Direct/iterative hybrid solver for scattering by inhomogeneous media

Oscar P. Bruno* Ambuj Pandey†

Abstract

This paper presents a fast high-order method for the solution of two-dimensional problems of scattering by penetrable inhomogeneous media, with application to high-frequency configurations containing (possibly) discontinuous refractivities. The method relies on a hybrid direct/iterative combination of 1) A differential volumetric formulation (which is based on the use of appropriate Chebyshev differentiation matrices enacting the Laplace operator) and, 2) A second-kind boundary integral formulation (which, once again, utilizes Chebyshev discretization, but, in this case, in the boundary-integral context). The approach enjoys low dispersion and high-order accuracy for smooth refractivities, as well as second-order accuracy (while maintaining low dispersion) in the discontinuous refractivity case. The solution approach proceeds by application of Impedance-to-Impedance (ItI) maps to couple the volumetric and boundary discretizations. The volumetric linear algebra solutions are obtained by means of a multifrontal solver, and the coupling with the boundary integral formulation is achieved via an application of the iterative linear-algebra solver GMRES. In particular, the existence and uniqueness theory presented in the present paper provides an affirmative answer to an open question concerning the existence of a uniquely solvable second-kind ItI-based formulation for the overall scattering problem under consideration. Relying on a modestly-demanding scatterer-dependent precomputation stage (requiring in practice a computing cost of the order of $O(N^\alpha)$ operations, with $\alpha \approx 1.07$, for an N-point discretization and for the relevant Chebyshev accuracy orders q used), together with fast ($O(N)$-cost) single-core runs for each incident field considered, the proposed algorithm can effectively solve scattering problems for large and complex objects possibly containing discontinuities and strong refractivity contrasts.

1 Introduction

This paper considers the problem of evaluation of wave scattering by penetrable inhomogeneous media in two dimensions. This is a problem of fundamental importance in a wide range of applications, including underwater acoustics, biological and medical imaging, seismology and geophysics, etc. In all of these applications, it is highly desirable to utilize efficient and accurate numerical methods which can deal with arbitrary scattering geometries and (often discontinuous) refractive index distributions, even in the high-frequency regime. As is well known [27, 47, 49], this problem presents a number of challenges, as it requires use of large numbers of discretization points and, for iterative solvers, increasingly large numbers of iterations as the frequencies and/or refractive-index values increase. This paper presents a hybrid iterative/direct linear algebra formulation for this problem, which, like the approach [34, 35], combines a volumetric differential formulation in a bounded region, and a surface boundary integral equation that provides the coupling to the complementary unbounded exterior domain. Unlike the previous volumetric/boundary formulation [34, 35], which tackles the volumetric problem via a finite-element discretization, further, the method proposed here utilizes (i) Polynomial approximation patches of accuracy of finite order q (with, e.g., $q = 10, 20, 40$); (ii) A high-order boundary integral formulation, as well as, both, (iii) A multifrontal direct linear solver (the Intel MKL implementation of the multifrontal solver Pardiso [12, 43, 44]);

*Computing and Mathematical Sciences, Caltech, Pasadena, CA 91125, USA, obruno@caltech.edu

†Indian Institute of Science Education and Research Bhopal (IISER Bhopal), ambuj@iiserb.ac.in

and (iv) The iterative linear solver GMRES. Leveraging a new version of the smoothing technique [31], finally, the proposed algorithm yields second-order convergence even for discontinuous refractivities. As a result of these innovations, the proposed algorithm can be quite effective: after a modestly-demanding precomputation stage, requiring a computing cost that grows nearly linearly with the number N of degrees of freedom used (Figure 4a demonstrates a growth of the order of $\approx O(N^\alpha)$ with $\alpha = 1.07$), and at a cost per GMRES iteration that grows essentially linearly with N, the proposed method can evaluate, with a favorable number of iterations, scattering by configurations including large and complex objects as well as strong refractivity contrasts and discontinuities—with high accuracy and in fast single-core runs. A variety of numerical experiments have shown (cf. Figure 4 and its caption) that, as may be expected in view of the algebraic character of the precomputation and iteration stages, for each order q and each discretization size N, the associated computing times are essentially constant asymptotically as the frequency κ grows.

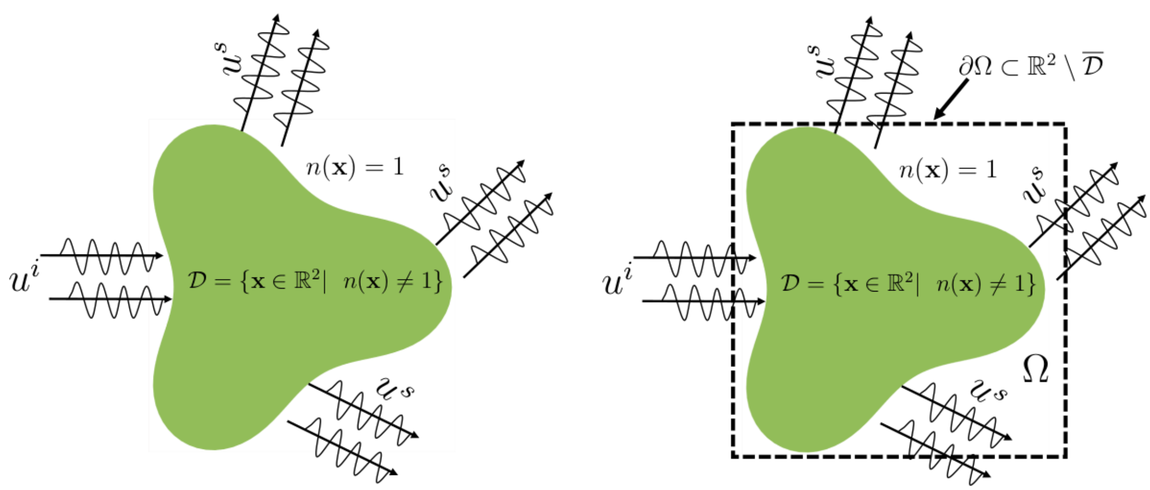

Figure 1: **Left:**: Scattering by an inhomogeneous region $\mathcal{D} = \{\mathbf{x} \in \mathbb{R}^2 | n(\mathbf{x}) \neq 1\}$. An incident wave u^i satisfying (1) impinges upon the inhomogeneity $\mathcal{D}$, and thereby the scattered field u^s (satisfying (3)) results. **Right:** The unique solution of the problem (1)-(3), i.e., the total wave u, which is equal to the sum of incident wave u^i and scattered wave u^s, is computed in a computational domain $\overline{\Omega}$ (the square region, enclosed by dotted black lines) containing $\mathcal{D}$, by solving an equivalent formulation (15)-(17) in $\overline{\Omega}$.

The problem we consider concerns scattering of an incident time-harmonic acoustic wave u^i by a bounded two-dimensional inhomogeneity $\mathcal{D} = \{\mathbf{x} : n(\mathbf{x}) \neq 1\} \subset \mathbb{R}^2$, where $n(\mathbf{x})$ denotes the (possibly discontinuous) index of refraction, which is assumed to equal unity in the complement $\mathbb{R}^2 \setminus \overline{\mathcal{D}}$ of the closure $\overline{\mathcal{D}}$ of the set $\mathcal{D}$, as depicted on the left portion of Figure 1. Throughout this paper it is assumed that u^i satisfies the free space Helmholtz equation

$$\Delta u^i(\mathbf{x}) + \kappa^2 u^i(\mathbf{x}) = 0, \quad \mathbf{x} \in \mathbb{R}^2, \tag{1}$$

where κ is the wave number of the incoming wave u^i. The total acoustic field u (which equals the sum $u = u^i + u^s$ of the incident and scattered fields) satisfies the equation [23]

$$\Delta u(\mathbf{x}) + \kappa^2 n^2(\mathbf{x}) u(\mathbf{x}) = 0, \quad \mathbf{x} \in \mathbb{R}^2, \tag{2}$$

and the scattered field u^s satisfies the Sommerfeld radiation condition

$$\lim_{r\to\infty} \sqrt{r} \left(\frac{\partial u^s}{\partial r} - i\kappa u^s \right) = 0, \tag{3}$$

where $r = (x_1^2 + x_2^2)^{1/2}$ and $i = \sqrt{-1}$ is the imaginary unit.

The simplest computational approaches to the problem (2)-(3) proceed by replacing the unbounded propagation region by a bounded computational domain containing the scatterer $\mathcal{D}$ in its interior (which results in the introduction of an artificial boundary), and then tackling the resulting bounded problem by means of finite element or finite difference discretizations. These approaches yield sparse linear systems, and, in order to satisfy the radiation condition (3), they rely on the use of absorbing boundary conditions. The classical absorbing-boundary techniques [28] and the more recent PML approaches [11,21,32] generally require, for accuracy, the use of a relatively large distance between scatterers and the absorbing boundary regions, and, thus, relatively large computational domains—leading to large number of unknowns and correspondingly large linear systems. In contrast, as illustrated in Example 4.4, the proposed approach can utilize computational boundaries that lie arbitrarily close to the scattering surfaces. Further, although square computational domains are considered in this paper for definiteness (as depicted on the right portion of Figure 1 and described in detail in Section 2), the proposed algorithm can be generalized in a straightforward manner to computational domains consisting of a union of disjoint square components covering the region $\{n(x) \neq 1\}$ occupied by the scatterer—thus leading, upon use of sufficiently small square components, to computational domains tightly covering the scattering regions where the refractivity is different from the free-space refractivity. Other absorbing-boundary approaches [30] allow for the use of computational boundaries that lie near the scattering boundaries—at the expense of a degree of algorithmic complexity. Additionally, the frequently used low-order finite-difference (FDM) and finite-element methods (FEM) for the problem (2) generally suffer from significant dispersion errors [8], also known as pollution error [6] (a problem which can be alleviated or even eliminated [40] by employing high-order finite elements), and they lead to linear systems which require large numbers of iterations if treated by means of iterative linear algebra solvers [26].

An alternative widely-used computational approach for the problem (2)-(3) relies on the equivalent *Lippmann-Schwinger* volumetric integral equation [23, 31]

$$u(\mathbf{x}) + \kappa^2 \int_{\mathcal{D}} G_\kappa(\mathbf{x} - \mathbf{y}) u(\mathbf{y}) m(\mathbf{y}) d\mathbf{y} = u^i(\mathbf{x}), \qquad \mathbf{x} \in \mathbb{R}^2, \tag{4}$$

where $G_\kappa(\mathbf{x}) = \frac{i}{4} H_0^1(\kappa|\mathbf{x}|)$ denotes the radiating fundamental solution of Helmholtz equation in free space and $m(\mathbf{x}) = 1 - n^2(\mathbf{x})$ is the contrast function. This formulation offers several advantages; notably this approach only requires discretization of the scattering region $\mathcal{D}$, and the solutions thus obtained automatically satisfy the Sommerfeld radiation condition (3). Additionally, equation (4) is a Fredholm equation of the second kind, and, therefore, upon discretization, the condition number of the resulting linear system remains essentially constant as the discretization is refined. Unfortunately, however, scattering solvers based on volumetric integral equation formulations give rise to certain difficulties, since 1) The resulting discrete linear systems, which are dense and non-Hermitian, cannot be effectively solved by means of classical direct linear-algebra techniques except for problems that are acoustically small; and 2) The use of iterative linear-algebra solvers for such volumetric formulations requires very large number of iterations for convergence whenever the frequency or the contrast function $m(\mathbf{x})$ (or both) are large. In recent years, a number of algorithms, including direct and iterative solvers, have been proposed for the solution of Lippmann-Schwinger equation, for instance, see [2, 5, 16, 17, 25, 42, 45] and references therein. The simplest fast algorithms in this context, which rely on the use of equidistant grids and FFTs, only provide first order convergence in presence of a discontinuous index of refraction. For instance, the scheme introduced in [25] provides a fast high-order FFT-based method for smooth refractivities, but it does not yield higher-order accuracy in presence of discontinuous refractive indices, and it requires large iteration numbers for high-frequencies or high refractivity contrast. The algorithm introduced in [16] exhibits second order convergence in the presence of discontinuous refractivity, but this approach does not address problem 2 above: the algorithm requires large iteration numbers for large frequencies. The recent fast algorithms [4,42] provide convergence-order higher than two via special treatment at discontinuity boundaries, but they also suffer from large iteration numbers at high frequency. Recently preconditioning techniques were introduced in [38,47,49], which were shown to reduce iteration numbers, even at high frequency. No reports have been provided in either theoretical, graphical or tabular form, on the numerical accuracy of the solutions provided by these methods. Further, the effectiveness of these methodologies is highly dependent on the smoothness of the refractive-index function. For example, reference [38, Sec. 2.5] indicates that "if the [velocity] field has... discontinuities neither will the Nyström method be able to give an accurate discretization scheme nor can the sweeping factorization provide... an accurate approximating solution. Thus, for our preconditioner to work, we require certain smoothness from the velocity fields".

Methods which, like the one proposed in this paper, are based on a combination of a volumetric differential formulation coupled with a boundary integral equation for *physically-exact* truncation of the computation domain, have been proposed previously. The first such contributions were provided in [34,35], and extensions to multi-domain iterative solvers in the context of finite-element discretizations can be found in [9,10,13,20]. In these contexts, the interior volumetric PDE is generally discretized by means of FEM of low order of accuracy, while boundary-element or Nyström discretizations are used in the discretization of the boundary integral equation. As mentioned above, the use of low-order FEM methods leads to accuracy degradation as the domain sizes grow, in view of the well known dispersion errors [6, 8] which requires increases in the number of points per wavelength in order to maintain fixed accuracy as the wavenumber κ grows. High-order methods greatly reduce dispersion and pollution errors, and they remain advantageous even in presence of discontinuous PDE coefficients (Tables 7, 9 and 10). Indeed, the improved second-order accurate spectral discretization we introduce for discontinuous-coefficient problems enjoys essentially dispersionless performance—an important feature that is not obtained from commonly used low-order finite-difference or finite-element methods.

A *direct solver* based on spectral discretizations of fixed order of accuracy, with computational complexity of order $O(N^{3/2})$, was introduced in [27]. The method achieves its operation count by decomposing the domain in a number of spectral square patches that are organized in a tree structure, with a subsequent aggregation process, whereby certain "Impedance-to-Impedance" (ItI) maps for individual cells are recursively merged into ItI maps for larger and larger rectangular groups of cells. Ultimately, when the computational domain boundary is reached a boundary integral equation is used in conjunction with the Dirichlet-to-Neumann map (DtN) of the complete domain to enact the interactions between the bounded scatterer and the exterior domain. This algorithm can effectively treat high-frequency problems for which the refractivity is smooth; the illustrations available in the literature do not include applications for which refractivity discontinuities exist, but it is expected that the first-order accuracy would ensue in such cases.

The approach proposed in this paper is a fast hybrid direct/iterative method which is demonstrated to run at a cost of $O(N^{\alpha})$ operations with $\alpha \approx 1.07$, and which, as illustrated in Section 4, enjoys a number of additional appealing features: the algorithm 1) Requires a small, essentially fixed, numbers of iterations as the refractive index (and, thus, the interior wavelength) is increased while keeping the exterior wavelength fixed (Table 10 below); 2) Requires significantly milder increases in iteration numbers (see Tables 7 and 9 and Remark 4 below) than other iterative solvers [16, 37] as the exterior frequency increases, in view of its resolution of all interior multiple scattering events via a direct solver; 3) Exhibits very low dispersion; and, 4) Converges with high-order accuracy for smooth refractivities, and with second-order accuracy (maintaining low dispersion) for discontinuous refractivities, as discretizations are refined. This solver relies on a general-purpose sparse direct solution technique for the volumetric interior problem that, in particular, enforces the PDE at spectral cell boundaries by matching "transmission values" (that is, the values of the solution and its normal derivative) at such boundaries; and it incorporates a second-kind integral formulation in conjunction with an ItI map at the computational domain boundary (instead of the possibly singular DtN map used in [27]). The algorithm is completed by means of the iterative linear-algebra solver GMRES. In particular, the existence and uniqueness theory presented in the present paper provides an affirmative answer to an open-question put forth in [27, Sec. 6], concerning the existence of a uniquely-solvable second-kind formulation—which involves only ItI maps, and no DtN maps.

As indicated in Section 3, the proposed hybrid direct/iterative strategy provides significant advantages over non-hybrid strategies in which either a fully iterative linear algebra solver is used, or a generic direct fast sparse solver such as [24] is utilized. Indeed, a fully iterative solver would necessarily require large numbers of iterations in order to account for the multiple scattering events that take place at boundaries of discontinuity of the refractive-index function n. As demonstrated in Section 4 (example 4.2), on the other hand, the coupling to the boundary integral solver destroys the sparsity inherent in the interior spectral matrix, and can thereby significantly hinder an overall direct solver strategy. The proposed hybrid strategy achieves the dual goal of maintaining a reduced iteration count (since the boundary integral equation, which requires reduced iteration numbers, is the only equation that is solved iteratively) while maintaining sparsity.

The overall proposed formulation can be used in conjunction with any adequate direct sparse linear algebra solver for the volumetric portion of the algorithm. If the specialized Helmholtz direct linear-algebra solver proposed in [27] were thus used, the resulting approach would accomplish three goals mentioned in that reference, namely 1) Use of an exterior solver based on the ItI (instead of the Dirichlet-to-Neumann map); 2) Employment of an overall formulation that is invertible for all frequencies; and 3) Use of an iterative strategy for the solution of the integral equation portion of the method. As indicated above, in this paper we utilize the Intel MKL implementation of the multifrontal solver Pardiso [12,43,44], which has shown to provide excellent performance, at nearly linear computing cost, to tackle the volumetric portion of the problem. In all, the proposed approach provides fast and essentially dispersionless solutions for high-frequency and/or high-contrast problems, with high-order accuracy for smooth refractivities, and it maintains second order accuracy for discontinuous refractive indexes n.

This paper is organized as follows. The proposed second-kind integro-differential formulation and the associated solution-uniqueness proof are presented in Section 2. Section 3 then presents a detailed description of the proposed algorithm, and Section 4 provides a variety of numerical results demonstrating the character of the proposed methodology. Concluding remarks, finally, are presented in Section 5.

2 Uniquely-solvable, second-kind integro-differential hybrid formulation

As discussed in the previous section, the proposed numerical method is based on a reformulation of the problem (1)-(3) as a combination of a differential equation formulation in a volumetric region and a boundary integral equation formulation on the boundary of the computational domain. To describe the method we consider an open bounded "computational" domain $\Omega \subset \mathbb{R}^2$ containing the inhomogeneity: $\overline{\mathcal{D}} \subset \Omega$. As mentioned in Section 1 and depicted on the right portion of Figure 1, throughout this paper the domain Ω is taken to equal a square for simplicity, but the algorithm can easily be generalized to allow for computational domains consisting of a union of disjoint square components tightly covering the region $\{n(x) \neq 1\}$. Then the complete problem (1)-(3) is reformulated in terms of two main elements: 1) A Helmholtz equation with variable coefficients in the volumetric region Ω, and; 2) A boundary integral equation on $\partial\Omega$ which couples the solution within Ω to the solution in the unbounded domain $\mathbb{R}^2 \setminus \overline{\Omega}$. In order to proceed with this plan the following section first discusses a certain impedance-to-impedance operators [27, 34, 35] associated with the Helmholtz problems in the interior and exterior of Ω.

2.1 Interior and Exterior Impedance-to-Impedance operators

Let $\Omega \subset \mathbb{R}^2$ denote a bounded open domain with a Lipschitz boundary $\partial\Omega$. Then, for each non-vanishing real constant β, the "exterior" impedance operator $T_{\text{ext}} : H^{-\frac{1}{2}}(\partial\Omega) \to H^{-\frac{1}{2}}(\partial\Omega)$ is defined by

$$T_{\text{ext}}[\psi](\mathbf{x}) = u_{\text{ext}}(\mathbf{x}) - i\beta \frac{\partial u_{ext}}{\partial \boldsymbol{\nu}}(\mathbf{x}), \tag{5}$$

where $\boldsymbol{\nu}$ is the unit outward normal vector at $\partial\Omega$ and where $u_{\text{ext}} \in H^1_{\text{loc}}\left(\mathbb{R}^2 \setminus \Omega\right)$ is the unique radiating solution of the exterior problem:

$$\begin{cases} \Delta u_{\text{ext}}(\mathbf{x}) + \kappa^2 u_{\text{ext}}(\mathbf{x}) = 0, & \text{if } \mathbf{x} \in \mathbb{R}^2 \setminus \overline{\Omega}, \\ u_{\text{ext}}(\mathbf{x}) + i\beta \frac{\partial u_{\text{ext}}}{\partial \boldsymbol{\nu}}(\mathbf{x}) = \psi(\mathbf{x}), & \text{if } \mathbf{x} \in \partial\Omega; \end{cases} \tag{6}$$

see [19, Theorem 2.3] and [39, Theorem 6.11] (cf. [33, Theorem 4.12] and [35, Sec. 3.2] where corresponding results for smooth boundaries are provided). The definition of the "interior" impedance operator $T_{\text{int}} : H^{-\frac{1}{2}}(\partial\Omega) \to H^{-\frac{1}{2}}(\partial\Omega)$ is analogous:

$$T_{\text{int}}[\phi](\mathbf{x}) = u_{\text{int}}(\mathbf{x}) - i\beta \frac{\partial u_{\text{int}}}{\partial \boldsymbol{\nu}}(\mathbf{x}), \tag{7}$$

where $u_{\text{int}} \in H^1(\Omega)$ is the unique solution of the problem

$$\begin{cases} \Delta u_{\text{int}}(\mathbf{x}) + \kappa^2 n^2(\mathbf{x}) u_{\text{int}}(\mathbf{x}) = 0, & \text{for } \mathbf{x} \in \Omega, \\ u_{\text{int}}(\mathbf{x}) + i\beta \frac{\partial u_{\text{int}}}{\partial \boldsymbol{\nu}}(\mathbf{x}) = \phi(\mathbf{x}), & \text{for } \mathbf{x} \in \partial\Omega. \end{cases} \tag{8}$$

2.2 Hybrid formulation

As is known [23, Theorem 2.5], any radiating solution u_{ext} of the Helmholtz equation over the exterior domain $\mathbb{R}^2 \setminus \overline{\Omega}$ may be represented by means of Green's formula

$$u_{\text{ext}}(\mathbf{x}) = \int_{\partial\Omega} \left(\frac{\partial G_\kappa(\mathbf{x}-\mathbf{y})}{\partial \boldsymbol{\nu}(\mathbf{y})} u_{\text{ext}}(\mathbf{y}) - G_\kappa(\mathbf{x}-\mathbf{y}) \frac{\partial u_{\text{ext}}}{\partial \boldsymbol{\nu}}(\mathbf{y}) \right) ds(\mathbf{y}), \quad \mathbf{x} \in \mathbb{R}^2 \setminus \overline{\Omega} \tag{9}$$

which, utilizing the jump relations [23, Theorem 3.1] of the single- and double-layer potentials on $\partial\Omega$ yields the relation

$$u_{\text{ext}}(\mathbf{x}) = \frac{u_{\text{ext}}(\mathbf{x})}{2} + \int_{\partial\Omega} \left(\frac{\partial G_\kappa(\mathbf{x}-\mathbf{y})}{\partial \boldsymbol{\nu}(\mathbf{y})} u_{\text{ext}}(\mathbf{y}) - G_\kappa(\mathbf{x}-\mathbf{y}) \frac{\partial u_{\text{ext}}}{\partial \boldsymbol{\nu}}(\mathbf{y}) \right) ds(\mathbf{y}), \quad \text{for} \quad \mathbf{x} \in \partial\Omega. \tag{10}$$

Similarly, an incident field u^i (a function that satisfies equation (1) throughout $\mathbb{R}^2$) satisfies

$$0 = \frac{u^i(\mathbf{x})}{2} + \int_{\partial\Omega} \left(\frac{\partial G_\kappa(\mathbf{x}-\mathbf{y})}{\partial \boldsymbol{\nu}(\mathbf{y})} u^i(\mathbf{y}) - G_\kappa(\mathbf{x}-\mathbf{y}) \frac{\partial u^i}{\partial \boldsymbol{\nu}}(\mathbf{y}) \right) ds(\mathbf{y}), \quad \text{for} \quad \mathbf{x} \in \partial\Omega. \tag{11}$$

In the case u_{ext} equals the scattered field u^s resulting from the incident field u^i, we may combine equations (10) and (11) and obtain the corresponding relation

$$\frac{u(\mathbf{x})}{2} - \int_{\partial\Omega} \left(\frac{\partial G_\kappa(\mathbf{x}-\mathbf{y})}{\partial \boldsymbol{\nu}(\mathbf{y})} u(\mathbf{y}) - G_\kappa(\mathbf{x}-\mathbf{y}) \frac{\partial u}{\partial \boldsymbol{\nu}}(\mathbf{y}) \right) ds(\mathbf{y}) = u^i(\mathbf{x}), \quad \text{for} \quad \mathbf{x} \in \partial\Omega, \tag{12}$$

for the total field $u = u^i + u^s$. Clearly, defining, for $\mathbf{x} \in \partial\Omega$, $\phi(\mathbf{x}) = \left(u(\mathbf{x}) + i\beta \frac{\partial u}{\partial \boldsymbol{\nu}}(\mathbf{x})\right)$ and

$$\mathcal{A}^{\text{int}}_{\text{ext}}[\phi](\mathbf{x}) = \int_{\partial\Omega} \left(\frac{1}{2} \frac{\partial G_\kappa(\mathbf{x}-\mathbf{y})}{\partial \boldsymbol{\nu}(\mathbf{y})} \left(I + T_{\text{int}}\right)[\phi](\mathbf{y}) - \frac{1}{2i\beta} G_\kappa(\mathbf{x}-\mathbf{y}) \left(I - T_{\text{int}}\right)[\phi](\mathbf{y}) \right) ds(\mathbf{y}), \tag{13}$$

equation (12) may be re-expressed in the form

$$\frac{1}{4}\left(I + T_{\text{int}}\right)[\phi](\mathbf{x}) - \mathcal{A}^{\text{int}}_{\text{ext}}[\phi](\mathbf{x}) = u^i(\mathbf{x}). \tag{14}$$

In particular it is easy to check that, given a solution u of (2)-(3) and defining $\phi = u + i\beta\frac{\partial u}{\partial \boldsymbol{\nu}}$ for $\mathbf{x} \in \partial\Omega$, the pair of functions (u, ϕ) is a solution of the problem

$$\Delta u(\mathbf{x}) + \kappa^2 n^2(\mathbf{x}) u(\mathbf{x}) = 0, \quad \text{if } \mathbf{x} \in \Omega, \tag{15}$$

$$\phi(\mathbf{x}) - \left(u(\mathbf{x}) + i\beta \frac{\partial u}{\partial \boldsymbol{\nu}}(\mathbf{x}) \right) = 0, \quad \text{if } \mathbf{x} \in \partial\Omega, \tag{16}$$

$$\frac{1}{4}\left(I + T_{\text{int}}\right)[\phi](\mathbf{x}) - \mathcal{A}_{\text{ext}}^{\text{int}}[\phi](\mathbf{x}) = u^i(\mathbf{x}) \;\; \text{for} \;\; \mathbf{x} \in \partial\Omega. \tag{17}$$

As shown in the following section, equation (17) (and, thus, the full problem (15)–(17)) is uniquely solvable—and the solution u must therefore coincide with the restriction to $\overline{\Omega}$ of the solution of the original inhomogeneous scattering problem (1)–(3). Once the solution u of (15)-(17) is obtained for $\mathbf{x} \in \overline{\Omega}$, the scattered field u^s (and hence the total field $u = u^i + u^s$) at any point $\mathbf{x} \in \mathrm{R}^2 \setminus \overline{\Omega}$ can be easily obtained by utilizing the representation formula (9) with $u_{\text{ext}} = u^s$. In other words, the hybrid integro-differential problem (15)–(17) is equivalent to the original inhomogeneous scattering problem (2)-(3), as claimed.

Remark 1. The density function ϕ, which, per Theorem 1 below, is the unique solution of equation (17), might in principle be expected to exhibit some sort of singularity at the corners of the square $\partial\Omega$; see e.g. [29, 48]. However, in view of (16), the solution ϕ under consideration is actually an infinitely differentiable (and, indeed, analytic) function along each one of the sides of the square $\partial\Omega$. This follows from the relation $\phi = u + i\beta\frac{\partial u}{\partial \boldsymbol{\nu}}$ and the fact that the solution u is infinitely smooth (and, in fact, analytic) in a certain neighborhood of $\partial\Omega$ within which the refractive index n is constantly equal to one.

2.3 Uniqueness

Theorem 1 (Uniqueness of solution for the second-kind hybrid volume-boundary formulation)**.** *Let* $\phi \in H^{-\frac{1}{2}}(\partial\Omega)$ *denote a solution of equation* (17) *with* $u^i = 0$*. Then* $\phi = 0$*.*

Proof. Letting $u_{\text{ext}} \in H^1_{\text{loc}}\left(\mathbb{R}^2 \setminus \Omega\right)$ denote the radiating solution of (6) corresponding to the impedance data $\psi = \phi$ on $\partial\Omega$, the Green relation (10), with integral expressions interpreted as in [39, Thm. 4.4], may be re-expressed in the form

$$\frac{1}{4}\left(I + T_{\text{ext}}\right)[\phi](\mathbf{x}) - \mathcal{A}_{\text{ext}}^{\text{ext}}[\phi](\mathbf{x}) = 0 \;\; \text{for} \;\; \mathbf{x} \in \partial\Omega, \tag{18}$$

where

$$\mathcal{A}_{\text{ext}}^{\text{ext}}[\phi](\mathbf{x}) = \int_{\partial\Omega} \left(\frac{1}{2}\frac{\partial G_\kappa(\mathbf{x}-\mathbf{y})}{\partial \boldsymbol{\nu}(\mathbf{y})}\left(I + T_{\text{ext}}\right)[\phi](\mathbf{y}) - \frac{1}{2i\beta} G_\kappa(\mathbf{x}-\mathbf{y})\left(I - T_{\text{ext}}\right)[\phi](\mathbf{y}) \right) ds(\mathbf{y}). \tag{19}$$

Equation (17) with $u^i = 0$ and (18) can be recast in the forms

$$\frac{T_{\text{int}}[\phi](\mathbf{x})}{2} - \int_{\partial\Omega} \left(\frac{\partial G_\kappa(\mathbf{x}-\mathbf{y})}{\partial \boldsymbol{\nu}(\mathbf{y})} - i\eta G_\kappa(\mathbf{x}-\mathbf{y}) \right) T_{\text{int}}[\phi](\mathbf{y}) ds(\mathbf{y}) = f(\mathbf{x}), \quad \mathbf{x} \in \partial\Omega, \tag{20}$$

$$\frac{T_{\text{ext}}[\phi](\mathbf{x})}{2} - \int_{\partial\Omega} \left(\frac{\partial G_\kappa(\mathbf{x}-\mathbf{y})}{\partial \boldsymbol{\nu}(\mathbf{y})} - i\eta G_\kappa(\mathbf{x}-\mathbf{y}) \right) T_{\text{ext}}[\phi](\mathbf{y}) ds(\mathbf{y}) = f(\mathbf{x}), \quad \mathbf{x} \in \partial\Omega, \tag{21}$$

where

$$f(\mathbf{x}) = -\frac{\phi(\mathbf{x})}{2} + \int_{\partial\Omega} \left(\frac{\partial G_\kappa(\mathbf{x}-\mathbf{y})}{\partial \boldsymbol{\nu}(\mathbf{y})} + i\eta G_\kappa(\mathbf{x}-\mathbf{y}) \right) \phi(\mathbf{y}) ds(\mathbf{y}), \tag{22}$$

and where $\eta = 1/\beta$. Clearly, equations (20) and (21) are identical combined field integral equation of second kind, with the same right hand side, for the unknowns $T_{\text{ext}}[\phi]$ and $T_{\text{int}}[\phi]$, respectively. Since, as is well known [23, p. 51], the combined field integral equation admits unique solutions, it follows that $T_{\text{int}}[\phi] = T_{\text{ext}}[\phi]$ or, equivalently,

$$u_{\text{int}}(\mathbf{x}) - i\beta \frac{\partial u_{\text{int}}}{\partial \boldsymbol{\nu}}(\mathbf{x}) = u_{\text{ext}}(\mathbf{x}) - i\beta \frac{\partial u_{\text{ext}}}{\partial \boldsymbol{\nu}}(\mathbf{x}) \quad \text{on} \quad \partial\Omega, \tag{23}$$

where u_{int} is the solution of (8). But, from (8) and (6) we have

$$u_{\text{int}}(\mathbf{x}) + i\beta \frac{\partial u_{\text{int}}}{\partial \boldsymbol{\nu}}(\mathbf{x}) = \phi(\mathbf{x}) = \psi(\mathbf{x}) = u_{\text{ext}}(\mathbf{x}) + i\beta \frac{\partial u_{\text{ext}}}{\partial \boldsymbol{\nu}}(\mathbf{x}) \qquad \text{on } \partial\Omega, \tag{24}$$

and, therefore, using (23) it follows that

$$u_{\text{int}}(\mathbf{x}) = u_{\text{ext}}(\mathbf{x}) \quad \text{and} \quad \frac{\partial u_{\text{int}}}{\partial \boldsymbol{\nu}}(\mathbf{x}) = \frac{\partial u_{\text{ext}}}{\partial \boldsymbol{\nu}}(\mathbf{x}) \quad \text{on} \quad \partial\Omega. \tag{25}$$

Let us now define

$$U_\phi(\mathbf{x}) = \begin{cases} u_{\text{int}}(\mathbf{x}) & \text{for } \mathbf{x} \in \overline{\Omega} \\ u_{\text{ext}}(\mathbf{x}) & \text{for } \mathbf{x} \in \mathbb{R}^2 \setminus \overline{\Omega}. \end{cases} \tag{26}$$

Since u_{ext} is the radiating solution of (6) and u_{int} is the solution of (8), on account of (25) it follows that U_ϕ is the radiating solution of the Helmholtz problem (1)-(3) throughout $\mathbb{R}^2$ with $u^i = 0$. Since this problem admits a unique solution in $H^2_{\text{loc}}(\mathbb{R}^2)$ [23, Theorem 8.7] we conclude that U_ϕ vanishes identically. In particular, it follows that $u_{\text{int}} = 0$ throughout $\overline{\Omega}$ and, thus, $\phi = 0$ in $\partial\Omega$ in view of (24). The proof is now complete. □

Having established the well posedness of the second-kind hybrid formulation (15)-(17) we now present, in the next section, the proposed numerical algorithm for the solution of this problem.

3 Numerical algorithm

The proposed algorithm relies on the formulation (15)–(17) in a computational domain Ω which, for definiteness, throughout this paper is taken to equal the square $\Omega = (-a, a)^2$ with a value of a selected in such a way that $\overline{\mathcal{D}} \subset \Omega$. The algorithm consists of two main components, namely 1) A spectral volumetric solver of fixed order of accuracy for the Boundary Value Problem (BVP) (8) in the domain Ω for given impedance data $\phi \in H^{-1/2}(\partial\Omega)$; and 2) A solver for the boundary integral equation (17) on $\partial\Omega$, which couples the solution within Ω to the solution in the exterior of that domain. In order to achieve second-order convergence for discontinuous scatterers the algorithm utilizes a filtered Fourier-smoothing technique outlined in Section 3.1.1. The overall hybrid approach is completed via an application of the iterative solver GMRES, as detailed in Section 3.3. As mentioned in Section 1, the overall hybrid method meets the dual goals of achieving reduced iteration numbers while maintaining the sparsity of the spectral matrix.

3.1 Volumetric boundary-value solver

This section describes our discretization and direct solution strategy for the BVP (8) for given values of the impedance ϕ on $\partial\Omega$. The presentation includes two subsections, covering the proposed filtered Fourier smoothing technique that enables second-order convergence even for discontinuous scatterers (Section 3.1.1), and the Chebyshev-based volumetric discretization used (Section 3.1.2).

3.1.1 Filtered Fourier smoothing (FFS) of discontinuous refractivities

As is well known, discontinuities in the refractive-index $n(\mathbf{x})$ give rise to severe restrictions on the order of accuracy of the numerical solutions of the scattering problem (1)-(3): in such cases only first-order accuracy is generally obtained. In the context of the volumetric Lippmann-Schwinger integral-equation solvers, however, Reference [31] shows that full second order convergence can be reinstated for such problems by means of an application of a certain Fourier-smoothing technique [17,31]. In detail, a quadratic convergence rate toward the solution for the exact refractivity $n(\mathbf{x})$ results in that context as the discontinuous contrast function $m(\mathbf{x}) = 1 - n^2(\mathbf{x})$ is replaced by truncations of its Fourier series of certain orders, with the additional requirement that sufficiently accurate values of the Fourier coefficients for the exact discontinuous function $m(\mathbf{x})$ be used; see Remark 2 below. Since (15)–(17) is equivalent to the corresponding Lippmann-Schwinger problem, the same conclusions hold in our present spectral context as well. In what follows we present a new version of the Fourier-smoothing approach, which, incorporating a new filtering component that eliminates a certain erratic convergence behavior in the un-filtered approach (see Table 8), is then applied to the differential equations considered in this paper. The properties of the resulting Filtered Fourier Smoothing (FFS) method are demonstrated in practice via a variety of numerical results in Section 4.

To introduce the method, letting $m = 1 - n^2(\mathbf{x})$ we re-express the Helmholtz equation (2) in the form

$$\Delta w(\mathbf{x}) + \kappa^2(1 - m(\mathbf{x}))w(\mathbf{x}) = 0; \tag{27}$$

in our context the resulting procedure will be applied to the problem (8) to obtain the intermediate solutions $w = u_{\text{int}}$, and, once convergence has been achieved for the impedance data ϕ, to produce the corresponding solution $w = u$ of (15)–(16).

As is well known, the Fourier series of the (possibly discontinuous) function $m(\mathbf{x})$ converges uniformly to $m(\mathbf{x})$ except on the discontinuity set, around which it suffers the well known Gibbs-ringing artifact. Assuming, for notational simplicity, a square domain Ω of side $2a$, the FFS approach proposed in this section utilizes the order-F *filtered* truncated Fourier expansion

$$m^F(\mathbf{x}) = \sum_{\ell_1=-F}^{F} \sum_{\ell_2=-F}^{F} c_{\ell_1,\ell_2} e^{\frac{\pi i}{a}(\ell_1 x_1 + \ell_2 x_2)} \tag{28}$$

of the $2a$-biperiodic Fourier series of m in Ω, where $\mathbf{x} = (x_1, x_2)$ and where the filtered Fourier coefficient c_{ℓ_1,ℓ_2} are given by

$$c_{\ell_1,\ell_2} = \left(\frac{1}{4a^2} \int_{-a}^{a} \int_{-a}^{a} m(x_1, x_2) e^{-\frac{\pi i}{a}(\ell_1 x_1 + \ell_2 x_2)} dx_1 dx_2 \right) \exp\left(-\alpha \left(\left(\frac{2\ell_1}{F} \right)^{2p} + \left(\frac{2\ell_2}{F} \right)^{2p} \right) \right). \tag{29}$$

Here, p and α are the parameters in the exponential filter used; following [3], throughout this paper the values $p = 4$ and $\alpha = 16 \log 10$ have been used.

Remark 2. Note that, for a discontinuous function m, evaluation of the integral (29) via an FFT would yield only first-order accurate coefficients—and would ultimately reduce the accuracy the overall solver to first order. A fast ($O(F^2 \log F)$) algorithm for highly accurate evaluation of these coefficients follows from application of the (one-dimensional) FC-based integration method presented in Appendix A to the integral (29) in the x_1 and x_2 directions.

An additional difficulty associated with the smoothing algorithm still needs to be tackled since, unlike the algorithm [31], our strategy relies on use of non-equispaced (Chebyshev) volumetric grids, and, therefore, a straightforward evaluation of the Fourier series of the function $m^F(\mathbf{x})$ at the required N discretization points, for which an FFT cannot be directly employed, generally requires an $O(NF^2)$ computational cost. Since, generically, $F^2 = O(N)$, the overall $O(N^2)$ cost of the straightforward approach is unacceptably large within our scheme. One can easily expedite this computation, however, by means of the FFT-refined trigonometric polynomial interpolation method presented in [15], which yields high-order accuracy while maintaining computational efficiency. In our context, once accurate values of the Fourier coefficients c_{ℓ_1,ℓ_2} have somehow been obtained, this interpolation approach can be performed as a two-step procedure:

1. Evaluate the Fourier series $m^F(\mathbf{x})$ on a sufficiently fine equispaced refinement of the associated F^2-point FFT grid. (In our examples the fine FFT grid is finer than the original grid by a factor of four in each dimension.) This step can be performed by means of an FFT on a zero-padded version of the sum (28), at a cost of $O(F^2 \log F)$ operations.

2. In order to evaluate $m^F(\mathbf{x}_0)$ for $\mathbf{x}_0 = (x_0, y_0) \in \overline{\Omega}$, obtain the value of $m^F(\mathbf{x})$ at a number R of points neighboring $\mathbf{x}_0$ in the fine grid mentioned in point 1., and interpolate to $\mathbf{x}_0$ by means of iterated one dimensional polynomial interpolation; see e.g. [15]. (In our examples we have used fifth order Lagrange polynomial interpolation.)

This procedure yields interpolating polynomials that accurately reproduce the exact values of the truncated Fourier series at an $O(N \log N)$ computational cost.

3.1.2 Volumetric discretization

This section presents the proposed direct solution strategy for the numerical solution of the BVP (8). As discussed in Section 3.1.1, discontinuities in the refractive index $n(\mathbf{x})$, if any, are dealt with by utilizing the modified BVP

$$\Delta w(\mathbf{x}) + \kappa^2 \left(1 - m^F(\mathbf{x})\right) w(\mathbf{x}) = 0, \quad \text{if } \mathbf{x} \in \Omega, \tag{30}$$

$$w(\mathbf{x}) + i\beta \frac{\partial w}{\partial \boldsymbol{\nu}}(\mathbf{x}) = \phi(\mathbf{x}) \quad \text{if } \mathbf{x} \in \partial\Omega \tag{31}$$

instead of the BVP (8)—a procedure that, according to [17, Corollary 3.9] (cf. also [31]) leads to second-order accurate approximations to the actual solutions of the original problem (8) instead of the first-order convergence that would otherwise result. (Note that, interestingly, the proof and illustrations presented in [17, Corollary 3.9] and [31] are given in the context of integral formulations of the problem. But, since the integral-equation and PDE solutions for the Fourier-smoothed problem coincide, the improved approximation order carries over, as indicated above and demonstrated in Section 4, to the present differential formulation.)

For the discussion in the present section we assume that the impedance data ϕ in equation (31) is known on $\partial\Omega$. We wish to utilize a general-purpose fast sparse direct solver, such as, e.g., the multifrontal algorithm [1, 12, 24], for the solution of our discrete version of (30)-(31). Naturally, the performance of sparse linear-algebra solvers is highly dependent on the sparsity pattern of the coefficient matrix of the linear system. In view of this fact, we seek to approximate all necessary differential operators in such a way that the resulting linear system is as sparse as possible while maintaining essentially dispersionless approximations and higher order accuracy.

To do this we approximate the unknown function w and its derivatives by means of local Chebyshev representations. In detail, assuming, for notational simplicity, a square computational domain Ω, the proposed BVP solver proceeds by first splitting Ω into a total of $P \times P$ mutually disjoint square patches $\Omega_{i,j}, 1 \leq i, j \leq P$, such that

$$\overline{\Omega} = \bigcup_{i,j=1}^{P} \overline{\Omega}_{i,j}.$$

Then, the solution of (30)-(31) is obtained by solving the equivalent set of coupled transmission problems

$$\Delta w_{i,j}(\mathbf{x}) + \kappa^2(1 - m^F(\mathbf{x}))w_{i,j}(\mathbf{x}) = 0, \quad \text{if } \ \mathbf{x} \in \Omega_{i,j}, \tag{32}$$

$$w_{i,j}(\mathbf{x}) = w_{r,s}(\mathbf{x}) \quad \text{and} \quad \frac{\partial w_{i,j}}{\partial \boldsymbol{\nu}_{i,j}}(\mathbf{x}) = \frac{\partial w_{r,s}}{\partial \boldsymbol{\nu}_{r,s}}(\mathbf{x}) \quad \text{if } \ \mathbf{x} \in (\Gamma_{i,j} \cap \Gamma_{r,s}) \setminus \partial\Omega, \tag{33}$$

$$w_{i,j}(\mathbf{x}) + i\beta \frac{\partial w_{i,j}}{\partial \boldsymbol{\nu}_{i,j}}(\mathbf{x}) = \phi(\mathbf{x}), \quad \text{if } \ \mathbf{x} \in \Gamma_{i,j} \cap \partial\Omega, \tag{34}$$

$(1 \leq i, r \leq P, 1 \leq j, s \leq P,)$, where $w_{i,j} = w\big|_{\Omega_{i,j}}$, and where $\boldsymbol{\nu}_{i,j}$ denotes the outward unit normal vector for the domain $\Omega_{i,j}$ on the boundary $\Gamma_{i,j} = \partial\Omega_{i,j}$. For any pair of patches that share a common boundary, the conditions (33) amount to a manifestation, valid for smooth solutions, of the (uniquely solvable) weak formulation of equations (30)-(31) in a multi-patch decomposition; see e.g. equations (1.7) and (1.8) in reference [35].

To obtain the desired solutions, for a given positive integer q we discretize the closure $\overline{\Omega}_{i,j} = [a_{i-1}, a_i] \times [b_{j-1}, b_j]$ of the patch $\Omega_{i,j}$ by means of the two-dimensional tensor product $\mathcal{N}_{i,j} = \left\{\mathbf{x}_{i,j,k,\ell} \middle| \, 0 \leq k, \ell \leq q\right\}$ Chebyshev mesh given by

$$\mathbf{x}_{i,j,k,\ell} = \left(\frac{a_{i-1} + a_i}{2} + \frac{a_i - a_{i-1}}{2}\cos\left(\frac{\pi k}{q}\right), \frac{b_{j-1} + b_j}{2} + \frac{b_j - b_{j-1}}{2}\cos\left(\frac{\pi \ell}{q}\right)\right).$$

Equations for the unknown values of $w_{i,j}(\mathbf{x})$ at the grid points $\mathbf{x} = \mathbf{x}_{i,j,k,\ell}$ $(1 \leq i, j \leq P,\ 0 \leq k \leq q, 0 \leq \ell \leq q)$ are obtained by enforcing discrete versions of equations (32), (33), and (34), as appropriate, at the discretization points $\mathbf{x}_{i,j,k,\ell}$ (see Remark 3), via approximation of the necessary differential operators $\partial/\partial\boldsymbol{\nu}_{i,j}$ and Δ by Chebyshev spectral differentiation matrices local to the relevant patch(es) $\Omega_{i,j}$. These Chebyshev-based approximations of derivatives remain accurate even for large wavenumbers, and, when used for discretization of the joint transmission problem (32)–(34), they give rise to a sparse linear systems of the form

$$\mathbf{A}\mathbf{w} = \mathbf{b}, \tag{35}$$

where the entries of the right hand side vector $\mathbf{b}$ associated with observation points $\mathbf{x}_{i,j,k,\ell} \in \Omega$ equal zero, and where the entries corresponding to boundary points $\mathbf{x}_{i,j,k,\ell} \in \partial\Omega$ equal $\phi\left(\mathbf{x}_{i,j,k,\ell}\right)$. The unknown vector $\mathbf{w}$, on the other hand, contains the N unknowns $w_{i,j,k,\ell}$, one corresponding to each point $\mathbf{x}_{i,j,k,\ell}$, where

$$N = (q+1)^2 P^2; \tag{36}$$

note that, in particular, different unknowns are used at single discretization points that are common to two subdomain boundaries. Owing to its sparse nature, this linear system is suitable for treatment by sparse linear solvers such as e.g. the multifrontal-based direct solver [1,12,24].

Remark 3. As illustrated in Figure 2, the only non-zero entries in the equation associated with the point $\mathbf{x}_{i,j,k,\ell}$ correspond to discretization points lying on the grid lines that pass through $\mathbf{x}_{i,j,k,\ell}$. Two separate unknowns are used at each patch-boundary discretization point, which are then set to be equal as part of the equation system. Similarly, four separate unknowns are used at each patch corner point that is not on $\partial\Omega$, and two separate unknowns are used at each patch corner point that is on $\partial\Omega$. This strategy is used so as to render each patch discretization independent of all other patch discretizations at a minimal increase in the number of unknowns. A question arises as to which of the two possible enforcements of the matching normal derivative conditions in (33), either using horizontal or vertical normal derivatives, are used at corner points. The indeterminacy is resolved in our algorithm by means of the arbitrary but acceptable selection of horizontal normal derivatives in (33) at all corner points.

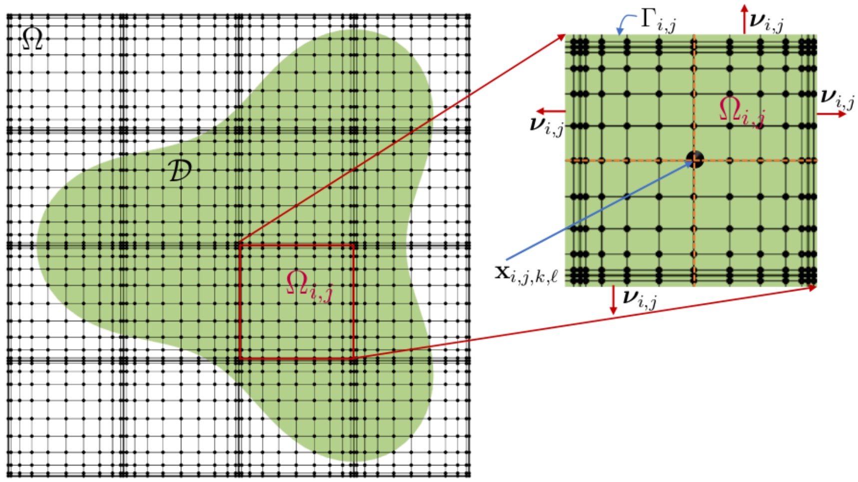

Figure 2: Domain-partitioning setup: the computational domain Ω containing the inhomogeneity $\mathcal{D}$ is split into $P \times P$ Chebyshev patches $\Omega_{i,j}$, $1 \leq i,j \leq P$, (with $P = 4$ in this illustration). Derivatives at a given point $\mathbf{x}_{i,j,k,\ell} \in \Omega_{i,j}$ are evaluated as derivatives of the Chebyshev expansions obtained from function values along the lines passing through $\mathbf{x}_{i,j,k,\ell}$. The unit normal vector $\boldsymbol{\nu}_{i,j}$ on the boundary $\Gamma_{i,j}$ of the patch $\Omega_{i,j}$ points to the exterior of the patch.

The discrete version of the impedance quantity $T_{\mathrm{int}}[\phi]$ (equation (7)) which, for a given ϕ, is necessary as part of the proposed iterative algorithm for the solution of (32)-(34) (see Section (3.3)), can readily be obtained by differentiating the solution of the linear system (35) on the basis of the Chebyshev representations introduced in Section 3.1.2. It is useful to note that, as indicated in Section 3.3, solutions of the system (35) with various right-hand sides (one solution per iteration) are required as part of the proposed iterative scheme. To obtain the necessary solutions at a reduced computing cost, in our algorithm the **LU** factorization of the sparse matrix **A**, which is obtained by means of the efficient implementation MKL Pardiso of the multi-frontal sparse solver [1, 12], is computed once and stored for repeated use in multiple GMRES iterations, or, even, for multiple right-hand sides. The computational cost of assembly of the matrix **A** and evaluation of its **LU** factorization, whose combination amounts to the most expensive portion of the overall hybrid volumetric solver, is studied in Section 4. In particular, Figure 4 and Table 4 in that section demonstrate a computing cost of $O(N^{\alpha})$ operations, with $\alpha \approx 1.07$, for this portion of the algorithm, with a total number of the order of $O(qN)$ of non-zero matrix entries.

3.2 High-Order Approximation of the Boundary Integral Operators

The proposed algorithm utilizes a fast, high-order Nyström integration algorithm for the evaluation of the integral operators in equation (17). We note without a detailed proof that, in view of the smoothness of the solution ϕ in a neighborhood of $\partial\Omega$ (Remark 1) together with stability theory (see e.g. [36, Th. 10.2]), the Chebyshev approximation of the density ϕ that we utilize in this section gives rise to high-order accuracy in the overall algorithm—as illustrated numerically in Table 6. In what follows we describe the associated integration scheme and certain connections with the overall volumetric iterative solver of which it is a component. Clearly, it is desirable for the underlying grid in the approximation of (17) to be a subset of the volumetric (piece-wise Chebyshev) interior grid: otherwise an additional fast and accurate interpolation procedure would be required for the evaluation of the integral density to the underlying quadrature points. To avoid such additional difficulties while preserving maximal accuracy, we use a two-dimensional analog of the rectangular polar integration scheme recently introduced in [14] for the solution of surface scattering problems in the three dimensions. The resulting procedure is described in what follows.

In a first stage, the entire integration domain $\partial\Omega$ is covered by a set of non-overlapping boundary patches $\{\gamma_p\}_{p=1}^{P}$ ($P = 4P$), each one of which is the image of the interval $[-1,1]$ via a smooth invertible mapping $\boldsymbol{\xi}_p$. Using this covering and the parameterizations $\boldsymbol{\xi}_p$, the integral operator that is used as part of (17) can be decomposed in the form

$$\int_{\partial\Omega}\left(\frac{\partial G_\kappa(\mathbf{x}-\mathbf{y})}{\partial\boldsymbol{\nu}(\mathbf{y})}\eta(\mathbf{y}) - G_\kappa(\mathbf{x}-\mathbf{y})\zeta(\mathbf{y})\right)ds(\mathbf{y}) = \sum_{p=1}^{P} I_p(\mathbf{x}),$$

where

$$I_p(\mathbf{x}) = \int_{-1}^{1}\left(\frac{\partial G_\kappa\left(\mathbf{x}-\boldsymbol{\xi}_p(t)\right)}{\partial\boldsymbol{\nu}\left(\boldsymbol{\xi}_p(t)\right)}\eta\left(\boldsymbol{\xi}_p(t)\right) - G_\kappa\left(\mathbf{x}-\boldsymbol{\xi}_p(t)\right)\zeta\left(\boldsymbol{\xi}_p(t)\right)\right)\left|\frac{\partial\boldsymbol{\xi}_p(t)}{\partial t}\right| dt. \tag{37}$$

An adequate choice of a methodology for the accurate evaluation of (37) depends on the relative position of the target point $\mathbf{x}$ with respect to the integration patch γ_p. If the target point $\mathbf{x}$ is sufficiently far from γ_p then the integrand in (37) is smooth and can be integrated with high-order accuracy by means of any high-order quadrature rule. On the other hand, if the target point is either close to or within the integration patch, the integrand is either singular or near singular, and hence a specialized quadrature rule must be used for its accurate evaluation. Thus, depending upon the distance from the target point to the integration patch, the overall integration approach relies on three different methods:

Evaluation of non-singular integrals: For target points $\mathbf{x}$ sufficiently far from the integration patch we use the Clenshaw-Curtis quadrature which, as is known, enjoys high-order convergence for smooth integrands [46], and whose discretization is taken to coincide with the restriction to γ_p of the volumetric discretization $\cup_{i,j=1}^{P}\mathcal{N}_{i,j}$.

Evaluation of singular integrals: For target points $\mathbf{x}$ in the integration patch, the accurate approximation of (37) becomes challenging in view of the integrand singularity. To deal with this difficulty, we first replace the density functions η and ζ in (37) by their Chebyshev expansions and we thus obtain

$$I_p(\mathbf{x}) = \sum_{\ell=0}^{M} c_\ell I_{p,\ell}^1(\mathbf{x}) + \sum_{\ell=0}^{M} d_\ell I_{p,\ell}^2(\mathbf{x}), \tag{38}$$

where

$$I_{p,\ell}^1(\mathbf{x}) = \int_{-1}^{1}\frac{\partial G_\kappa\left(\mathbf{x}-\boldsymbol{\xi}_p(t)\right)}{\partial\boldsymbol{\nu}\left(\boldsymbol{\xi}_p(t)\right)}T_\ell(t)\left|\frac{\partial\boldsymbol{\xi}_p(t)}{\partial t}\right| dt, \tag{39}$$

$$I_{p,\ell}^2(\mathbf{x}) = \int_{-1}^{1} G_\kappa\left(\mathbf{x}-\boldsymbol{\xi}_p(t)\right)T_\ell(t)\left|\frac{\partial\boldsymbol{\xi}_p(t)}{\partial t}\right| dt, \tag{40}$$

and where T_ℓ is the Chebyshev polynomial of degree ℓ. The Chebyshev coefficients c_ℓ, d_ℓ can be obtained accurately and efficiently by means of FFTs. Note that the integrals in equations (39) and (40) do not depend on the density, and therefore, may be computed only once and stored for repeated use. In addition to this, evaluation of these integrals does not require interpolation, even if refined meshes are used for their evaluation, as the corresponding integrands are known analytically in the complete domain of integration. However, evaluation of these integrals present certain difficulties owing to the weakly singular character of the integral kernel. To resolve the integrand singularity in equations (39) and (40) we utilize changes of variable whose Jacobian vanishes along with several of its derivatives at the singularity point. The idea is not limited to the specific kernel presently under consideration, and it can be readily incorporated for a general class of weakly singular kernels. Thus, we present our discussion in that general context.

Letting

$$I_\ell(\mathbf{x}) = \int_{-1}^{1} H_\kappa\left(\boldsymbol{\xi}_p(t_0)-\boldsymbol{\xi}_p(t)\right)T_\ell(t)\left|\frac{\partial\boldsymbol{\xi}_p(t)}{\partial t}\right| dt, \tag{41}$$

where $\mathbf{x} = \boldsymbol{\xi}_p(t_0)$ and where $H_\kappa\left(\boldsymbol{\xi}_p(t_0) - \boldsymbol{\xi}_p(t)\right)$ is any weakly singular kernel, we re-express I_ℓ in the form

$$I_\ell(\mathbf{x}) = \int_{-1}^{t_0} H_\kappa\left(\boldsymbol{\xi}_p(t_0) - \boldsymbol{\xi}_p(t)\right) T_\ell(t) \left|\frac{\partial \boldsymbol{\xi}_p(t)}{\partial t}\right| dt + \int_{t_0}^{1} H_\kappa\left(\boldsymbol{\xi}_p(t_0) - \boldsymbol{\xi}_p(t)\right) T_\ell(t) \left|\frac{\partial \boldsymbol{\xi}_p(t)}{\partial t}\right| dt. \tag{42}$$

Both the first and second integrands in (42) are singular at $t = t_0$. To resolve the singularity we use the changes of variables [23]

$$t = t_0 - \frac{1+t_0}{\pi}\omega_k\left[\frac{\pi}{2}(-\tau+1)\right] \quad \text{and} \quad t = t_0 + \frac{1-t_0}{\pi}\omega_k\left[\frac{\pi}{2}(\tau+1)\right]$$

in the first and second integrals in (42), respectively, which, roughly speaking, distributes half of the discretization points near the singular point t_0, and the other half fairy uniformly throughout the integration interval [23, p. 84]. Here, for $0 \le s \le 2\pi$ and for a given integer $k > 1$ we have set

$$\omega_k(s) = 2\pi\frac{[v(s)]^k}{[v(s)]^k + [v(2\pi - s)]^k}, \qquad \text{where} \quad v(s) = \left(\frac{1}{k} - \frac{1}{2}\right)\left(\frac{\pi - s}{\pi}\right)^3 + \frac{1}{k}\left(\frac{s-\pi}{\pi}\right) + \frac{1}{2}.$$

It is easy to check that the Jacobians of these changes of variables vanish up to order $k-1$ at the singular point $t = t_0$, which renders smooth integrands that can be integrated with high-order accuracy by means of the Clenshaw-Curtis quadrature.

Evaluation of near-singular integrals: This case arises when the target point $\mathbf{x}$ is "very close" to, but outside the integration patch γ_p. In this case, while the integrand in (37) is, strictly speaking, non-singular, its numerical integration poses similar challenges to the singular case. To effectively treat this issue we project the target point to the closest point to it on the integration patch and then follow the same strategy used for singular integration by treating the projection point as the singular point.

3.3 Overall hybrid solver

As discussed in the Section 2, the proposed method obtains the solution of the scattering problem (2)-(3) by solving the equivalent formulation (15)–(17). If the impedance data ϕ in (16) were known on $\partial\Omega$ then the solution of the scattering problem (2)-(3) could be readily obtained by solving the BVP (15)-(16) using the direct solution algorithm discussed in Section 3.1. To obtain ϕ on $\partial\Omega$, equation (17) is solved iteratively, where, for each iteration, the integral operator $\mathcal{A}^{\mathrm{int}}_{\mathrm{ext}}[\phi]$ (defined in (13)) is evaluated via the algorithm discussed in Section 3.2 in conjunction with the direct solution technique presented in Section 3.1 for the evaluation of the interior impedance operator $T_{\mathrm{int}}[\phi]$ (equation (7)). Note that, per the first three sentences in Remark 3 and in view of (36), the sparse $N \times N$ matrix $\mathbf{A}$ associated with the interior problem (equation (35)) contains only $2(q+1)N$ non-zero entries.

The main lines of the proposed overall hybrid solver are as follows:

1. Replace the discontinuous refractivity $n^2(\mathbf{x})$ in equation (15) by its filtered Fourier-smoothed version $1 - m^F(\mathbf{x})$ as discussed in Section 3.1.1.

2. Using either an initial guess (e.g. $\phi = u^i(\mathbf{x})$) or any improved guess for ϕ produced by the linear-algebra solver GMRES, obtain the solution $u = u_{\mathrm{int}}$ of the problem (15)-(16), and then use equation (7) to evaluate $T_{\mathrm{int}}[\phi](\mathbf{x})$, and, thus, $(I + T_{\mathrm{int}})[\phi]$ and $(I - T_{\mathrm{int}})[\phi]$ on $\partial\Omega$.

3. Evaluate the left hand side of equation (17), on the discretization of $\partial\Omega$, by applying the methods in Section 3.2 to integral densities equal to $(I + T_{\mathrm{int}})[\phi]$ and $(I - T_{\mathrm{int}})[\phi]$.

4. Pass the resulting residual (equal to the difference between the left-hand and the right-hand sides in (17)) to the GMRES algorithm, to obtain a new approximation for the density ϕ.

5. Check for convergence of the density ϕ to a given prescribed residual tolerance, and iterate by returning to step 2 until convergence is achieved.

6. Solve the BVP (15)-(16) for the converged impedance function ϕ obtained per point 5. If desired, use equation (9) with $u_{\rm ext} = \frac{1}{2}(I + T_{\rm int})[\phi]$ and $\partial u_{\rm ext}/\partial\boldsymbol{\nu} = \frac{1}{2i\beta}(I - T_{\rm int})[\phi]$ to produce u in the exterior of Ω and/or, using the Green function asymptotics [23, Theorem 2.5], far field values for the solution u.

4 Numerical results

This section presents results of numerical tests and examples that demonstrate the performance of the scattering solvers introduced in the Section 3, with an emphasis on problems containing discontinuous refractivities. All numerical results presented in this section were produced by means of a C++ implementation of the algorithms described in Section 3 on a single core of an Intel i7-4600M processor. The relative error (in the near field) reported here was computed according to the expression

$$\varepsilon_\infty^N = \frac{\max\limits_{1\leq i\leq N} \left|u^{\rm ref}(\mathbf{x}_i) - u^{\rm approx}(\mathbf{x}_i)\right|}{\max\limits_{1\leq i\leq N} \left|u^{\rm ref}(\mathbf{x}_i)\right|},$$

where $\{\mathbf{x}_i \in \Omega : 1 \leq i \leq N\}$ is a listing of all volumetric Chebyshev discretization points $\mathbf{x}_{i,j,k,\ell}$ considered in Section 3.1.2, over all subdomains $\Omega_{i,j}$, and where $u^{\rm ref}$ is either a closed form solution, when available, or a highly accurate numerical solution produced by the proposed algorithm on a fine discretization. GMRES tolerances were prescribed in each case to achieve the desired solution error. Values of the coupling parameter β in the range $10^{-5} \leq \beta \leq 10^{-3}$ were typically used: as shown in Figure 3 use of such values of β suffices to eliminate difficulties arising from resonance. (Typically smaller values of β tend to give rise to smaller numbers of iterations, while slightly larger values of β can result in somewhat higher accuracies; we have found that use of large values of β, say, in the range $1 \leq \beta \leq 100$, however, can significantly increase the iteration numbers required to meet a prescribed GMRES tolerance and/or solution accuracy.) In all of the tabulated results the acronyms "numIt" and "Order" denote the number of GMRES iterations required to achieve the desired accuracy and the numerical order of convergence $\log\left(\varepsilon_\infty^N/\varepsilon_\infty^{2N}\right)/\log(2)$ respectively. In accordance with Section 3.1.2, $P \times P$ denotes the total number of Chebyshev patches used in the discretization of the computational domain Ω, each one of which contains $q \times q$ discretization points; cf. Figure 2.

Example 4.1. *(High-order Convergence for the Boundary Integral Operator)*

This example illustrates the high-order convergence of the singular integration technique introduced in Section 3.2. For our example we let $\kappa = 5\pi$, $v(\mathbf{x}) = u^i(\mathbf{x}) = e^{i\kappa x_1}$ and $\overline{\Omega} = \{(x_1, x_2)| -1.5 \leq x_1, x_2 \leq 1.5\}$, and we evaluate numerically the integral

$$2\int\limits_{\partial\Omega} \left\{ G_\kappa(\mathbf{x}-\mathbf{y})\frac{\partial v(\mathbf{y})}{\partial\boldsymbol{\nu}(\mathbf{y})} - \frac{G_\kappa(\mathbf{x}-\mathbf{y})}{\partial\boldsymbol{\nu}(\mathbf{y})}v(\mathbf{y})\right\} d\mathbf{y} \tag{43}$$

for $\mathbf{x} \in \partial\Omega$—whose exact value, in view of Green's theorem, is $e^{i\kappa x_1}$. The corresponding results over successive discretizations are presented in Table 1, clearly demonstrating high-order accuracy.

The proposed integration scheme additionally remains accurate for large frequencies. To illustrate this, we have computed the integral (43) for various wavenumbers; the corresponding results are presented in Table 2 for experiments with a fixed number of points per wavelength. Table 2 shows that, as claimed, the proposed scheme does not deteriorate as the wavenumber is increased while keeping a constant number of points per wavelength.

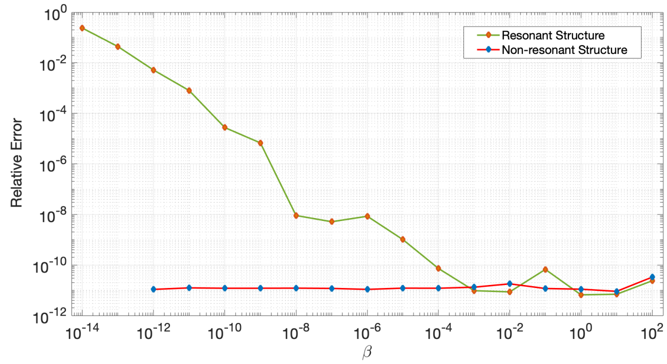

Figure 3: Accuracy of the proposed hybrid solver for the problem (15)–(17), as a function of the impedance parameter β defined in Section 2.1, with $\kappa = \sqrt{2}\pi$ and $n = 1$ in the domain $\Omega = (-a, a) \times (-a, a)$ with $a = 1$ ("resonant structure", with interior eigenfunction $\sin(\pi x_1)\sin(\pi x_2)$) and $a = 1.1$ ("non-resonant structure"). In all cases, resonant and non-resonant, the GMRES algorithm achieved the 10^{-12} tolerance imposed. As illustrated in the figure, in non-resonant cases the error is essentially independent of β as $\beta \to 0$—since in such cases the ItI map T_{int} is well defined for all β, up to and including $\beta = 0$. In the resonant case, in contrast, use of a nonzero value of β is necessary to ensure the map T_{int} is well defined and the algorithms accuracy does not deteriorate.

κ	P	q	ϵ_∞^N	Order
5π	8	6	2.7×10^{-0}	-
5π	8	12	5.1×10^{-1}	2.4
5π	8	24	2.7×10^{-3}	7.5
5π	8	48	3.9×10^{-8}	16.1
5π	8	96	4.9×10^{-11}	9.6
5π	8	192	1.2×10^{-13}	8.7

Table 1: Convergence study for the singular integration method introduced in Section 3.2. Numerical errors were obtained by comparison against closed-form values of the integral (43).

κ	P	q	PPW	ϵ_∞^N
10π	12	30	6	2.3×10^{-5}
20π	24	30	6	2.6×10^{-5}
40π	48	30	6	2.7×10^{-5}
80π	96	30	6	2.9×10^{-5}
160π	192	30	6	3.0×10^{-5}
320π	384	30	6	3.1×10^{-5}

Table 2: Illustration of the proposed high-order integration scheme for large wavenumbers with a fixed number of points per wavelength.

Example 4.2. *(Sparsity and Efficiency of the Hybrid Approach)*

As discussed in the introduction, use of the hybrid direct/iterative strategy, in which the boundary integral equation is treated iteratively, provides a significant advantage over the corresponding direct non-hybrid approach, in which a matrix is constructed for the (complete) coupled volume and boundary discretization. This advantage arises mainly from sparsity: the matrix $\mathbf{A}$ (equation (35)) associated with the hybrid approach is significantly sparser than the corresponding non-hybrid matrix, as the coupling induced by the boundary integral operator introduces large numbers of nonzero matrix entries. As a result (and as demonstrated below in this section) the hybrid approach lends itself much more effectively to treatment via multifrontal linear-algebra solvers. To visualize the source of the sparsity enjoyed by the matrix $\mathbf{A}$ we note that, in the non-hybrid approach, each boundary entry gives rise to an equation that links all $4P(q+1)$ boundary unknowns and, additionally, in view of equations (15) through (17), $(q-1)$ interior unknowns per boundary unknown (as needed to compute the normal derivative at each boundary point)—so that, in total, each equation resulting from a boundary point contains $4Pq(q+1)$ nonzero entries. This is in contrast to the equations arising from interior unknowns, each one of which contains merely $2(q+1)$ non-zero entries. The benefit provided by the hybrid method is that it decomposes the problem into two parts: a first one that uses a direct solver for the sparse matrix associated with interior unknowns, and a second one which treats the boundary unknowns by means of an iterative procedure.

Table 3, which displays the total number "NNZ" of non-zero matrix entries contained in the matrices treated by means of a direct solver for the hybrid and non-hybrid methods, demonstrates the sparsity patterns achieved in practice by the proposed hybrid approach. As illustrated in Figure 4, further, such sparsity patterns translate into fast pre-computation and solution times—which, in fact, grow nearly linearly with the discretization size.

$P \times P$	$q \times q$	# Bdry. Unknowns	Bdry. unknowns NNZ	
		$4P(q+1)$	Non-hybrid	Hybrid
16×16	10×10	704	$5,451,776$	7744
32×32	10×10	1408	$21,807,104$	15488
64×64	10×10	2816	$87,228,416$	30976
128×128	10×10	5632	$348,913,664$	61952
256×256	10×10	11264	$1,395,654,656$	123904

Table 3: Number NNZ of non-zero matrix entries associated with each boundary unknown for the non-hybrid and hybrid algorithms, respectively, for various discretization sizes. The greatly enhanced sparsity pattern associated with the hybrid method enables efficient use of multi-frontal linear-algebra solvers.

$P \times P$	N	Pre-comp. (sec.)		Per-it. time (sec.)			Memory required (MB)		
		A/BIE	**LU**-D	**LU**-inv/It	BIE/It	Tot.	**A** storage	**LU**-D	Ratio
16×16	30,976	-/0.6	0.5	.04	.01	.05	-	-	-
32×32	123,904	.03/3	2	.15	.06	0.21	-	-	-
64×64	495,616	.12/12	10	.6	.3	0.9	502	1,497	2.98
128×128	$1,982,464$	.46/46	46	3	1	4	12,99	8,104	6.24
175×175	$3,705,625$	.88/87	93	5	2	7	2,222	15,617	7.03
200×200	$4,840,000$	1.11/113	131	7	3	10	2,830	2,1250	7.50
256×256	$7,929,856$	2.00/182	247	11	4	15	4,484	34,624	7.72
350×350	$14,822,500$	4/340	541	22	7	29	8,175	66,410	8.12

Table 4: Computing times and memory required by the various portions of the hybrid algorithm on the computational domain $\Omega = [-.5, .5]^2$ with $\kappa = 800$. Each one of the $P \times P$ Chebyshev patches used was discretized by means of a $q \times q$ Chebyshev-point discretization with $q = 10$. The titles used are defined in the text. The BIE precomputation cost can be essentially eliminated if an accelerated Green function method [7, 22] is utilized. As indicated in the text, after the precomputation stages, small additional memory costs suffice to perform even very large numbers of GMRES iterations, if needed.

Table 4 and its caption, in turn, report computing times and memory required to perform each one of the various operations associated with the hybrid method. Thus, in particular, for a problem involving nearly 15 million unknowns, the single-core precomputation and per-iteration computing times amount to 344+541 $\approx$ 900 sec. and $22 + 7 \approx 30$ sec. respectively, with a corresponding memory cost of $(8,175 + 66,410 + 7,500)$ MB $\approx$ 82 GB. (The BIE precomputation time and memory cost, the latter one of which is not listed in Table 4, but which amounts to e.g. 7,500 MB for the $\approx$ 15 million unknown problem, can be essentially eliminated if an accelerated Green function method [7, 22] is utilized.) An additional (small) memory cost is associated with each GMRES iteration: in the $\approx$ 15 million unknown test case, for example, after an initial integral equation setup memory cost of $1,732$ MB, every 100 iterations require a mere 25 MB of additional memory. Thus, in view of Table 10 below, using this discretization a solution with an error of the order of 10^{-3} for a domain spanning *350 wavelengths* in diameter containing a *discontinuous* refractive index can be obtained, on the basis of 12 iterations, in a *single-core* CPU time of $\approx 900 + 12 \cdot 30 \ = 1,260$ secs. = 21 mins.

(The titles used in Table 4 and 5 are defined as follows. The "Pre-comp" columns list the costs of the various precomputation stages, namely "**A**": computing time required to produce the interior matrix; "BIE": computing time required to evaluate all the necessary values of the Green function; and "LU-D": computing time required to obtain the LU decomposition of **A** by means of the multifrontal linear-algebra software "Intel MKL PARDISO". The "Per-it" columns lists costs necessary to perform each iteration, namely "LU-inv/It": computing time required at each iteration of the iterative hybrid algorithm to solve equation (17) on the basis of the precomputed LU decomposition; and "BIE/It": computing time required to apply the discrete version of the integral operator T_{int} in equation (16), respectively; the column "Tot." lists total computing time per iteration. The "Memory required" column in Table 4 lists memory costs, including"**A** storage": Memory required for storage of the matrix **A** and associated data required by the software"Intel MKL PARDISO"; and "**LU**-D": Total memory required by the solver Pardiso for the precomputation of the **LU** decomposition of the matrix **A**; the column "Ratio" lists the ratio of the memory requirements in the two previous columns.)

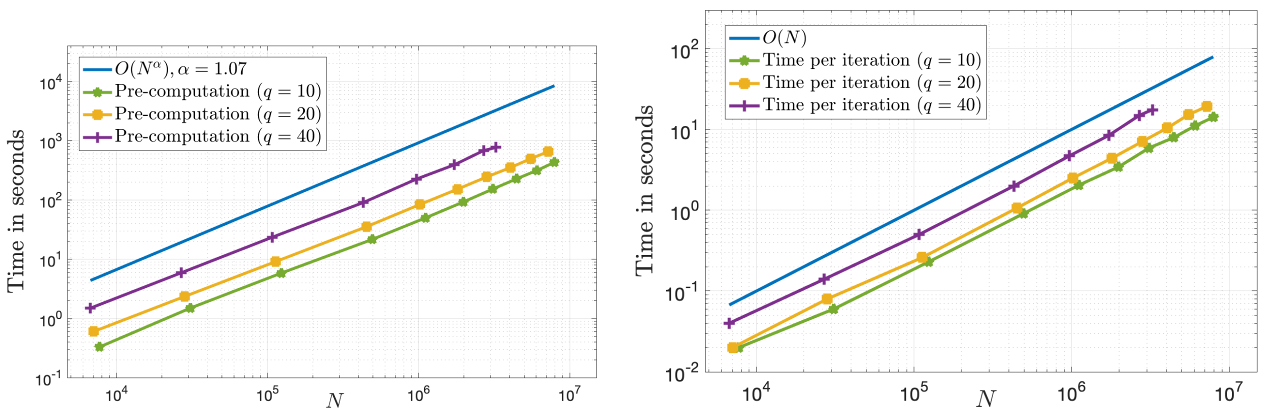

(a) Pre-computation costs for $q = 10$, $q = 20$, and $q = 40$.

(b) Time per GMRES iteration for $q = 10, 20$, and $q = 40$.

Figure 4: **Left:** Pre-computation time in seconds (required to build the matrix **A**, to obtain its **LU** decomposition and to produce the BIE Green-function precomputation), for $q = 10$, $q = 20$ and $q = 40$, as a function of N, with $\kappa = 800$ and for all discretizations allowable within the available memory, vs. a curve $O(N^\alpha)$ with $\alpha = 1.07$. **Right:** Per-iteration time (in sec.) required by the hybrid algorithm with $q = 10$, $q = 20$ and $q = 40$, as a function of N, vs. a line $O(N)$). As illustrated in Table 5, a variety of numerical experiments have shown that these computing times are essentially constant asymptotically as κ grows.

$P \times P$	$q \times q$	κ	Pre-comp. (in sec.)			Per it. time
			A/BIE	**LU**-D	Tot.	(sec.)
256×256	10×10	10	2.00/134.74	246.52	383.26	15.20
256×256	10×10	100	2.00/174.89	246.74	423.63	15.28
256×256	10×10	200	2.00/181.32	246.56	429.88	15.42
256×256	10×10	400	2.00/180.03	246.18	428.21	15.49
256×256	10×10	800	2.00/181.89	247.74	431.63	15.49

Table 5: Pre-computation time for $q = 10$ for various values of κ with fixed $N = P^2 \times (q+1)^2 = 7,929,856$.

Example 4.3. *(Scattering by a Smooth Gaussian Bump)*

This example demonstrates the high-order convergence enjoyed by the proposed algorithm when applied to smooth contrast functions $m(\mathbf{x})$. In detail, we consider the total field u that arises under plane wave excitation incident from the positive x_1 axis, for the contrast function given by the smooth Gaussian bump $m(\mathbf{x}) = -1.5e^{-60|\mathbf{x}|^2}$. Numerical results for the domain $\Omega = (-0.5, 0.5)^2$ and $\kappa = 20\pi$, at various discretization levels, are displayed in Table 6—clearly demonstrating the high-order convergence of the proposed algorithm for smooth scattering media. The extremely low dispersion provided by the proposed algorithm is demonstrated in Table 7 (see also Tables 9 and 10)—which shows that, for the same domain Ω, the accuracy is maintained while keeping the number of points per wavelength fixed—even for large frequencies.

κ	$P \times P$	$q \times q$	ϵ_∞^N	Order	# Iter
20π	2×2	10×10	1.39×10^{-0}	-	28
20π	4×4	10×10	5.81×10^{-1}	1.22	28
20π	8×8	10×10	9.50×10^{-3}	5.93	28
20π	16×16	10×10	2.11×10^{-5}	8.81	28
20π	32×32	10×10	1.08×10^{-7}	7.61	28
20π	64×64	10×10	1.50×10^{-10}	9.49	28

Table 6: Convergence study for the smooth Gaussian bump test case. For these experiments the GMRES residual tolerance and the coupling parameter β were set to 10^{-12} and 10^{-5}, respectively.

Remark 4. It is important to note the relatively mild (roughly linear) increases in iteration numbers demonstrated in Tables 7 and 9 as the incident frequency grows (cf. references [16, 37]). The observed linear growth is purely associated with the spectral character of the boundary integral operators used, and it results as the algorithm bypasses, by means of its interior direct solver component, the iterative resolution of all interior multiple-scattering events that would otherwise require significantly larger iteration numbers.

κ	$P \times P$	$q \times q$	N/Γ_N	ϵ_∞^N	# Iter.	Time (sec.)	
						pre-comp	per. It.
50	10×10	10×10	14641/484	2.75×10^{-5}	14	.63	0.03
100	22×22	10×10	58564/968	6.66×10^{-5}	35	2.35	0.1
200	44×44	10×10	234256/1936	1.09×10^{-4}	76	9.68	0.41
400	88×88	10×10	937024/3872	2.12×10^{-4}	162	43	1.62
800	176×176	10×10	3748096/7744	3.54×10^{-4}	291	194	6.87

Table 7: Numerical solution for a problem of scattering by the smooth Gaussian bump example for a range of frequencies, including high-frequency cases. Approximately 9.4 points per shortest wavelength (which occurs at $\mathbf{x} = 0$) were used for the $\kappa = 50$ through $\kappa = 800$ examples (for $\kappa = 800$ the computational domain is two-hundred five shortest wavelengths in size). For these experiments both the GMRES residual and the coupling parameter β were set to 10^{-5}.

Example 4.4. *(Fourier Smoothing and scattering by a discontinuous refractive index distribution)*

This example demonstrates the character of the proposed Filtered Fourier Smoothing strategy (Section 3.1.1) for penetrable inhomogeneous media with discontinuous refractivity—via an application to the canonical problem of scattering by a circular scatterer. For this experiment we considered a circular scatterer $\mathcal{D}$ of diameter $d = 1$, and with discontinuous refractive index given by $n^2(\mathbf{x}) = 2$ for $\mathbf{x} \in \mathcal{D}$ and

$n^2(\mathbf{x}) = 1$ for $\mathbf{x} \notin \mathcal{D}$. The computational domain $\Omega = (-.51, .51) \times (-.51, .51)$ was used. An incident wave of the form $u^i(\mathbf{x}) = J_0(\kappa|\mathbf{x}|)$ was assumed, where J_0 is the Bessel function of the first kind of order zero. With this incident wave a closed form expression for the solution of the problem (2)-(3) is known [5]. Table 8 presents errors obtained in the numerical solution with and without Fourier smoothing, and including regular Fourier smoothing (FS) and filtered Fourier smoothing (FFS), for various discretization levels, clearly demonstrating the quadratic convergence of the FFS-based approach, the slower and rather erratic convergence in absence of Fourier smoothing, and the improvements resulting from the use of filtering. Table 9, in turn, concerns the character of the FFS method under high-frequency illumination, displaying fixed accuracies (of the order of three digits in this case), for the $n^2(\mathbf{x}) = 2$ scatterer $\mathcal{D}$ just considered and for problems up to $276 \cdot \lambda_{\text{int}}$ in diameter, where $\lambda_{\text{int}} = \frac{2\pi}{n\kappa} d$ denotes the wavelength in the interior of $\mathcal{D}$. We note that a fixed accuracy is maintained using 11 points per wavelength, demonstrating, additionally, the dispersionless character of the algorithm even under a discontinuous index of refraction, for which the accuracy of the algorithm is reduced to second order. Table 10, finally, presents numerical results for highly refractive scatterers. In contrast with the behavior observed in the case of high-frequency illumination, in the present case, in which high-frequencies result from corresponding large refractive indexes, the iteration numbers required to maintain accuracy remain fixed as the refractivity values are increased—on account of the direct matrix solution used for the interior problem, and in spite of the resulting high-frequency interior scattering phenomenology.

κd	$P \times P$	$q \times q$	Without FS		With FS (without filter)		With FFS (incl. filter)	
			ϵ_∞^N	Order	ϵ_∞^N	Order	ϵ_∞^N	Order
10π	3×3	10×10	8.80×10^{-2}	-	1.66×10^{-1}	-	2.14×10^{-1}	-
10π	6×6	10×10	1.27×10^{-2}	2.79	1.60×10^{-2}	3.37	4.60×10^{-2}	2.22
10π	12×12	10×10	1.22×10^{-2}	0.06	5.49×10^{-4}	4.86	1.43×10^{-3}	5.00
10π	24×24	10×10	3.30×10^{-3}	1.88	2.50×10^{-4}	1.13	3.18×10^{-4}	2.17
10π	48×48	10×10	2.66×10^{-3}	0.31	8.87×10^{-5}	1.49	5.09×10^{-5}	2.64
10π	96×96	10×10	5.72×10^{-4}	2.22	4.82×10^{-5}	0.87	1.23×10^{-5}	2.04
10π	192×192	10×10	1.69×10^{-4}	1.76	1.08×10^{-5}	2.15	3.00×10^{-6}	2.03

Table 8: Demonstration of the quadratic convergence of the FS and FFS-based hybrid solvers for a problem of scattering by a circular inclusion of diameter $d = 1$ with $u^i(\mathbf{x}) = J_0(\kappa|\mathbf{x}|)$ and with discontinuous refractive index is given by $n^2(\mathbf{x}) = 2$ for $\mathbf{x} \in \mathcal{D}$ and $n^2(\mathbf{x}) = 1$ for $\mathbf{x} \notin \mathcal{D}$. For these experiments the GMRES residual tolerance and the coupling parameter β were set to 10^{-6} and 10^{-5}, respectively. The beneficial effects of Fourier smoothing and filtering, which lead to higher accuracies and a more predictable convergence behavior, can be clearly appreciated.

κ	$P \times P$	$q \times q$	N/Γ_N	ϵ_∞^N	# Iter.	Time (sec.)	
						pre-comp	per. It.
50	10×10	10×10	14641/484	2.02×10^{-3}	15	0.61	0.03
100	22×22	10×10	58564/968	2.55×10^{-3}	37	3	0.09
200	44×44	10×10	245025/1980	2.78×10^{-3}	69	11	0.42
400	88×88	10×10	980100/3960	3.80×10^{-3}	161	46	1.72
800	176×176	10×10	3920400/7920	2.81×10^{-3}	281	200	7.25
1200	264×264	10×10	8433216/11616	3.90×10^{-3}	400	470	16.3

Table 9: High-frequency scattering problem. Numerical solution, using FFS, for a problem of scattering by a circular inclusion of diameter $d = 1$, with $u^i(\mathbf{x}) = J_0(\kappa|\mathbf{x}|)$, and with $n^2(\mathbf{x}) = 2$ for $\mathbf{x} \in \mathcal{D}$ and $n^2(\mathbf{x}) = 1$ otherwise. For these experiments the GMRES residual tolerance and coupling parameter β were set to 10^{-5}.

$P \times P$	$q \times q$	N/Γ_N	ref-index $n(\mathbf{x})$	# λ_{int}	ϵ_∞^N	# Iter.	Time (sec.)	
							pre-comp	per. It.
125×125	10×10	1890625/5500	50	125	3.43×10^{-3}	10	84	3.4
150×150	10×10	2722500/6600	60	150	4.52×10^{-3}	10	125	5.0
200×200	10×10	4840000/8800	80	200	4.15×10^{-3}	12	232	9.3
250×250	10×10	7562500/11000	100	250	3.73×10^{-3}	12	403	14.0
350×350	10×10	14822500/15400	140	350	3.34×10^{-3}	12	848	29.0

Table 10: Large contrast scattering problem. Numerical solution, using FFS, for a problem of scattering by a circular inclusion of diameter $d = 2$ refractive index $n(\mathbf{x})$ (resp. refractive index 1) in the interior (resp. the exterior) of the inclusion, with $u^i(\mathbf{x}) = J_0(\kappa|\mathbf{x}|)$, where $\kappa = 5\pi$. Three digit accuracy is maintained using 11 points per wavelength. For these experiments the GMRES residual tolerance and coupling parameter β were set to 10^{-5}.

Figure 5 provides a graphical depiction of the scattering pattern obtained for a circular scatterer $\mathcal{D}$ of diameter $d = 2$, under incident illumination given by $u^i(\mathbf{x}) = \exp(i\kappa x_1)$, with $\kappa = 100$ and with $n^2(\mathbf{x}) = 3$ for $\mathbf{x} \in \mathcal{D}$ and $n^2(\mathbf{x}) = 1$ for $\mathbf{x}$ outside $\mathcal{D}$—for which we have $d = 55\lambda_{\text{int}}$. Using 12 points per wavelength the method achieves three digits of accuracy for this problem in the near field in a two and half minutes single-core computation, including both, precomputation and all necessary iterations.

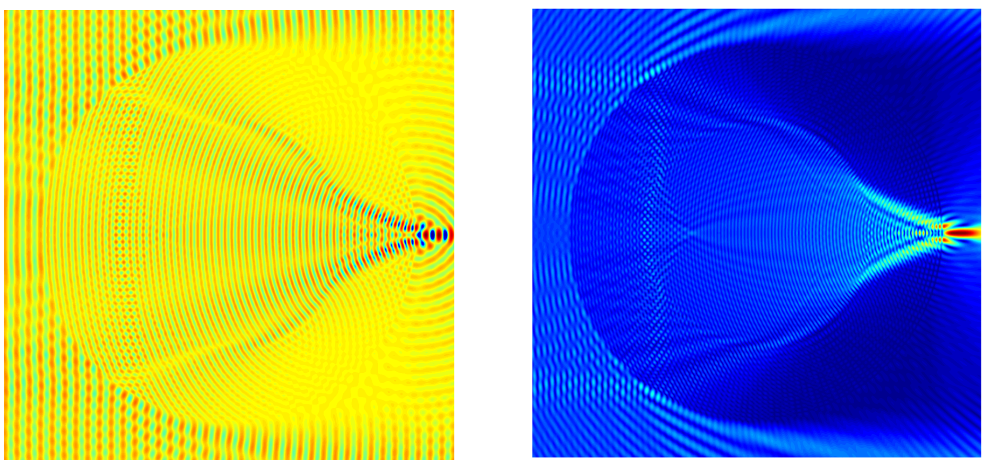

(a) Real part of the total Field u.　　(b) Absolute value of the total field u.

Figure 5: Scattering of the plane wave $\exp(i\kappa x_1)$ with $\kappa = 100$ by a penetrable circular inclusion $\mathcal{D}$ of diameter $d = 55\lambda_{\text{int}}$ with $n^2(\mathbf{x}) = 3$ for $\mathbf{x} \in \mathcal{D}$ and $n^2(\mathbf{x}) = 1$ otherwise. Using 12 points per wavelength and relying on the FFS method, the algorithm produced this three-digit accurate solution in a two and half minutes single-core computation.

Example 4.5. *(Scattering by variable discontinuous refractivity)*

Our next example demonstrates the properties of the solver, including FFS, when applied to a scatterer containing continuously variable material properties as well as discontinuities across a material interface. We thus consider the problem of evaluation of the total field u that results for the refractive-index distribution

$$n^2(\mathbf{x}) = \begin{cases} 3 + 2e^{-4|\mathbf{x}|^2} & \text{if } \mathbf{x} \in \mathcal{D}, \\ 1 & \text{otherwise,} \end{cases} \tag{44}$$

where $\mathcal{D}$ is circular inclusion of unit radius, under the plane wave incidence $u^i(\mathbf{x}) = \exp(i\kappa x_1)$. Since analytical solutions are not available in this case, we use numerical solution obtained on a finer grids for reference. The numerical results reported in Table 11 display errors that in fact decrease faster than the quadratic rate expected from use of the FFS approach.

κd	$P \times P$	$q \times q$	Without FFS		With FFS	
			ϵ_∞^N	Order	ϵ_∞^N	Order
10π	3×3	10×10	1.15×10^{0}	-	1.06×10^{-0}	-
10π	6×6	10×10	1.45×10^{-1}	2.98	1.18×10^{-1}	3.16
10π	12×12	10×10	5.48×10^{-2}	1.40	1.28×10^{-2}	3.20
10π	24×24	10×10	1.48×10^{-2}	1.88	1.53×10^{-3}	3.06
10π	48×48	10×10	7.73×10^{-3}	0.93	3.34×10^{-4}	2.19
10π	96×96	10×10	4.24×10^{-3}	.86	7.62×10^{-5}	2.13
10π	192×192	10×10	8.32×10^{-4}	2.34	1.42×10^{-5}	2.42

Table 11: Convergence Study: Convergence of the proposed algorithm for scattering for a discontinuous-refractivity problem as in (44). An incident plane wave incoming from the positive x-axis was used in this case. As noted in the text, the errors decrease somewhat faster than the quadratic rate expected from use of the FFS approach.

Figure 6 displays near fields obtained for the discontinuous refractive-index (44) under wavenumbers $\kappa = 50$ and $\kappa = 150$, for which the diameters of inhomogeneity are $36\lambda_{\text{min}}$ and $108\lambda_{\text{min}}$, respectively, where λ_{min} denotes the smallest interior wavelength. In both the cases, the algorithm achieved three-digit accuracy in the near field by using 11 points per λ_{min} in single-core computations requiring one and thirteen minutes, respectively.

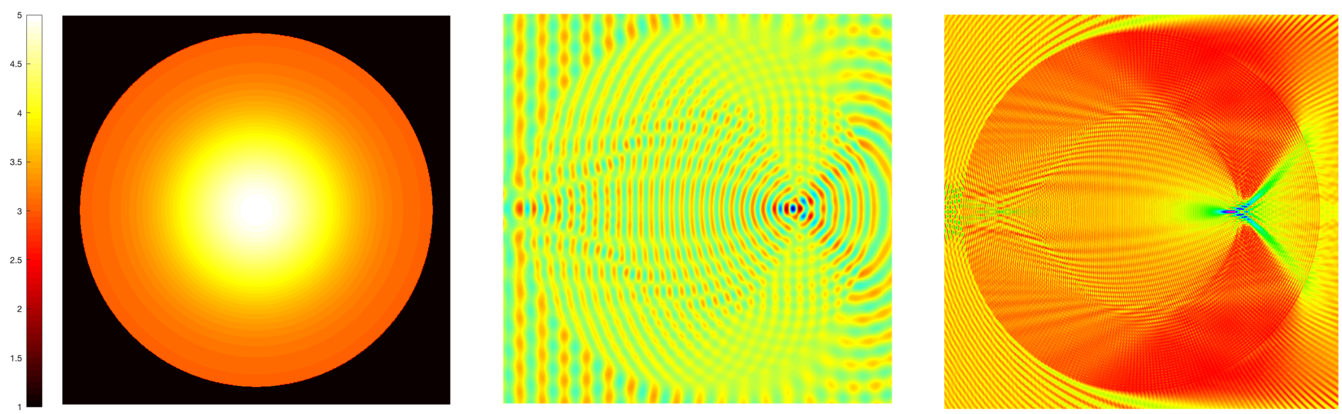

(a) Gaussian Refractivity. (b) $\kappa = 50$, Real part of u (c) $\kappa = 150$, $|u|$.

Figure 6: Scattering of a plane wave $\exp(i\kappa x_1)$ by a the Gaussian refractivity profile (44) for $\kappa = 50$ and $\kappa = 150$. In both cases, using the FFS method and 11 points per wavelength the algorithm produced three-digit accuracy in a one and thirteen-minute computation respectively.

Example 4.6. *(Scattering by geometries containing corners and cusps)*

None of the algorithmic components, nor the resulting accuracies in the proposed method, are constrained in any way by the geometry of the scatterer. Without any additional effort, the approach can easily deal with arbitrarily complicated geometries. To demonstrate this, we consider two additional geometries, containing corner- and cusp-singularities, respectively. Once again the accuracy of any one solution is evaluated by comparison with results obtained on finer grids. In both cases we compute the near field solution u under the plane wave incidence $u^i(\mathbf{x}) = \exp(i\kappa x_1)$.

κd	$P \times P$	$q \times q$	ϵ_∞^N	Order
12π	3×3	10×10	1.33×10^0	-
12π	6×6	10×10	6.95×10^{-2}	4.26
12π	12×12	10×10	9.75×10^{-3}	2.83
12π	24×24	10×10	2.12×10^{-3}	2.20
12π	48×48	10×10	4.66×10^{-4}	2.18
12π	96×96	10×10	9.22×10^{-5}	2.34
12π	192×192	10×10	1.92×10^{-5}	2.26

Table 12: Convergence Study: Illustration of quadratic convergence of the proposed algorithm for a geometry containing a corner singularity.

κd	$P \times P$	$q \times q$	ϵ_∞^N	Order
12π	3×3	10×10	1.42×10^0	-
12π	6×6	10×10	9.86×10^{-2}	3.84
12π	12×12	10×10	1.21×10^{-2}	3.02
12π	24×24	10×10	2.26×10^{-3}	2.42
12π	48×48	10×10	6.23×10^{-4}	1.85
12π	96×96	10×10	1.46×10^{-4}	2.09
12π	192×192	10×10	2.15×10^{-5}	2.76

Table 13: Convergence Study: Illustration of quadratic convergence of the proposed algorithm for a geometry containing a cusp singularity.

Table 12 presents numerical results for the scatterer $\mathcal{D}$ depicted in Figure 7(a), with $n^2(\mathbf{x}) = 2$ for $\mathbf{x} \in \mathcal{D}$ and one otherwise. The computed near field for $\kappa = 200$, which was determined to be accurate up to three digits, is displayed in Figure 7(b). Table 13, in turn, presents numerical results for the scatterer $\mathcal{D}$ depicted in Figure 8(a), which equals the region contained between the four unit discs centered at $(1, 1)$, $(1, -1)$, $(-1, 1)$ and $(-1, -1)$, with $\kappa d = 12\pi$, and with $n^2(\mathbf{x}) = 2$ for $\mathbf{x} \in \mathcal{D}$ and $n^2(\mathbf{x}) = 1$ otherwise. Figure 8(b), finally, displays the near field for this geometry, but with $\kappa = 20\pi$ and $n^2(\mathbf{x}) = 16$ for $\mathbf{x} \in \mathcal{D}$ and one otherwise—thus yielding a scatterer $80\lambda_{\text{int}}$ in size. A two-digit solution was obtained using merely nine points per wavelength and computing time of seven minutes.

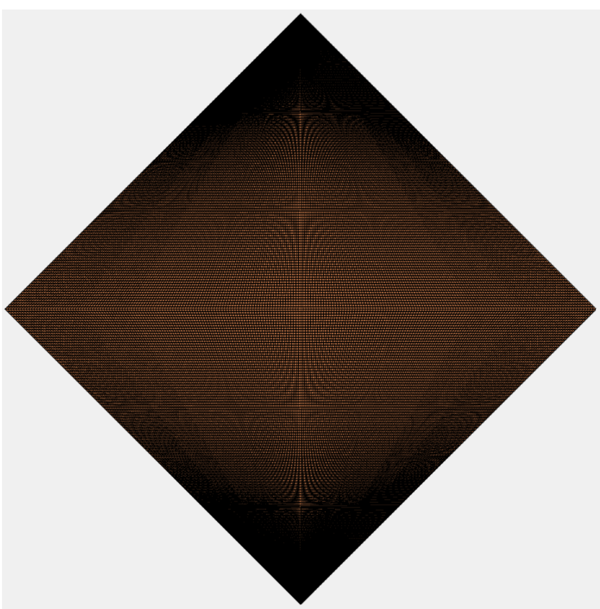

(a) Square scatterer

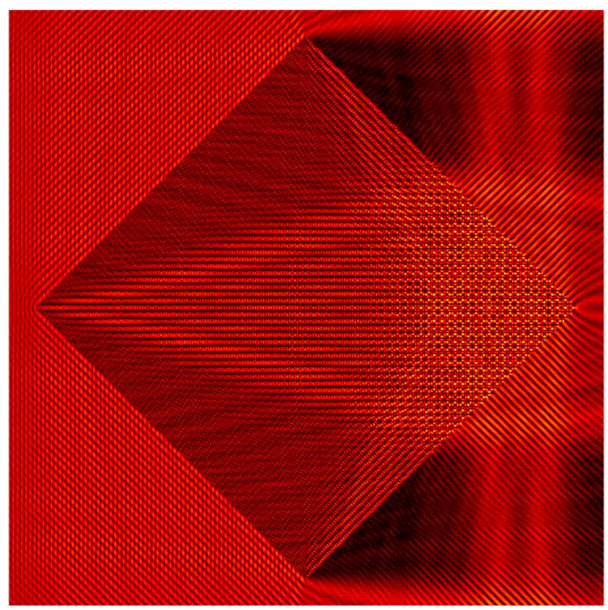

(b) Absolute value of the total field u.

Figure 7: Scattering by a geometry containing corner singularities, with $n^2(\mathbf{x}) = 3$ for $\mathbf{x} \in \mathcal{D}$ and one otherwise. For this experiment the incident field $u^i(\mathbf{x}) = \exp(i\kappa x_1)$ with $\kappa = 200$ was used. Errors of the order of 10^{-3} were obtained in the near field solution and the total computing time is fourteen minutes.

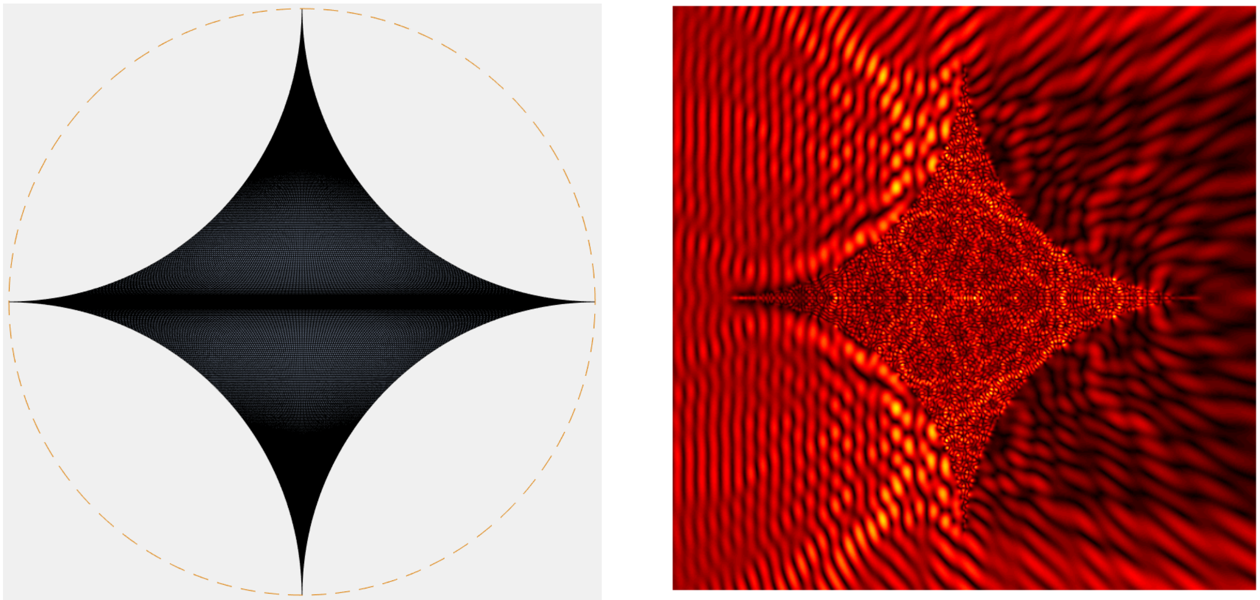

(a) Star-shaped geometry with cusp

(b) Absolute value of total field u

Figure 8: Scattering of the incident field $u^i(\mathbf{x}) = \exp(i\kappa x_1)$, with $\kappa = 20\pi$, by a geometry containing cusp singularities, with $n^2(\mathbf{x}) = 16$ for $\mathbf{x} \in \mathcal{D}$ and $n^2(\mathbf{x}) = 1$ otherwise. Errors of the order of 10^{-2} were obtained in the near field solution on the basis of nine points per interior wavelength.

5 Conclusions

This paper introduced a new methodology for solutions of two-dimensional problems of scattering by penetrable inhomogeneous media with possibly discontinuous refractivity. The solver achieves high-order convergence for smooth refractivities at nearly-linear computing cost, and, to the best the of our knowledge, it is the first hybrid direct/iterative solver which yields second order convergence for discontinuous refractivities, and for low- or high-frequencies alike. The method additionally enjoys very low dispersion for either smooth or discontinuous refractive indexes, and it can natively and easily handle scatterers with complicated geometric singularities, including e.g. as corners and cusps. Extensions of the proposed approach to electromagnetic and elastic wave scattering problems, as well as three-dimensional configurations are envisioned.

Acknowledgments

This work was supported by NSF, DARPA and AFOSR through contracts DMS-2109831 and HR00111720035, and FA9550-21-1-0373, and by the NSSEFF Vannevar Bush Fellowship under contract number N00014-16-1-2808.

A Appendix: Fast and accurate computation of Fourier coefficients of discontinuous functions

In order to enable fast and accurate evaluation of the Fourier coefficients of a given, possibly discontinuous, function f in the interval $[0, 2\pi]$, as needed in Section 3.1.1 (see Remark 2), we rely on the Fourier continuation (FC) approach [3, 18]. For our description we assume that the function f has only one discontinuity, say, at $x = a \in (0, 2\pi)$, but an arbitrary number of discontinuities may be treated in similar fashion.

Let now f_j^c $(j = 1, 2)$ denote d_j-periodic Fourier continuation functions of the restrictions of the function f to the intervals $[0, a]$ and $[a, 2\pi]$, respectively. We thus have

$$f_j^c(x) = \sum_{k=-F}^{F} c_k^j e^{\frac{2\pi ikx}{d_j}}, \tag{45}$$

where, following e.g. [3], the Fourier coefficients c_k^j are obtained in $O(F \log F)$ operations by means of the FC procedure and associated FFTs, and where the resulting functions f_j^c with $j = 1, 2$ approximate the restrictions of the function f to the intervals $[0, a]$ and $[a, 2\pi]$, respectively, with high-order accuracy. Let

$$f_\ell = \frac{1}{2\pi} \int_0^{2\pi} f(t) e^{-i\ell t} dt = \frac{1}{2\pi} \int_0^{a} f(t) e^{-i\ell t} dt + \frac{1}{2\pi} \int_a^{2\pi} f(t) e^{-i\ell t} dt \tag{46}$$

denote the desired Fourier coefficient of f in the interval $[0, 2\pi]$. The two integrals on the right-hand side of (46) can be computed with high accuracy by substituting f by f_j^c and exchanging integration and summation. In the case of the first integral, for example, we have

$$\int_0^a f(t) e^{-i\ell t} dt \approx \int_0^a f_1^c(t) e^{-i\ell t} dt = \sum_{k=-F}^{F} c_k^1 \int_0^a e^{\frac{2\pi k - \ell d_1}{d_1} it} dt = \sum_{k=-F}^{F} c_k^1 b_{2\pi k - \ell d_1}, \tag{47}$$

where

$$b_{2\pi k - \ell d_1} = \begin{cases} \frac{d_1}{i(2\pi k - \ell d_1)} \left(e^{\frac{2\pi k - \ell d_1}{d_1} ia} - 1 \right) & \text{if } (2\pi k - \ell d_1) \neq 0, \\ a & \text{otherwise.} \end{cases}$$

The summation in (47) is a discrete scaled convolution and can be obtained for all ℓ in $O(F \log F)$ operations by using FFT [41]. Thus highly-accurate values of the Fourier coefficients f_ℓ of the discontinuous function f, for all ℓ, $-F \leq \ell \leq F$, can be produced in $O(F \log F)$ operations.

References

[1] C. Alappat, A. Basermann, A. R. Bishop, H. Fehske, G. Hager, O. Schenk, J. Thies, and G. Wellein. A recursive algebraic coloring technique for hardware-efficient symmetric sparse matrix-vector multiplication. *ACM Transactions on Parallel Computing (TOPC)*, 7(3):1–37, 2020.

[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 dimensions. *SIAM Journal on Scientific Computing*, 38(3):A1770–A1787, 2016.

[3] F. Amlani and O. P. Bruno. An FC-based spectral solver for elastodynamic problems in general three-dimensional domains. *Journal of Computational Physics*, 307:333–354, 2016.

[4] A. Anand, A. Pandey, B. R. Kumar, and J. Paul. An efficient high-order Nyström scheme for acoustic scattering by inhomogeneous penetrable media with discontinuous material interface. *Journal of Computational Physics*, 311:258–274, 2016.

[5] F. Andersson and A. Holst. A fast, bandlimited solver for scattering problems in inhomogeneous media. *Journal of Fourier Analysis and Applications*, 11(4):471–487, 2005.

[6] I. M. Babuska 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(6):2392–2423, 1997.

[7] C. Bauinger and O. P. Bruno. "interpolated factored green function" method for accelerated solution of scattering problems. *Journal of Computational Physics*, 430:110095, 2021.

[8] A. Bayliss, C. I. Goldstein, and E. Turkel. On accuracy conditions for the numerical computation of waves. *Journal of Computational Physics*, 59(3):396–404, 1985.

[9] J.-D. Benamou and B. Desprès. A domain decomposition method for the helmholtz equation and related optimal control problems. *Journal of Computational Physics*, 136(1):68–82, 1997.

[10] A. Bendali, Y. Boubendir, and M. Fares. A feti-like domain decomposition method for coupling finite elements and boundary elements in large-size problems of acoustic scattering. *Computers & structures*, 85(9):526–535, 2007.

[11] A. Bermúdez, L. Hervella-Nieto, A. Prieto, R. Rodrı, et al. An optimal perfectly matched layer with unbounded absorbing function for time-harmonic acoustic scattering problems. *Journal of computational Physics*, 223(2):469–488, 2007.

[12] M. BollhÃ¶fer, A. Eftekhari, S. Scheidegger, and O. Schenk. Large-scale sparse inverse covariance matrix estimation. *SIAM Journal on Scientific Computing*, 41(1):A380–A401, 2019.

[13] Y. Boubendir, A. Bendali, and M. Fares. Coupling of a non-overlapping domain decomposition method for a nodal finite element method with a boundary element method. *International journal for numerical methods in engineering*, 73(11):1624–1650, 2008.

[14] O. P. Bruno and E. Garza. A Chebyshev-based rectangular-polar integral solver for scattering by general geometries described by non-overlapping patches. *arXiv preprint arXiv:1807.01813*, 2018.

[15] O. P. Bruno and C. A. Geuzaine. An O(1) integration scheme for three-dimensional surface scattering problems. *Journal of Computational and Applied Mathematics*, 204(2):463–476, 2007.

[16] 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(2):670–694, 2004.

[17] O. P. Bruno and E. M. Hyde. Higher-order Fourier approximation in scattering by two-dimensional, inhomogeneous media. *SIAM Journal on Numerical Analysis*, 42(6):2298–2319, 2005.

[18] O. P. Bruno and M. Lyon. High-order unconditionally stable FC-AD solvers for general smooth domains i. basic elements. *Journal of Computational Physics*, 229(6):2009–2033, 2010.

[19] F. Cakoni, D. Colton, and P. Monk. The direct and inverse scattering problems for partially coated obstacles. *Inverse problems*, 17(6):1997, 2001.

[20] B. Caudron, X. Antoine, and C. Geuzaine. Optimized weak coupling of boundary element and finite element methods for acoustic scattering. *Journal of Computational Physics*, 421:109737, 2020.

[21] R. Cimpeanu, A. Martinsson, and M. Heil. A parameter-free perfectly matched layer formulation for the finite-element-based solution of the helmholtz equation. *Journal of Computational Physics*, 296:329–347, 2015.

[22] R. Coifman, V. Rokhlin, and S. Wandzura. The fast multipole method for the wave equation: A pedestrian prescription. *IEEE Antennas and Propagation magazine*, 35(3):7–12, 1993.

[23] D. Colton and R. Kress. *Inverse acoustic and electromagnetic scattering theory*, volume 93. Springer, 2013.

[24] T. A. Davis. Algorithm 832: Umfpack v4. 3—an unsymmetric-pattern multifrontal method. *ACM Transactions on Mathematical Software (TOMS)*, 30(2):196–199, 2004.

[25] R. Duan and V. Rokhlin. High-order quadratures for the solution of scattering problems in two dimensions. *Journal of Computational Physics*, 228(6):2152–2174, 2009.

[26] 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*, pages 325–363. Springer, 2012.

[27] 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*, 55(1):141–170, 2014.

[28] D. Givoli. High-order local non-reflecting boundary conditions: a review. *Wave motion*, 39(4):319–326, 2004.

[29] P. Grisvard. *Elliptic problems in nonsmooth domains.* SIAM, 2011.

[30] T. Hagstrom and T. Warburton. A new auxiliary variable formulation of high-order local radiation boundary conditions: corner compatibility conditions and extensions to first-order systems. *Wave motion*, 39(4):327–338, 2004.

[31] E. M. Hyde and O. P. Bruno. A fast, higher-order solver for scattering by penetrable bodies in three dimensions. *Journal of Computational Physics*, 202(1):236–261, 2005.

[32] S. G. Johnson. Notes on perfectly matched layers (pmls). *arXiv preprint arXiv:2108.05348*, 2021.

[33] A. Kirsch. Remarks on some notions of weak solutions for the Helmholtz equation. *Applicable Analysis*, 47(1-4):7–24, 1992.

[34] A. Kirsch and P. Monk. Convergence analysis of a coupled finite element and spectral method in acoustic scattering. *IMA journal of numerical analysis*, 10(3):425–447, 1990.

[35] A. Kirsch and P. Monk. An analysis of the coupling of finite-element and Nyström methods in acoustic scattering. *IMA Journal of numerical analysis*, 14(4):523–544, 1994.

[36] R. Kress. *Linear Integral Equations*, volume 82. Springer Science & Business Media, 2013.

[37] A. L. Laird and M. Giles. Preconditioned iterative solution of the 2d Helmholtz equation. *Oxford University Computing Laboratory*, 2002.

[38] F. Liu and L. Ying. Sparsify and sweep: An efficient preconditioner for the Lippmann–Schwinger equation. *SIAM Journal on Scientific Computing*, 40(2):B379–B404, 2018.

[39] W. C. H. McLean. *Strongly elliptic systems and boundary integral equations.* Cambridge university press, 2000.

[40] J. M. Melenk and S. Sauter. Wavenumber explicit convergence analysis for galerkin discretizations of the helmholtz equation. *SIAM Journal on Numerical Analysis*, 49(3):1210–1243, 2011.

[41] V. Nascov and P. C. Logofătu. Fast computation algorithm for the Rayleigh-Sommerfeld diffraction formula using a type of scaled convolution. *Applied optics*, 48(22):4310–4319, 2009.

[42] A. Pandey and A. Anand. Improved convergence of fast integral equation solvers for acoustic scattering by inhomogeneous penetrable media with discontinuous material interface. *Journal of Computational Physics*, 376:767–785, 2019.

[43] O. Schenk and K. Gärtner. Solving unsymmetric sparse systems of linear equations with PARDISO. *Future Generation Computer Systems*, 20(3):475–487, 2004.

[44] O. Schenk, K. Gärtner, and W. Fichtner. Efficient sparse LU factorization with left-right looking strategy on shared memory multiprocessors. *BIT Numerical Mathematics*, 40(1):158–176, 2000.

[45] G. Vainikko. Fast solvers of the Lippmann-Schwinger equation. In *Direct and inverse problems of mathematical physics*, pages 423–440. Springer, 2000.

[46] J. Waldvogel. Fast construction of the Fejér and Clenshaw–Curtis quadrature rules. *BIT Numerical Mathematics*, 46(1):195–202, 2006.

[47] L. Ying. Sparsifying preconditioner for the Lippmann–Schwinger equation. *Multiscale Modeling & Simulation*, 13(2):644–660, 2015.

[48] S. Zargaryan and V. Maz'ya. The asymptotic form of the solutions of the integral equations of potential theory in the neighbourhood of the corner points of a contour. *Journal of Applied Mathematics and Mechanics*, 48(1):120–124, 1984.

[49] L. Zepeda-Núñez and H. Zhao. Fast alternating bidirectional preconditioner for the 2d high-frequency Lippmann–Schwinger equation. *SIAM Journal on Scientific Computing*, 38(5):B866–B888, 2016.