ell,$$

irradiates from the interior of  $\Omega^\ell$  to the exterior of it. If we sample  $\mathbf{v}^\ell$  at the bottom of  $\Omega^\ell$ , we can physically interpret it as the trace of a wavefield propagating downwards.

Following this physical interpretation, the wavefield propagating downwards from  $\Omega^\ell$  propagates inwards  $\Omega^{\ell-1}$ . By using the Green's representation formula we can propagate the wavefield from  $\Omega^\ell$  to  $\Omega^{\ell-1}$  by artificial sources located at the top of  $\Omega^{\ell-1}$  following

$$(3.14) \quad \mathbf{C}^{\ell-1} \mathbf{v}^{\ell-1} = -\delta_{n^{\ell-1}} \mathbf{C}_{n^{\ell-1}, n^{\ell-1}+1}^{\ell-1} \mathbf{v}_1^\ell + \delta_{n^{\ell-1}+1} \mathbf{C}_{n^{\ell-1}+1, n^{\ell-1}}^{\ell-1} \mathbf{v}_0^\ell + \mathbf{g}^{\ell-1},$$

which is an alternate expression for the incomplete Green's representation formula introduced in [64].

We can observe that  $\mathbf{v}^{\ell-1}$  contains the energy radiating from  $\mathbf{g}^{\ell-1}$  and  $\mathbf{g}^\ell$ . If we sample  $\mathbf{v}^{\ell-1}$  at the bottom of  $\Omega^{\ell-1}$ , it can be interpreted as the downwards wavefield consisting of the downwards wavefield generated by  $\mathbf{g}^{\ell-1}$  and the wavefield generated by  $\mathbf{g}^\ell$  that propagated through  $\Omega^{\ell-1}$ .

We can continue the propagation by adding the local wavefields in a manner that is usually called a multiplicative Schwartz iteration (see [36]). Analogously, the same procedure can be performed for the wavefields propagating upwards. The application of both sweeps, downwards and upwards can be easily understood as a block Jacobi iteration as presented in [64]. Furthermore, the polarizing conditions can be further exploited to reduce the number of iterations by using a Gauss-Seidel approach described by Alg. 1, which in practical terms reduces the number of iterations by a factor of two (see [66]).

All the intuitive arguments provided in the current exposition can be made rigorous from an algebraic point of view, we skip such explanation and redirect the interested readers to [64].

ALGORITHM 1. *Gauss-Seidel Preconditioner*

```

1: function  $\mathbf{v} = \text{PRECONDITIONER}(\mathbf{g})$ 
2:   for  $\ell = 1 : L$  do                                      $\triangleright$  Downwards Sweep
3:      $\mathbf{g}^\ell = \mathbf{g}_{\chi_{\Omega^\ell}}$                                 $\triangleright$  Partition of the Source
4:     if  $\ell \neq 1$  then
5:        $\mathbf{g}^\ell_+ = -\delta_{n^\ell} \mathbf{C}_{n^\ell, n^{\ell+1}}^\ell \mathbf{w}_1^{\ell-1} + \delta_{n^{\ell+1}} \mathbf{C}_{n^{\ell+1}, n^\ell}^\ell \mathbf{w}_0^{\ell-1}$ 
6:     end if
7:      $\mathbf{w}^\ell = (\mathbf{C}^\ell)^{-1} \mathbf{g}^\ell$                                 $\triangleright$  Local Solve
8:   end for
9:   for  $\ell = L : -1 : 1$  do                                    $\triangleright$  Upwards Sweep
10:    if  $\ell \neq L$  then
11:       $\mathbf{g}^\ell_+ = -\delta_1 \mathbf{C}_{1,0}^\ell \mathbf{v}_n^{\ell+1} + \delta_0 \mathbf{C}_{0,1}^\ell \mathbf{v}_n^{\ell+1}$ 
12:    end if
13:     $\mathbf{v}^\ell = (\mathbf{C}^\ell)^{-1} \mathbf{g}^\ell$                                 $\triangleright$  Local Solve
14:  end for
15:   $\mathbf{v} = (\mathbf{v}^1, \mathbf{v}^2, \mathbf{v}^3, \dots, \mathbf{v}^{L-1}, \mathbf{v}^L)$             $\triangleright$  Concatenation
16: end function

```

**3.5. Bi-directional algorithm.** In this Section we first discuss briefly the shortcomings of Alg. 1 when applied to the sparsified Lippmann-Schwinger equation and then introduce the bi-directional preconditioner in order to achieve the optimal complexity.

The algorithm as presented in Alg. 1, may not work optimally in our setting for the Lippmann-Schwinger equation, in which the perturbation is localized using a smooth cut-off. The main issue is that waves propagating nearly tangential to the artificial boundaries between subdomains are perturbed due to the cut-off and depending on the perturbation they can be bended inwards and thus generate spurious grazing waves. Even though the spurious grazing waves are dampened by the introduction of a complex shift in the cut-off region to some extent, they can still cause significant deviations of the local Green's functions with respect to the global ones. Thus, they hinder the fast convergence of the iterations, especially when the perturbation is positive, i.e., the wave-guide situation, as explained in Section 3.3 and shown in Figures 2 and 3.

Moreover, a closer inspection to Alg. 1 reveals that within the algorithm pipeline, the residual given by  $\mathbf{r} = \mathbf{C} \mathbf{v}_{precond} - \mathbf{g}$ , where  $\mathbf{v}_{precond} = \text{Preconditioner}(\mathbf{g})$  are concentrated at the boundaries between subdomains. In particular,  $\mathbf{r}$  can be interpreted in the continuous case as a measure supported at the boundaries between subdomains with most singular term of the form  $\delta'(\mathbf{x})$ . Then, as explained in Page 243 of [56], the restriction of such measure to a subdomain (line 3 in Alg. 1) is not well defined, resulting in a slow convergence rate. One way to solve the residual problem is presented in [56, 58, 67], where an overlapping domain decomposition is used to ensure that at each iteration the residual lies at the interior of a subdomain.

In this work, we solve both problems simultaneously by introducing a simple alternating bi-directional preconditioner, that consists on applying an up-down sweep performed by Alg. 1 followed by an left-right pass sweep, performed by Alg. 1 but applied to the transpose of  $\mathbf{C}$ , which is denoted  $\text{Preconditioner}^T$  in Alg. 2. The introduction of the bi-directional preconditioner in Alg. 2 handles effectively the spurious grazing waves in one sweep by performing another sweep in the orthogonal direction.

Stage	Complexity
Off-line	$\mathcal{O}(N)$
On-line	$\mathcal{O}\left(n_{\text{iter}}^{\text{SP}}(N \log N + N \cdot n_{\text{iter}})\right)$

TABLE 1

Complexity for each state of the solver, supposing the  $q$ , the thickness of each extended local subdomain, is  $\mathcal{O}(1)$ .

In addition, the application of the preconditioner in the orthogonal direction, after a step of iterative refinement confines the residual to the interior of the subdomains, thus eliminating the ambiguity and resulting in fast convergence.

In particular, as it will be shown in the numerical tests, the bi-directional preconditioner is more efficient and less sensitive to the geometry, direction of incoming waves and sign of the perturbation, albeit with a higher memory footprint. The sweep in the orthogonal direction can easily be implemented as preconditioning the transpose of the matrix  $\mathbf{C}$  using Alg. 1.

ALGORITHM 2. *Bi-directional Preconditioner*

```

1: function  $\mathbf{v} = \text{BIDIRECTIONALPRECONDITIONER}(\mathbf{g})$ 
2:    $\mathbf{u}_1 = \text{Preconditioner}(\mathbf{g})$  ▷ Up-Down pass
3:    $\mathbf{e} = \mathbf{g} - \mathbf{C}\mathbf{u}_1$  ▷ Iterative refinement
4:    $\mathbf{u}_2 = \text{Preconditioner}^T(\mathbf{e})$  ▷ Left-Right pass
5:    $\mathbf{v} = \mathbf{u}_1 + \mathbf{u}_2$ 
6: end function

```

**4. Two-level Preconditioner.** Now we summarize the two-level preconditioner:

- the inner level, in which  $u_{\text{approx}}$  in (2.8) is computed approximately using the bi-directional preconditioner Alg. 2;
- and, the outer level, in which (2.1) is solved iteratively using GMRES preconditioned with Alg. 3.

ALGORITHM 3. *Lippmann-Schwinger Preconditioner*

```

1: function  $\mathbf{u} = \text{LIPPMANN-SCHWINGER PRECONDITIONER}(\mathbf{f})$ 
2:    $u = \text{GMRES}(\mathbf{C}, \mathbf{A}\mathbf{f}, \text{BiDirectionalPreconditioner}, \text{tol} = 10^{-3})$ 
3: end function

```

There is a balance to the choice of the tolerance for the inner GMRES iteration (in Alg. 3): if the tolerance is chosen too small, then the constants are large; on the other hand, if the tolerance is too big, the constants will decrease but the number of outer iterations will grow with the frequency, thus increasing the asymptotic cost. The best choice in our case, is to set the inner loop tolerance to  $10^{-3}$ . This choice ensures that the amount of outer iterations will be the same as using a direct multi-frontal method and produces reasonable constants.

As mentioned in the introduction, the algorithm presented in this paper has two phases: an off-line phase performed only once, and an on-line phase performed for each right-hand-side, or incoming wave. The complexity of the solver is summarized in Table 1; in the sequel we provide further details for the complexity for each phase, in which we suppose the number of subdomains,  $L$ , is  $\mathcal{O}(n)$  and that  $q = \mathcal{O}(1)$ .

**4.1. Off-line complexity.** For the off-line or setup phase, we need to build  $\mathbf{C}$  for which we need to compute  $\mathbf{A}$ . As explained in [63]  $\mathbf{A}$  can be constructed by solving only a hand-full of optimization problems. This property arises mainly by the translation invariance of the convolution kernel, meaning that we can reuse the same stencil for many grid-points. Solving the minimization problem needs a SVD of a

$9 \times N$  matrix, which can be performed in  $\mathcal{O}(N)$  complexity. To create  $\mathbf{C}$  we need to multiply  $\mathbf{A}$  times  $\mathbf{H}$ ; exploiting the sparsity of  $\mathbf{A}$  and performing the multiplication in the correct order, the product can be carried in  $\mathcal{O}(N)$  operations.

We can build all the local matrices  $\mathbf{C}^\ell$  in  $\mathcal{O}(N)$  complexity by reusing  $\mathbf{A}$ , we build  $L = \mathcal{O}(n)$  matrices with  $\mathcal{O}(n)$  elements each. Finally, the factorization of all the local problems can be performed in  $\mathcal{O}(N)$  complexity. Each local matrix is a circulant matrix and their factorization can be performed in  $\mathcal{O}(n)$  time, and we have  $L = \mathcal{O}(n)$  systems to factorize, which results in  $\mathcal{O}(N)$  complexity. We point out that assembly and factorization of the local matrices is a embarrassingly parallel operation which can be effortlessly parallelized.

**4.2. On-line complexity.** For the outer GMRES we need to apply  $\mathbf{H}$ , which consists of multiplications by sparse matrices and a convolution that can be performed fast via a FFT (see section 2.3 of [26]). The cost of applying  $\mathbf{H}$  is dominated by the application of the FFT, resulting on an overall  $\mathcal{O}(N \log N)$  complexity. For Alg. 1 each local solve can be performed in  $\mathcal{O}(n)$  time resulting in a complexity  $\mathcal{O}(N)$ . We have an overall complexity of  $\mathcal{O}(N)$  for each application of Alg. 1.

Finally, the overall on-line complexity depends on the number of outer and inner iterations to converge. From [63] the number of outer iterations is almost independent of the frequency,  $n_{\text{iter}}^{\text{SP}} = \mathcal{O}(1)$ , and we provide in the next section numerical evidence that under the assumptions mentioned in the introduction, the number of inner iterations grows logarithmically with the frequency,  $n_{\text{iter}} = \mathcal{O}(\log \omega)$ , and given that we are in the high-frequency regime  $n_{\text{iter}} = \mathcal{O}(\log N)$ , resulting in the overall complexity mentioned in the introduction.

**5. Numerical Experiments.** The numerical experiments were implemented in Julia v0.4.3 linked against Intel MKL 11.2, and were performed in a server with dual socket Intel Xeon CPU E5-2670, and 386 Gb in RAM. We used the FFT library FFTW3 [34] for the application of  $\mathbf{H}$ . For the local solves in the inner iteration we used the PARDISO [54] and UMFPACK [23] solvers.

One of the main advantages of the method of polarized traces lies on its modularity. We used the state-of-the-art library FFTW3 for the application of  $\mathbf{H}$  but it is possible to further accelerate the application of  $\mathbf{H}$  by using GPU-accelerated libraries (such as cuFFT [51]), or a distributed-memory FFT libraries. Within the main bottle neck of the algorithm, which is the application of the preconditioner, it is possible to effortlessly exchange the linear solver. In our experiments we used PARDISO [54] and UMFPACK [23], but it is easy to exchange them by solvers using compressed linear algebra [52, 3, 43, 38], or by massively parallel solvers such as MUMPS [5] (or its compressed linear algebra variant [4]) or SuperLU-DIST [24], among many others.

In this section we corroborate the complexity claims for our method; in particular, present 3 different kinds of numerical experiments:

1. we study the asymptotic cost of solving sparsified systems from discretizations of the Lippmann-Schwinger equation using the bi-directional preconditioner;
2. we solve a wave scattering example in three different medium profiles shown in Fig. 4 using the full two level preconditioner,
3. we solve an example arising from reflectometry of plasmas confined in a fusion reactor [62].

**5.1. Sparsified system.** First, we assemble two linear systems by sparsifying the Lippmann-Schwinger equation for the perturbation given in Fig. 4 (*left*) and demonstrate the performance of the original method of polarized traces, and the new




FIG. 4. (left) Perturbation used for the wave scattering. (right) Wave-speed for 64 randomly placed Gaussian bumps

N	$\omega$	L	$T_{PT}^+$ [s]	$T_{BD}^+$ [s]	$T_{PT}^-$ [s]	$T_{BD}^-$ [s]
$200 \times 200$	200	4	0.5 ( <b>3.6</b> )	0.2 ( <b>2.7</b> )	0.5 ( <b>3.7</b> )	0.2 ( <b>2.9</b> )
$400 \times 400$	400	8	2.1 ( <b>4.2</b> )	0.7 ( <b>3.1</b> )	2.2 ( <b>4.3</b> )	0.6 ( <b>3.0</b> )
$500 \times 500$	500	10	3.5 ( <b>4.6</b> )	1.1 ( <b>3.2</b> )	3.3 ( <b>4.5</b> )	1.0 ( <b>3.0</b> )
$800 \times 800$	800	16	11.5 ( <b>5.7</b> )	3.2 ( <b>3.8</b> )	8.62 ( <b>4.8</b> )	2.7 ( <b>3.1</b> )
$1000 \times 1000$	1000	20	20.4 ( <b>6.6</b> )	5.1 ( <b>4.1</b> )	14.0 ( <b>5.2</b> )	4.2 ( <b>3.1</b> )
$1250 \times 1250$	1250	25	30.5 ( <b>7.9</b> )	8.7 ( <b>4.3</b> )	22.5 ( <b>5.4</b> )	7.3 ( <b>3.3</b> )
$1400 \times 1400$	1400	28	43.5 ( <b>8.9</b> )	11.1 ( <b>4.7</b> )	29.1 ( <b>5.6</b> )	8.7 ( <b>3.4</b> )
$1600 \times 1600$	1600	32	62.8 ( <b>10.1</b> )	15.8 ( <b>4.7</b> )	41.8 ( <b>5.8</b> )	12.6 ( <b>3.4</b> )
$2000 \times 2000$	2000	40	120.5 ( <b>12.6</b> )	25.0 ( <b>5.3</b> )	63.8 ( <b>6.1</b> )	18.1 ( <b>3.5</b> )

TABLE 2

Average execution time and number of iterations (in parenthesis) to solve the sparsified system arising from the Lippmann-Schwinger equation using the positive (+) and negative (-) perturbation, using either the method of polarized traces (PT), or the new bi-directional preconditioner (BD). We use GMRES with a tolerance of  $10^{-6}$ , in this case  $N = n^2$  is the number of degrees of freedom,  $\omega$  is the frequency and  $L$  is the number of subdomains.

bi-directional preconditioner. The two sparsified Lippmann-Schwinger systems (3.1) correspond to the same perturbation shown in Fig. 4 (left) but with a different sign, to which we refer to as the positive and negative perturbation. We solve the resulting sparsified systems with GMRES (tolerance of  $10^{-6}$ ) preconditioned with the method of polarized traces [66] and the bi-directional preconditioner in Alg. 2. We solve (3.1) for 64 different right-hand sides representing 64 different incident waves, and we compute an average time. We repeat the process for a set of increasing frequencies, such that we have 10 points per wavelength in the background medium. Moreover, the domain is decomposed in  $L = n/50$  subdomains, such that each subdomain is roughly 5 wavelengths thick, and each subdomain is extended by 10 grid points in both directions (see Fig 1). The average execution times are presented in Table 2.

Table 2 clearly depicts the different behavior between the matrix-free version of the method of polarized traces, and its bi-directional variant. We can clearly observe that the execution time increases faster for the original method of polarized traces. This trend is more marked in the case of the positive perturbation as explained in Section 3.3. In particular, we observe that the execution time (due to an increase in the number of iterations) grows as  $\mathcal{O}(N\omega^{1/2}) = \mathcal{O}(N^{5/4})$  instead of  $\mathcal{O}(N \log N)$  and thus increasing overall the complexity of the global algorithm. In addition, we

can observe that the constants are lower for the bi-directional preconditioner. We point out that the bottleneck of both preconditioners is the local solves, which are performed using a multifrontal solver; the main difference between both methods is the number of such solves at each iteration: for the method of polarized traces we need 7 solves per layer and per iteration, whereas the bi-directional preconditioner needs only 4; in addition, the bi-directional preconditioner converges in a lower number of iterations.




FIG. 5. *Left: execution time of the off-line stage; Right: execution time of the on-line stage of solving the Sparsified Lippmann-Schwinger equation using GMRES preconditioned with Alg. 2. The tolerance was fixed to  $10^{-6}$ , and  $\omega h = \mathcal{O}(1)$ .*

Fig. 5 (*left*), shows the scaling for the factorization of the local problems and the assembly of the sparsified system. In addition, Fig. 5 (*right*), shows the execution time for the solution of the sparsified system using the bi-directional algorithm as reported in Table 2 but in logarithmic scale. We can observe from Fig. 5 that the assembly and factorization cost is  $\mathcal{O}(N)$  and that the online cost of one inner iteration is  $\mathcal{O}(N \log N)$ , and, in particular,  $n_{\text{iter}} = \mathcal{O}(\log \omega)$ .




FIG. 6. *Real part of the total wavefield generated by the scattering of a plane wave impinged on the negative perturbation (left) and on the positive perturbation (right), in both examples we have  $\omega = 500$ .*

**5.2. Wave scattering.** We test the performance of the two-level preconditioner by solving the Lippmann-Schwinger equation (2.1) associated to three problems on wave scattering. We use three different perturbation profiles, the negative smooth



perturbation, the positive smooth perturbation and Gaussian bumps as shown in Fig. 4 (*left*) and (*right*). The profiles were chosen to depict the behavior of the method in three generic settings for wave propagation: a defocussing medium (negative perturbation), a focusing medium (positive perturbation) and a medium with multiple scattering (Gaussian bumps). We can observe the total wavefield scattered by a plane-wave impinging on each perturbation in Fig. 6 (*left*), Fig. 6 (*right*) and Fig. 10 (*left*), respectively.




FIG. 7. Average execution time for the on-line stage of solving the high-frequency Lippmann-Schwinger equation with GMRES (tolerance of  $10^{-10}$ ) preconditioned with Alg. 3 with an inner tolerance of  $10^{-3}$ , for the three different perturbation profiles;  $\omega h = \mathcal{O}(1)$ .

For each perturbation profile we solve a system of increasing size, keeping  $\omega h$  constant. Each system is solved using GMRES with a tolerance of  $10^{-10}$ , preconditioned with Alg. 3 with a inner tolerance of  $10^{-3}$ ; the domain is decomposed in  $L = n/50$  subdomains, and each subdomain is extended 10 grid points in each direction (see Fig. 1). At each frequency, we compute the wavefield generated by a set of 64 different incident waves impinging on the perturbation by solving the Lippmann-Schwinger equation, and we report the average execution time in Fig. 7. From Fig. 7 we can observe that the execution time scales as  $\mathcal{O}(N \log N)$  for each perturbation profile, which corroborates the claims made in the introduction. For the off-line stage we obtain the same execution times as in Fig. 5 (*left*).




FIG. 8. Left: average execution time for the on-line stage of solving the Lippmann-Schwinger with the Gaussian bumps perturbation with the original sparsifying preconditioner, i.e., using multifrontal methods, and using the two-level preconditioner with  $\omega h = \mathcal{O}(1)$ ; Right: execution time of the off-line stage for the two-level preconditioner and the original sparsifying preconditioner.

Moreover, for the sake of comparison, we time the execution time of solving the Lippmann-Schwinger equation using the original sparsifying preconditioner [63] (i.e., solving (3.1) via a multifrontal solver [23]). For the comparison we used the multiple

scattering perturbation in Fig. 4 (*right*), with the same set of frequencies and incoming waves as before. We use the implementation of the sparsifying preconditioner in [65]. The average execution times of the on-line stage are reported in Fig. 8 (*left*) together with the averages times for the two-level preconditioner. From Fig. 8 (*left*) we can observe that, unsurprisingly, the original sparsifying preconditioner has the same asymptotic on-line cost as the two-level preconditioner presented in this paper; however, the main difference lies in the off-line stage, in which the multifrontal solver has a complexity of  $\mathcal{O}(N^{3/2})$  compared to  $\mathcal{O}(N)$  for the two-level preconditioner as shown in Fig. 8 (*right*).




FIG. 9. *Perturbation for the plasma physics example (taken from [2]) we can observe the at the interior the perturbation becomes negative, thus yielding a strongly elliptic problem locally, thus becoming an evanescent medium.*




FIG. 10. *Real part of the total wavefield generated by the scattering of a plane wave in the Gaussian bumps model (left) and in the plasma model (right). In both examples  $\omega = 500$ .*

**5.3. Plasma physics.** Finally, we test the performance of the two-level preconditioner using a problem arising from plasma physics, in particular, a radio-frequency (RF) wave propagating in a fusion reactor. Under some simplifications, the propaga-




FIG. 11. Average execution time for solving the Lippmann-Schwinger equation using the plasma perturbation profile (left), and the average number of outer iterations needed for convergence (right). The problem was solved with GMRES (tolerance of  $10^{-10}$ ) preconditioned with the two-level preconditioner with an inner tolerance of  $10^{-3}$  and  $\omega h = \mathcal{O}(1)$ . The extra logarithmic factor comes from the increasing number of outer iterations.

tion of a RF wave can be described using a cold plasma model [62]. We borrow from [2] a model<sup>4</sup> corresponding to an idealized plasma density profile shown in Fig. 9. The physics in this case are different, an incoming wave propagates until it hits a cut-off region in which the problem becomes strictly elliptic or evanescent, thus causing the waves to reflect, as depicted in Fig. 10 (right).

We solve a system of increasing size, keeping  $\omega h$  constant; we decompose the domain such that we have roughly 5 wavelengths inside each subdomain, and each subdomain is extended 10 grid points in each direction. At each frequency, we compute the wavefield generated by a set of 64 different incident waves impinging on the perturbation by solving the Lippmann-Schwinger equation and report the average execution times in Fig. 11 (left).

In the case of the plasma example, we observe from Fig. 11 (left) a slightly bigger scaling. The main reason for this scaling is the fact that the number of outer iteration grows logarithmically, as shown in Fig. 11 (right), which is presumably caused by the transition of the physics at the cut-off region. Nonetheless, the number of inner iterations grows like  $\mathcal{O}(\log N)$ , the same as in the wave scattering case, resulting in an overall complexity of  $\mathcal{O}(N \log^2 N)$ , which, to the our knowledge, compares favorably to most of current solvers.

**6. Conclusions.** We have presented a new fast iterative method for the high-frequency Lippmann-Schwinger equation in 2D. The method has an asymptotic cost of  $\mathcal{O}(N \log N)$  for the wave scattering problems presented in this paper.

Some future work and extensions are: the design of more efficient ABC, in order to use only one directional preconditioner, thus decreasing the memory footprint; extensions to 3D, in which a nested approach can be used to obtain low complexity algorithms; parallelization of the solvers using state-of-the-art linear solvers; reduction of the constant via compressed linear algebra, which can provide extra savings in the plasma example, where the local ellipticity of the problem can be exploited.

Another line of research is to further reduce the on-line complexity by using the original reduction to a SIE of the method of polarized traces, coupled to the precomputation of integral operators and their compression.

We point out that it is envisageable to change the outer GMRES loop with the

<sup>4</sup>We refer the interested reader to Example 4.3 of [2] for further details.

flexible GMRES [53] algorithm, and reduce the inner-level tolerance to accelerate the algorithm, without a penalty on the asymptotic complexity.

**7. Acknowledgments.** We would like to thank Lexing Ying, and Leslie Greengard for fruitful discussions, and Laurent Demanet for computational resources.

#### REFERENCES

- [1] M. B. AMAR AND FARNOUX F. C., *Numerical solution of the Lippmann-Schwinger equations in photoemission: application to Xenon*, Journal of Physics B: Atomic and Molecular Physics, 16 (1983), p. 2339.
- [2] 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*, ArXiv e-prints, [math.NA] 1505.07157 (2015).
- [3] S. AMBIKASARAN AND E. DARVE, *An  $\mathcal{O}(n \log n)$  fast direct solver for partial hierarchically semi-separable matrices*, Journal of Scientific Computing, 57 (2013), pp. 477–501.
- [4] P. AMESTOY, C. ASHCRAFT, O. BOITEAU, A. BUTTARI, J.-Y. L’EXCELLENT, AND C. WEISBECKER, *Improving multifrontal methods by means of block low-rank representations*, SIAM Journal on Scientific Computing, 37 (2015), pp. A1451–A1474.
- [5] P. R. AMESTOY, I. S. DUFF, J.-Y. L’EXCELLENT, AND J. KOSTER, *A fully asynchronous multifrontal solver using distributed dynamic scheduling*, SIAM Journal on Matrix Analysis and Applications, 23 (2001), pp. 15–41.
- [6] F. ANDERSSON AND A. HOLST, *A fast, bandlimited solver for scattering problems in inhomogeneous media*, Journal of Fourier Analysis and Applications, 11, pp. 471–487.
- [7] 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*, Computer Methods in Applied Mechanics and Engineering, 128 (1995), pp. 325–359.
- [8] I. M. BABUKA AND S. A. SAUTER, *Is the pollution effect of the FEM avoidable for the Helmholtz equation considering high wave numbers?*, SIAM Journal on Numerical Analysis, 34 (1997), pp. 2392–2423.
- [9] A. BAYLISS, C. I. GOLDSTEIN, AND TURKEL E., *On accuracy conditions for the numerical computation of waves*, Journal of Computational Physics, 59 (1985), pp. 396 – 404.
- [10] M. BEBENDORF, *Hierarchical LU decomposition-based preconditioners for BEM*, Computing, 74 (2005), pp. 225–247.
- [11] M. BEBENDORF AND W. HACKBUSCH, *Existence of  $\mathcal{H}$ -matrix approximants to the inverse fematrix of elliptic operators with  $l^\infty$ -coefficients*, Numerische Mathematik, 95 (2002), pp. 1–28.
- [12] J.-P. BÉRENGER, *A perfectly matched layer for the absorption of electromagnetic waves*, Journal of Computational Physics, 114 (1994), pp. 185–200.
- [13] A. BERMÚDEZ, L. HERVELLA-NIETO, A. PRIETO, AND R. RODRÍGUEZ, *An optimal perfectly matched layer with unbounded absorbing function for time-harmonic acoustic scattering problems*, Journal of Computational Physics, 223 (2007), pp. 469 – 488.
- [14] ———, *Perfectly matched layers for time-harmonic second order elliptic problems*, Archives of Computational Methods in Engineering, 17 (2010), pp. 77–107.
- [15] G. BEYLKIN, C. KURCZ, AND L. MONZÓN, *Fast convolution with the free space Helmholtz Greens function*, Journal of Computational Physics, 228 (2009), pp. 2770 – 2791.
- [16] O. P. BRUNO, *Wave scattering by inhomogeneous media: efficient algorithms and applications*, Physica B: Condensed Matter, 338 (2003), pp. 67 – 73. Proceedings of the Sixth International Conference on Electrical Transport and Optical Properties of Inhomogeneous Media.
- [17] O. P. BRUNO AND E. M. HYDE, *An efficient, preconditioned, high-order solver for scattering by two-dimensional inhomogeneous media*, Journal of Computational Physics, 200 (2004), pp. 670 – 694.
- [18] Z. CHEN AND X. XIANG, *A source transfer domain decomposition method for Helmholtz equations in unbounded domain*, SIAM Journal on Numerical Analysis, 51 (2013), pp. 2331–2356.
- [19] ———, *A source transfer domain decomposition method for Helmholtz equations in unbounded domain part II: Extensions*, Numerical Mathematics: Theory, Methods and Applications, 6 (2013), pp. 538–555.
- [20] H. CHENG, J. HUANG, AND T. J. LEITERMAN, *An adaptive fast solver for the modified Helmholtz equation in two dimensions*, Journal of Computational Physics, 211 (2006), pp. 616 – 637.

- [21] J. W. COOLEY AND J. W. TUKEY, *An algorithm for the machine calculation of complex Fourier series*, Mathematics of Computation, 19 (1965), pp. 297–301.
- [22] E. CORONA, P.-G. MARTINSSON, AND D. ZORIN, *An direct solver for integral equations on the plane*, Applied and Computational Harmonic Analysis, 38 (2015), pp. 284 – 317.
- [23] T. A. DAVIS, *Algorithm 832: UMFPACK v4.3—an unsymmetric-pattern multifrontal method*, ACM Transactions on Mathematical Software, 30 (2004), pp. 196–199.
- [24] J. W. DEMMEL, S. C. EISENSTAT, J. R. GILBERT, X. S. LI, AND J. W. H. LIU, *A supernodal approach to sparse partial pivoting*, SIAM J. Matrix Analysis and Applications, 20 (1999), pp. 720–755.
- [25] B. DESPRÉS, *Domain decomposition method and the Helmholtz problem*, Mathematical and Numerical Aspects of Wave Propagation Phenomena, (1991), p. 4452.
- [26] R. DUAN AND V. ROKHLIN, *High-order quadratures for the solution of scattering problems in two dimensions*, Journal of Computational Physics, 228 (2009), pp. 2152 – 2174.
- [27] I. S. DUFF AND J. K. REID, *The multifrontal solution of indefinite sparse symmetric linear*, ACM Trans. Math. Softw., 9 (1983), pp. 302–325.
- [28] B. ENGQUIST AND A. MAJDA, *Absorbing boundary conditions for the numerical simulation of waves*, Math. Comp., 31 (1977), pp. 629–651.
- [29] B. ENGQUIST AND L. YING, *Sweeping preconditioner for the Helmholtz equation: Hierarchical matrix representation*, Communications on Pure and Applied Mathematics, 64 (2011), pp. 697–735.
- [30] B. ENGQUIST AND L. YING, *Sweeping preconditioner for the Helmholtz equation: moving perfectly matched layers*, Multiscale Modeling & Simulation, 9 (2011), pp. 686–710.
- [31] B. ENGQUIST AND H.-K. ZHAO, *Absorbing boundary conditions for domain decomposition*, Applied Numerical Mathematics, 27 (1998), pp. 341 – 365. Special Issue on Absorbing Boundary Conditions.
- [32] ———, *Approximate separability of Green’s function for high frequency Helmholtz equations*, March 2014.
- [33] 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, Ivan G. Graham, Thomas Y. Hou, Omar Lakkis, and Robert Scheichl, eds., vol. 83 of Lecture Notes in Computational Science and Engineering, Springer Berlin Heidelberg, 2012, pp. 325–363.
- [34] M. FRIGO AND S. G. JOHNSON, *The design and implementation of FFTW3*, Proceedings of the IEEE, 93 (2005), pp. 216–231. Special issue on “Program Generation, Optimization, and Platform Adaptation”.
- [35] M.J. GANDER, *Optimized schwarz methods*, SIAM Journal on Numerical Analysis, 44 (2006), pp. 699–731.
- [36] M. J. GANDER, F. MAGOULÈS, AND F. NATAF, *Optimized Schwarz methods without overlap for the Helmholtz equation*, SIAM Journal on Scientific Computing, 24 (2002), pp. 38–60.
- [37] A. GEORGE, *Nested dissection of a regular finite element mesh*, SIAM Journal on Numerical Analysis, 10 (1973), pp. 345–363.
- [38] A. GILLMAN, A.H. BARNETT, AND P.G. MARTINSSON, *A spectrally accurate direct solution technique for frequency-domain scattering problems with variable media*, BIT Numerical Mathematics, (2014), pp. 1–30.
- [39] A. GILLMAN AND P.-G. MARTINSSON, *A direct solver with  $\mathcal{O}(n)$  complexity for variable coefficient elliptic PDEs discretized via a high-order composite spectral collocation method*, SIAM Journal on Scientific Computing, 36 (2014), pp. A2023–A2046.
- [40] I. G. GRAHAM, E. A. SPENCE, AND E. VAINIKKO, *Domain decomposition preconditioning for high-frequency Helmholtz problems using absorption*, ArXiv e-prints, (2015).
- [41] K. L. HO AND L. GREENGARD, *A fast direct solver for structured linear systems by recursive skeletonization*, SIAM Journal on Scientific Computing, 34 (2012), pp. A2507–A2532.
- [42] K. L. HO AND L. YING, *Hierarchical interpolative factorization for elliptic operators: Integral equations*, Communications on Pure and Applied Mathematics, (2015), pp. n/a–n/a.
- [43] M. GU J. XIA, S. CHANDRASEKARAN AND X. S. LI, *Superfast multifrontal method for large structured linear systems of equations*, SIAM Journal on Matrix Analysis and Applications, 31 (2010), pp. 1382–1411.
- [44] S. JOHNSON, *Notes on perfectly matched layers (PMLs)*, March 2010.
- [45] F. LANZARA, V. MAZ’YA, AND G. SCHMIDT, *Numerical solution of the Lippmann–Schwinger equation by approximate approximations*, Journal of Fourier Analysis and Applications, 10 (2004), pp. 645–660.
- [46] W. LENG, *A fast propagation method for the Helmholtz equation*, ArXiv e-prints, (2015).
- [47] W. LENG AND L. JU, *An overlapping domain decomposition preconditioner for the Helmholtz equation*, ArXiv e-prints, (2015).

- [48] F. LIU AND L. YING, *Additive sweeping preconditioner for the Helmholtz equation*, ArXiv e-prints, (2015).
- [49] ———, *Recursive sweeping preconditioner for the 3D Helmholtz equation*, ArXiv e-prints, (2015).
- [50] P.-G. MARTINSSON AND V. ROKHLIN, *A fast direct solver for boundary integral equations in two dimensions*, *Journal of Computational Physics*, 205 (2005), pp. 1 – 23.
- [51] J. NICKOLLS, I. BUCK, M. GARLAND, AND K. SKADRON, *Scalable parallel programming with CUDA*, *Queue*, 6 (2008), pp. 40–53.
- [52] F.-H. ROUET, X. S. LI, P. GHYSELS, AND A. NAPOV, *A distributed-memory package for dense Hierarchically Semi-Separable matrix computations using randomization*, ArXiv e-prints, (2015).
- [53] Y. SAAD, *A flexible inner-outer preconditioned GMRES algorithm*, *SIAM Journal on Scientific Computing*, 14 (1993), pp. 461–469.
- [54] O. SCHENK AND K. GRTNER, *Solving unsymmetric sparse systems of linear equations with PARDISO*, *Future Generation Computer Systems*, 20 (2004), pp. 475 – 487. Selected numerical algorithms.
- [55] J. SIFUENTES, *Preconditioned Iterative Methods for Inhomogeneous Acoustic Scattering Applications*, PhD thesis, Rice University, Houston TX, USA, 2010.
- [56] C. STOLK, *A rapidly converging domain decomposition method for the Helmholtz equation*, *Journal of Computational Physics*, 241 (2013), pp. 240–252.
- [57] C. C. STOLK, *A dispersion minimizing scheme for the 3-D Helmholtz equation with applications in multigrid based solvers*, ArXiv e-prints, (2015).
- [58] ———, *An improved sweeping domain decomposition preconditioner for the Helmholtz equation*, ArXiv e-prints, (2015).
- [59] E. TURKEL, D. GORDON, R. GORDON, AND S. TSYNKOV, *Compact 2D and 3D sixth order schemes for the Helmholtz equation with variable wave number*, *Journal of Computational Physics*, 232 (2013), pp. 272 – 287.
- [60] G. VAINIKKO, *Fast solvers of the Lippmann-Schwinger equation*, in *Direct and Inverse Problems of Mathematical Physics*, RobertP. Gilbert, Joji Kajiwara, and YongzhiS. Xu, eds., vol. 5 of *International Society for Analysis, Applications and Computation*, Springer US, 2000, pp. 423–440.
- [61] A. VION AND C. GEUZAIN, *Double sweep preconditioner for optimized Schwarz methods applied to the Helmholtz problem*, *Journal of Computational Physics*, 266 (2014), pp. 171–190.
- [62] H. WEITZNER, *Lower hybrid waves in the cold plasma model*, *Communications on Pure and Applied Mathematics*, 38 (1985), pp. 919–932.
- [63] L. YING, *Sparsifying preconditioner for the Lippmann–Schwinger equation*, *Multiscale Modeling & Simulation*, 13 (2015), pp. 644–660.
- [64] L. ZEPEDA-NÚÑEZ AND L. DEMANET, *The method of polarized traces for the 2d helmholtz equation*, *Journal of Computational Physics*, 308 (2016), pp. 347 – 388.
- [65] L. ZEPEDA-NUNEZ. <https://github.com/Forgotten/Lippmann-Schwinger>, 2016.
- [66] L. ZEPEDA-NUNEZ AND L. DEMANET, *Nested domain decomposition with polarized traces for the 2D Helmholtz equation*, ArXiv e-prints, [math.NA] 1510.01831 (2015).
- [67] 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.