

A HIGH-ORDER ACCURATE ACCELERATED DIRECT SOLVER FOR ACOUSTIC SCATTERING FROM SURFACES

JAMES BREMER^{‡,§}, ADRIANNA GILLMAN^{*}, AND PER-GUNNAR MARTINSSON[†]

ABSTRACT. We describe an accelerated direct solver for the integral equations which model acoustic scattering from curved surfaces. Surfaces are specified via a collection of smooth parameterizations given on triangles, a setting which generalizes the typical one of triangulated surfaces, and the integral equations are discretized via a high-order Nyström method. This allows for rapid convergence in cases in which high-order surface information is available. The high-order discretization technique is coupled with a direct solver based on the recursive construction of scattering matrices. The result is a solver which often attains $O(N^{1.5})$ complexity in the number of discretization nodes N and which is resistant to many of the pathologies which stymie iterative solvers in the numerical simulation of scattering. The performance of the algorithm is illustrated with numerical experiments which involve the simulation of scattering from a variety of domains, including one consisting of a collection of 1000 ellipsoids with randomly oriented semiaxes arranged in a grid, and a domain whose boundary has 12 curved edges and 8 corner points.

1. INTRODUCTION

The manuscript describes techniques based on boundary integral equation formulations for numerically solving certain linear elliptic boundary value problems associated with acoustic scattering in three dimensions. There are three principal components to techniques of this type:

1. *Reformulation.* The boundary value problem is first reformulated as a system of integral equations, preferably involving operators which are well-conditioned when considered on spaces of integrable functions.
2. *Discretization.* The resulting integral operators are then discretized. High-order discretization methods are to be preferred to low-order schemes and the geometry of the domain under consideration is the principal consideration in producing high-order convergent approaches.
3. *Accelerated Solver.* The large dense system of linear algebraic equations produced through discretization must be solved. Accelerated solvers which exploit the mathematical properties of the underlying physical problems are required to solve most problems of interest.

For boundary value problems given on two-dimensional domains — which result in integral operators on planar curves — satisfactory approaches to each of these three tasks have been described in the literature. Uniquely solvable integral equation formulations of boundary value problems for Laplace's equation and the Helmholtz equation are available [4, 30, 1, 19, 14]; many mechanisms for the high-order discretization of singular integral operators on planar curves are known [14, 25] and even operators given on planar curves with corner singularities can be discretized efficiently and to near machine precision accuracy [6, 5, 10, 23, 22]; finally, fast direct solvers which often have run times which are asymptotically optimal in the number of discretization nodes N have been constructed [27, 15, 33, 3, 16, 12, 20, 24].

For three-dimensional domains, the state of affairs is much less satisfactory. In this case, the integral operators of interest are given on curved surfaces which greatly complicates problems (2) and (3) above.

^{*}Department of Mathematics, Dartmouth College

[†]Department of Applied Mathematics, University of Colorado, Boulder

[‡]Department of Mathematics, University of California, Davis.

[§]Corresponding author. *E-mail address:* bremer@math.ucdavis.edu.

Date: August 29, 2013.

Problem (2) requires the evaluation of large numbers of singular integrals given on curved regions. Standard mechanisms for addressing this problem, like the Duffy transform followed by Gaussian quadrature, are limited in their applicability (see [7] for a discussion). Most fast solvers (designed to address problem (3)) are based on iterative methods and can fare poorly when the surface under consideration exhibits complicated geometry. For example, a geometry consisting of a $2 \times 2 \times 2$ grid of ellipsoids (1536 discretization points per ellipsoid) bounded by a box two wavelengths in size takes over 1000 GMRES iterations to get the residual below 10^{-3} (see Figure 6 and Remark 6.1). Moreover, iterative methods are unable to efficiently handle problems with multiple right hand sides that arise frequently in design problems.

In this manuscript, we describe our contributions towards providing a remedy to problems (2) and (3). We consider scattering problems that arise in applications such as sonar and radar where the goal is to find the scattered field at a set of locations Γ_{out} generated when an incident wave sent from locations Γ_{in} hit an object or collections of objects (whose surface we denote Γ), cf. Fig. 1(a). The locations of the source and target points are known and typically given by the user. In this context, we present an efficient high-order accurate method for constructing a scattering matrix \mathbf{S} which maps the incident field generated by the given sources on Γ_{in} to a set of induced charges on Γ that generate the scattered field on Γ_{out} . A crucial component to constructing \mathbf{S} is solving a boundary integral equation defined on Γ . A recently developed high order accurate discretization technique for integral equations on surfaces [7, 8] is used. This discretization has proven to be effective for Laplace and low frequency scattering problems. Since many scattering problems of interest involve higher frequencies and complicated geometries, a large number of discretization points are typically required to achieve high accuracy, even for high-order methods. In consequence, it would be infeasible to solve the resulting linear systems using dense linear algebra (e.g., Gaussian elimination). As a robust alternative to the commonly used iterative solvers (whose convergence problems in the present context was previously discussed), we propose the use of a direct solver that exploits the underlying structure of the matrices of interest to construct an approximation to the scattering matrix \mathbf{S} . Direct solvers have several benefits over iterative solvers: (i) the computational cost does not depend on the spectral properties of the system, (ii) each additional incoming field can be processed quickly, (iii) the amount of memory required to store the scattering matrix scales linearly with the number of discretization points.

1.1. Model problems. The solution technique presented is applicable to many boundary integral equations that arise in mathematical physics but we limit our discussion to the boundary value problems presented in this section.

Let Ω denote a scattering body, or a union of scattering bodies; Figures 6, 7, 8, 9, and 11 of Section 6 illustrate some specific domains of interest. Let $\Gamma = \partial\Omega$ denote the surface of the body (or bodies). We are interested in finding the scattered field u off Ω for a given incident wave u_{in} . Specifically this involves solving either an exterior Dirichlet problem for *sound-soft* scatterers

$$\begin{cases} \Delta u + \kappa^2 u = 0 & \text{in } \Omega^c, \\ u = g & \text{on } \Gamma, \\ \frac{\partial u}{\partial r} - i\kappa u = O\left(\frac{1}{r}\right) & \text{as } r \rightarrow \infty, \end{cases} \quad (\text{DBVP})$$

where $g = -u_{\text{in}}$, or an exterior Neumann problem for *sound-hard* scatterers

$$\begin{cases} \Delta u + \kappa^2 u = 0 & \text{in } \Omega^c, \\ \frac{\partial u}{\partial \eta} = g & \text{on } \Gamma, \\ \frac{\partial u}{\partial r} - i\kappa u = O\left(\frac{1}{r}\right) & \text{as } r \rightarrow \infty, \end{cases} \quad (\text{NBVP})$$

where η is the outwards point unit normal vector to Γ , and where $g = -\partial_{\eta} u_{\text{in}}$.

1.2. Outline. The paper proceeds by introducing the integral formulations for each choice of boundary condition in Section 2. Section 3 presents a high order accurate technique for discretizing the integral equations. Section 4 formally defines a scattering matrix and describes a method for constructing an

approximation of the scattering matrix. Section 5 describes an accelerated technique for constructing the scattering matrix that often has asymptotic cost of $O(N^{3/2})$ (providing the scattering body and the wave-length of the radiating field are kept constant as N increases). Section 6 presents the results of several numerical experiments which demonstrate the properties of the proposed solver.

2. INTEGRAL EQUATION FORMULATIONS OF LINEAR ELLIPTIC BOUNDARY VALUE PROBLEMS

The reformulation of the boundary value problems in Section 1.1 as integral equations involves the classical single and double layer operators

$$\begin{aligned} S_\kappa f(x) &= \iint_\Gamma G_\kappa(x, y) f(y) ds(y), \\ D_\kappa f(x) &= \iint_\Gamma \eta_y \cdot \nabla_y G_\kappa(x, y) f(y) ds(y) \quad \text{and} \\ D_\kappa^* f(x) &= \iint_\Gamma \eta_x \cdot \nabla_x G_\kappa(x, y) f(y) ds(y) \end{aligned} \tag{2.1}$$

where $G_\kappa(x, y)$ is the free space Green's function

$$G_\kappa(x, y) = \frac{\exp(i\kappa|x-y|)}{|x-y|}$$

associated with the Helmholtz equation at wavenumber κ and where η_p denotes the outward-pointing unit normal to the surface Γ at the point p .

The reformulation of (DBVP) and (NBVP) is not entirely straightforward because of the possibility of spurious resonances. That is, for certain values of the wavenumber κ , the naïve integral equation formulations of (DBVP) and (NBVP) lead to operators with nontrivial null-spaces, even though the original boundary value problems are well-posed. The wave-numbers κ for which this occurs are called *spurious resonances*. Complications arise due to the existence of multiple solutions at a resonance. An additional (and perhaps more serious) problem is that for values of κ close to resonant values, the integral operators (2.1) become ill-conditioned.

2.1. The boundary integral formulation for Dirichlet boundary values problems. The double layer representation

$$u(x) = D_\kappa \sigma(x).$$

of the solution of (DBVP) leads to the second kind integral equation

$$\frac{1}{2}\sigma(x) + D_\kappa \sigma(x) = g(x).$$

As is well known, this integral equation is not be uniquely solvable in the event that the eigenproblem

$$\begin{aligned} \Delta u + \kappa^2 u &= 0 \quad \text{in } \Omega \\ \frac{\partial u}{\partial \eta} &= 0 \quad \text{on } \Gamma \end{aligned}$$

admits a nontrivial solution. A uniquely solvable integral equation can be obtained by choosing to represent the solution as

$$u(x) = D_\kappa \sigma(x) + i|\kappa| S_\kappa \sigma(x); \tag{2.2}$$

see, for instance, [28, 14]. This leads to the integral equation

$$\frac{1}{2}\sigma(x) + D_\kappa \sigma(x) + i|\kappa| S_\kappa \sigma(x) = g(x), \tag{2.3}$$

which is sometimes referred to as the Combined Field Integral Equation (CFIE). We make use of the CFIE to solve the problem (DBVP) in this article.

2.2. The boundary integral formulation for Neumann boundary values problems. The single layer representation

$$u = S_\kappa \sigma.$$

of the solution of (NBVP) leads to the second kind integral equation

$$\frac{1}{2}\sigma(x) + D_\kappa^* \sigma(x) = g(x). \quad (2.4)$$

Once again the equation (2.4) is not necessarily uniquely solvable — when the eigenproblem

$$\begin{aligned} \Delta u + \kappa^2 u &= 0 && \text{in } \Omega \\ u &= 0 && \text{on } \Gamma \end{aligned} \quad (2.5)$$

has nontrivial solutions (that is, when κ is a resonant wavenumber), the operator appearing on the right-hand side of (2.4) is not invertible. Since the focus of this article is on the discretization of integral operators and the rapid inversion of the resulting matrices, we simply avoid wavenumbers κ for which spurious resonances exist. Appendix A presents a robust integral equation formulation which is suitable for use by the solver of this paper. We forgo its use here in order to reduce the complexity of the presentation.

3. A NYSTRÖM METHOD FOR THE DISCRETIZATION OF INTEGRAL EQUATIONS ON SURFACES

In this work, we use a modified Nyström method for the discretization of the integral operators (2.1) of the preceding section. Our approach differs from standard Nyström methods in two important ways:

- (1) It captures the L^2 action of an operator rather than its pointwise behavior. Among other advantages, this means that the approach results in well-conditioned operators in the case of Lipschitz domains.
- (2) A highly accurate mechanism for evaluating the singular integrals which arise in Nyström discretization whose cost is largely independent of the geometry of the surface is employed. This is in contrast to standard methods for the evaluation of the singular integrals arising from the Nyström discretization of integral operators on surfaces which involve constants which depend on the geometry of the domain.

This manuscript gives a cursory account of the method. A detailed description can be found in [7] and a discussion of the advantages of L^2 discretization over standard Nyström and collocation methods can be found in [6].

3.1. Decompositions and discretization quadratures. We define a decomposition D of a surface $\Gamma \subset \mathbb{R}^3$ to be a finite sequence

$$\{\rho_j : \Delta^1 \rightarrow \Gamma\}_{j=1}^m$$

of smooth mappings given on the simplex

$$\Delta^1 = \{(x, y) \in \mathbb{R}^2 : 0 \leq x \leq 1, 0 \leq y \leq 1 - x\}$$

with non-vanishing Jacobian determinants such that the sets

$$\rho_1(\Delta^1), \rho_2(\Delta^1), \dots, \rho_m(\Delta^1)$$

form a disjoint union of Γ .

We call a quadrature rule $\{x_1, \dots, x_l, w_1, \dots, w_l\}$ on Δ^1 with positive weights which integrates all elements of the space \mathcal{P}_N of polynomials of degree less than or equal to $2N$ on the simplex Δ^1 a *discretization quadrature*. Such a rule is called Gaussian if the length l of the rule is equal to $(N+1) \cdot (N+2)/2$. That is, a Gaussian rule is one whose length is equal to the dimension of the space of polynomials of degree less than or equal to N but which integrates polynomials of degree less than or equal to $2N$. For the most part, Gaussian quadratures on triangles are not available. In the experiments of this paper, we used quadrature rules constructed via a generalization of the procedure of [9]. Table 1 lists the properties of the discretization quadrature rules used in this paper and compare their lengths to the size of a true Gaussian quadrature.

Discretization order	N	4	6	8	10	12	16
Integration order	$2N$	8	12	16	20	24	32
Length of a true Gaussian quadrature	$(N+1)(N+2)/2$	15	28	45	66	91	153
Length of the quadrature that we use	l	17	32	52	82	112	192

TABLE 1. The properties of the quadratures used in this paper. We compare the length of an ideal (hypothetical) Gaussian quadrature to the quadrature that we actually have.

3.2. Discretizations of integral operators. We now associate with a decomposition

$$D = \{\rho_j : \Delta^1 \rightarrow \Gamma\}_{j=1}^m$$

and a discretization quadrature $\{x_1, \dots, x_l, w_1, \dots, w_l\}$ a scheme for the discretization of certain integral operators. We begin by letting S be the subspace of $L^2(\Gamma)$ consisting of all functions f which are pointwise defined on Γ and such that for each $j = 1, \dots, m$ the function

$$f(\rho_j(x)) |d\rho_j(x)^* d\rho_j(x)|^{1/2}$$

is a polynomial of order N on Δ^1 . Denote by P be the projection of $L^2(\Gamma)$ onto the subspace S and let $\Phi : S \rightarrow \mathbb{C}^{ml}$ be a mapping which takes $f \in S$ to a vector with entries

$$f(\rho_j(x_i)) \sqrt{w_i} |d\rho_j(\rho_j(x_i))^* d\rho_j(\rho_j(x_i))|^{1/4}, \quad j = 1, \dots, m, \quad i = 1, \dots, l.$$

The ordering of the entries of $\Phi(f)$ is irrelevant. Suppose that $T : L^2(\Gamma) \rightarrow L^2(\Gamma)$ is of the form

$$Tf(x) = \iint_{\Gamma} K(x, y) f(y) ds(y), \quad (3.1)$$

with K a linear combination of the kernels listed in (2.1). Let A denote the $ml \times ml$ matrix which commutes the discretization of the operator T induced by the specified decomposition and discretization quadrature as specified in the diagram (3.2).

$$\begin{array}{ccc} S \subset L^2(\Gamma) & \xrightarrow{PT} & S \subset L^2(\Gamma) \\ \downarrow \Phi & & \downarrow \Phi \\ \mathbb{C}^{ml} & \xrightarrow{A} & \mathbb{C}^{ml} \end{array} \quad (3.2)$$

3.3. Quadrature. Forming the matrix A which appears the diagram (3.2) is an exercise in numerical quadrature. Let x be a discretization node on Δ and let w be the corresponding weight. The l entries of the matrix A in (3.2) associated with the mapping ρ_j and the point $\rho_i(x)$ must map the scaled values

$$f(\rho_j(x_1)) |d\rho_j^*(x_k) d\rho_j(x_k)|^{1/4} \sqrt{w_k}, \quad k = 1, 2, \dots, l, \quad (3.3)$$

of the function f at the images of the discretization quadrature nodes under ρ_j to the value of the integral

$$\iint_{\Delta^1} K(\rho_i(x), \rho_j(y)) f(\rho_j(y)) |d\rho_j^*(y) d\rho_j(y)|^{1/2} \sqrt{w} dy. \quad (3.4)$$

The manner in which this vector is formed depends on the location of the point $\rho_i(x)$ vis-à-vis the patch $\rho_j(\Delta^1)$. Let B denote the ball of minimum radius containing Γ and $2B$ denote the ball with the same center as B but with twice the radius. Then we say that a point x is well-separated from a surface Γ if x is outside of $2B$, and we say that x is near Γ if x is inside of the ball $2B$ but not on the surface Γ .

3.3.1. Smooth regime. When the point $\rho_i(x)$ is well-separated from $\rho_j(\Delta^1)$ the kernel K is smooth and the discretization quadrature can be used to evaluate the integral (3.4). In this case, we take the l entries of the matrix A to be

$$K(\rho_i(x), \rho_j(x_k)) |d\rho_j^*(x_k) d\rho_j(x_k)|^{1/4} \sqrt{w} \sqrt{w_k}, \quad k = 1, 2, \dots, l. \quad (3.5)$$

3.3.2. *Near regime.* When $\rho_i(x)$ is near $\rho_j(\Delta^1)$, formula (3.5) may not achieve sufficient accuracy. Thus, we adaptively form a quadrature with nodes and weights

$$u_1, \dots, u_m, r_1, \dots, r_m$$

sufficient to evaluate the nearly singular integrals which arise and apply an interpolation matrix I which takes the scaled values (3.3) of the function f at the nodes of the discretization quadrature to its scaled values at the nodes of the adaptive quadrature. The vector of entries of A is given by the product of the vector of values of the kernel evaluated at the adaptive quadrature nodes with the matrix I . As in the scheme of [7], the adaptive quadrature is formed in such a way that the matrix I has special structure — it is a product of block diagonal matrices — which allows it to be applied efficiently.

3.3.3. *Singular regime.* When the target node lies on the surface $\rho_j(\Delta^1)$, the integral (3.4) is singular and its evaluation becomes quite difficult. The typical approach to the evaluation of integrals of this form is to apply the Duffy transform or polar coordinate transform in order to remove the singularity of the kernel. Once this has been done, quadratures for smooth functions (for instance, product Legendre quadratures) are typically applied.

While the such schemes are exponentially convergent, they can be extremely inefficient. The difficulty is that the functions resulting from the change of variables can have poles close to the integration domain. For instance, if a polar coordinate transform is applied, the singular integrals which must be evaluated take the form

$$\int_0^{2\pi} \int_0^{R(\theta)} (q_{-1}(\theta) + q_0(\theta)r + q_1(\theta)r^2 + q_2(\theta)r^3 + \dots) dr d\theta \quad (3.6)$$

where the function $R(\theta)$ gives a parameterization of the integration domain Δ^1 in polar coordinates and each $q_j(\theta)$ is a quotient form

$$\frac{r_j(\theta)}{l(\theta)^\alpha},$$

where α is a positive constant dependent on the order j and the kernel under consideration, $r_j(\theta)$ is a trigonometric polynomial of finite order, and $l(\theta)$ is a function which depends on the parameterization ρ_j . More specifically, if

$$\xi_1 = \frac{\partial \rho_j(x_1, x_2)}{\partial x_1} \quad \text{and} \quad \xi_2 = \frac{\partial \rho_j(x_1, x_2)}{\partial x_2},$$

then the function $l(\theta)$ is given by

$$l(\theta) = \sqrt{\lambda \cos^2(\theta) + 2\gamma \cos(\theta) \sin(\theta) + \lambda^{-1} \sin^2(\theta)}$$

where

$$\cos(\gamma) = \frac{\xi_1 \cdot \xi_2}{|\xi_1| |\xi_2|} \quad \text{and} \quad \lambda = \frac{|\xi_1|}{|\xi_2|}.$$

Obviously, the function $l(\theta)$ can have zeros close to the real axis.

The approach used in this paper — and described in [7] — calls for first applying a change of variables in order to ensure that the mapping ρ_j is conformal at the target node. Then the function $l(\theta)$ is a constant which does not depend on θ and the integrand simplifies considerably. The unfortunate side effect of applying this change of variables is that the integration domain is no longer the standard simplex Δ^1 but rather an arbitrary triangle. Now the integral which must be evaluated takes the form

$$\int_0^{2\pi} \int_0^{\tilde{R}(\theta)} (\tilde{q}_{-1}(\theta) + \tilde{q}_0(\theta)r + \tilde{q}_1(\theta)r^2 + \tilde{q}_2(\theta)r^3 + \dots) dr d\theta \quad (3.7)$$

with the \tilde{q}_j trigonometric polynomials of finite order. The function $\tilde{R}(\theta)$ can have poles close to the integration domain, however. We combat this problem by precomputing a collection of quadrature rules designed to integrate smooth functions on arbitrary triangles and apply them to evaluate the integrals (3.7). The performance of this approach is described in [7]; in one example of that paper, the standard approach of applying a change of variables followed by Gauss-Legendre quadrature required more than

100,000 quadrature nodes to accurately evaluate the singular integrals arising from a boundary value problem while the precomputed rules used here and in that paper required less than 1,000 nodes.

4. SCATTERING MATRICES AS A NUMERICAL TOOL

This section describes the concept of a scattering matrix, first from a physical perspective, and then the corresponding linear algebraic interpretation. It also informally describes how to build approximate scattering matrices computationally. A rigorous description is given in Section 5.

4.1. The scattering operator. We consider a “sound-soft” scattering formulation illustrated in Figure 1. We are given a charge distribution s_{in} on a contour Γ_{in} (the radiation “source”) that generates an incoming field g on the scatterer Γ . The incoming field on the scatterer induces an outgoing field u that satisfies (DBVP). Our objective is to computationally construct u_{out} , the restriction of the outgoing field to the specified contour Γ_{out} , which represents the “antenna.” To solve (DBVP), we use the combined field formulation (2.2). In other words, we represent the outgoing field via a distribution σ of “induced sources” on the contour Γ . The map we seek to evaluate can be expressed mathematically as

$$\begin{array}{ccc} u_{\text{out}} & = & T_{\text{out}} T_{\Gamma}^{-1} T_{\text{in}} s_{\text{in}}, \\ \text{outgoing field on } \Gamma_{\text{out}} & & \text{given charges on } \Gamma_{\text{in}} \end{array} \quad (4.1)$$

where $T_{\text{in}} : L^2(\Gamma_{\text{in}}) \rightarrow L^2(\Gamma)$ is the operator

$$[T_{\text{in}} s_{\text{in}}](x) = \int_{\Gamma_{\text{in}}} G_{\kappa}(x, y) s_{\text{in}}(y) ds(y), \quad x \in \Gamma,$$

$T_{\Gamma} : L^2(\Gamma) \rightarrow L^2(\Gamma)$ is the operator

$$[T_{\Gamma} g](x) = \frac{1}{2}g(x) + D_{\kappa}g(x) + i|\kappa|S_{\kappa}g(x), \quad x \in \Gamma, \quad (4.2)$$

and $T_{\text{out}} : L^2(\Gamma) \rightarrow L^2(\Gamma_{\text{out}})$ is the operator

$$[T_{\text{out}} \sigma](x) = \int_{\Gamma} G_{\kappa}(x, y) \sigma(y) ds(y) + \int_{\partial\Omega} \partial_{\nu_y} G_{\kappa}(x, y) \sigma(y) ds(y), \quad x \in \Gamma_{\text{out}}.$$

Observe that since T_{Γ} is an invertible second kind integral operator, its inverse has singular values bounded away from zero.

Given a finite precision $\varepsilon > 0$, we say that an operator $S : L^2(\Gamma) \rightarrow L^2(\Gamma)$ is a “scattering operator” for Γ to within precision ε when

$$\|T_{\text{out}} T_{\Gamma}^{-1} T_{\text{in}} - T_{\text{out}} S T_{\text{in}}\| \leq \varepsilon.$$

It turns out that the operators T_{in} and T_{out} have rapidly decaying singular values. Let $P_{\text{in}}^{(k)}$ denote the orthogonal projection onto the set of k leading left singular vectors of T_{in} , and let $P_{\text{out}}^{(k)}$ denote the orthogonal projection onto the set of k leading right singular vectors of T_{out} . Then

$$\|T_{\text{out}} P_{\text{out}}^{(k)} - T_{\text{out}}\| \leq \sigma_{k+1}(T_{\text{out}}) \quad \text{and} \quad \|T_{\text{in}} P_{\text{in}}^{(k)} - T_{\text{in}}\| \leq \sigma_{k+1}(T_{\text{in}}).$$

It follows that if we define an approximate scattering operator S via

$$S = P_{\text{out}}^{(k)} T_{\Gamma}^{-1} P_{\text{in}}^{(k)}, \quad (4.3)$$

then the approximation error

$$\|T_{\text{out}} T_{\Gamma}^{-1} T_{\text{in}} - T_{\text{out}} P_{\text{out}}^{(k)} T_{\Gamma}^{-1} P_{\text{in}}^{(k)} T_{\text{in}}\|$$

converges to zero very rapidly as k increases.

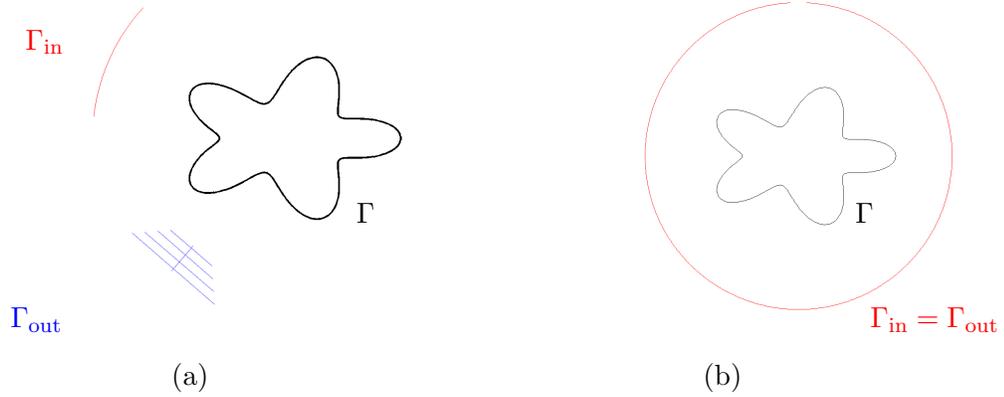


FIGURE 1. Geometry of the scattering problem in Section 4. The scattering surface is Γ . (a) The radiation source Γ_{in} is different from the antenna Γ_{out} . (b) The typical geometry considered with $\Gamma_{\text{in}} = \Gamma_{\text{out}}$. The scattering matrix generated by this geometry is the “complete” scattering matrix in the sense that any incoming field can be handled (as long as the radiation source is not inside Γ_{in}), and the full radiated field can be obtained.

4.2. Definition of a discrete scattering matrix. The continuum scattering problem described in Section 4.1 has a direct linear algebraic analog for the discretized equations. Forming the Nyström discretization of the three operators in (4.1), we obtain the map

$$\begin{array}{ccccc} \mathbf{u}_{\text{out}} & = & \mathbf{A}_{\text{out}} & \mathbf{A}^{-1} & \mathbf{A}_{\text{in}} & \mathbf{s}_{\text{in}}. \\ N_{\text{out}} \times 1 & & N_{\text{out}} \times N & N \times N & N \times N_{\text{in}} & N_{\text{in}} \times 1 \end{array} \quad (4.4)$$

Suppose that \mathbf{A} is an $N \times N$ matrix, and let k_1 denote a bound on the ε -ranks of \mathbf{A}_{out} and \mathbf{A}_{in} . Then form matrices $\hat{\mathbf{U}}_1$ and $\hat{\mathbf{V}}_1$ (the reason for the subscripts will become clear shortly) of size $N \times k_1$ such that

$$\|\mathbf{A}_{\text{in}} - \hat{\mathbf{U}}_1 \hat{\mathbf{U}}_1^\dagger \mathbf{A}_{\text{in}}\| \leq \varepsilon, \quad \text{and} \quad \|\mathbf{A}_{\text{out}} - \mathbf{A}_{\text{out}} (\hat{\mathbf{V}}_1 \hat{\mathbf{V}}_1^\dagger)^*\| \leq \varepsilon$$

where a matrix with the superscript \dagger denotes the pseudo-inverse of the matrix. In other words, the columns of $\hat{\mathbf{U}}_1$ span the column space of \mathbf{A}_{in} and the columns of $\hat{\mathbf{V}}_1$ span the row space of \mathbf{A}_{out} . Then the map (4.4) can be approximated as

$$\begin{array}{ccccc} \mathbf{u}_{\text{out}} & = & (\mathbf{A}_{\text{out}} (\hat{\mathbf{V}}_1^\dagger)^*) & \mathbf{S}_1 & (\hat{\mathbf{U}}_1^\dagger \mathbf{A}_{\text{in}}) & \mathbf{s}_{\text{in}}, \\ N_{\text{out}} \times 1 & & N_{\text{out}} \times k_1 & k_1 \times k_1 & k_1 \times N_{\text{in}} & \end{array} \quad (4.5)$$

where

$$\mathbf{S}_1 = \hat{\mathbf{V}}_1^* \mathbf{A}^{-1} \hat{\mathbf{U}}_1. \quad (4.6)$$

The $k_1 \times k_1$ matrix \mathbf{S}_1 is the discrete analog of the continuum scattering matrix S for Γ defined by (4.3).

4.3. Hierarchical construction of discrete scattering matrices. Observe that direct evaluation of the definition (4.6) of \mathbf{S}_1 via, e.g., Gaussian elimination, would be very costly since \mathbf{A} is dense, and N could potentially be large (much larger than k_1). To avoid this expense, we will build \mathbf{S}_1 via an accelerated hierarchical procedure. As a first step, we let the entire domain $\Gamma = \Gamma_1$ denote the root of the tree, and split Γ_1 into two disjoint halves, cf. Figure 2(b),

$$\Gamma_1 = \Gamma_2 \cup \Gamma_3. \quad (4.7)$$

Let n_2 and n_3 denote the number of collocation points in Γ_2 and Γ_3 , respectively. The idea is now to form scattering matrices \mathbf{S}_2 and \mathbf{S}_3 for the two halves, and then form \mathbf{S}_1 by “merging” these smaller scattering matrices. To formalize, split the matrix \mathbf{A} into blocks to conform with the partitioning (4.7),

$$\mathbf{A} = \begin{bmatrix} \mathbf{A}_{2,2} & \mathbf{A}_{2,3} \\ \mathbf{A}_{3,2} & \mathbf{A}_{3,3} \end{bmatrix}.$$

The off-diagonal blocks $A_{2,3}$ and $A_{3,2}$ have very rapidly decaying singular values, which means that they can be factored

$$\begin{array}{rcccl} A_{2,3} & = & \hat{U}_2 & \tilde{A}_{2,3} & \hat{V}_3^*, \\ n_2 \times n_3 & & n_2 \times k_2 & k_2 \times k_3 & k_3 \times n_3 \end{array} \quad (4.8)$$

$$\begin{array}{rcccl} A_{3,2} & = & \hat{U}_3 & \tilde{A}_{3,2} & \hat{V}_2^*, \\ n_3 \times n_2 & & n_3 \times k_3 & k_3 \times k_2 & k_2 \times n_2 \end{array} \quad (4.9)$$

The matrices $\{\hat{U}_2, \hat{U}_3\}$ and $\{\hat{V}_2, \hat{V}_3\}$ are additionally assumed to span the column spaces of A_{in} and A_{out}^* , in the sense that there exist “small” matrices U_1 and V_1 such that

$$\begin{array}{rcccl} \hat{U}_1 & = & \begin{bmatrix} \hat{U}_2 & 0 \\ 0 & \hat{U}_3 \end{bmatrix} & U_1, & \text{and} & \hat{V}_1 & = & \begin{bmatrix} \hat{V}_2 & 0 \\ 0 & \hat{V}_3 \end{bmatrix} & V_1. \\ N \times k_1 & & N \times (k_2 + k_3) & (k_2 + k_3) \times k_1 & & N \times k_1 & & N \times (k_2 + k_3) & (k_2 + k_3) \times k_1 \end{array} \quad (4.10)$$

Inserting (4.8), (4.9), and (4.10) into (4.6), we find

$$S_1 = V_1^* \begin{bmatrix} \hat{V}_2^* & 0 \\ 0 & \hat{V}_3^* \end{bmatrix} \begin{bmatrix} A_{2,2} & \hat{U}_2 \tilde{A}_{2,3} \hat{V}_3^* \\ \hat{U}_3 \tilde{A}_{3,2} \hat{V}_2^* & A_{3,3} \end{bmatrix}^{-1} \begin{bmatrix} \hat{U}_2 & 0 \\ 0 & \hat{U}_3 \end{bmatrix} U_2. \quad (4.11)$$

Using the Woodbury formula for matrix inversion (see Lemma 5.1), it can be shown that

$$\begin{bmatrix} \hat{V}_2^* & 0 \\ 0 & \hat{V}_3^* \end{bmatrix} \begin{bmatrix} A_{2,2} & \hat{U}_2 \tilde{A}_{2,3} \hat{V}_3^* \\ \hat{U}_3 \tilde{A}_{3,2} \hat{V}_2^* & A_{3,3} \end{bmatrix}^{-1} \begin{bmatrix} \hat{U}_2 & 0 \\ 0 & \hat{U}_3 \end{bmatrix} = \begin{bmatrix} I & S_2 \tilde{A}_{2,3} \\ S_3 \tilde{A}_{3,2} & I \end{bmatrix}^{-1} \begin{bmatrix} S_2 & 0 \\ 0 & S_3 \end{bmatrix} \quad (4.12)$$

where S_2 and S_3 are the scattering matrices for Γ_2 and Γ_3 , respectively. The matrices S_2 and S_3 are given by the following formulas

$$S_2 = \hat{V}_2^* A_{2,2}^{-1} \hat{U}_2, \quad \text{and} \quad S_3 = \hat{V}_3^* A_{3,3}^{-1} \hat{U}_3. \quad (4.13)$$

Combining (4.11) and (4.12), we obtain the formula

$$S_1 = V_1^* \begin{bmatrix} I & S_2 \tilde{A}_{2,3} \\ S_3 \tilde{A}_{3,2} & I \end{bmatrix}^{-1} \begin{bmatrix} S_2 & 0 \\ 0 & S_3 \end{bmatrix} U_1. \quad (4.14)$$

The upside here is that in order to evaluate (4.14), we only need to invert a matrix of size $(k_2 + k_3) \times (k_2 + k_3)$, whereas (4.6) requires inversion of a matrix of size $N \times N$. The downside is that we are still left with the task of inverting $A_{2,2}$ and $A_{3,3}$ to evaluate S_2 and S_3 via (4.13). Now, to reduce the cost of constructing S_2 and S_3 , we further sub-divide

$$\Gamma_2 = \Gamma_4 \cup \Gamma_5, \quad \text{and} \quad \Gamma_3 = \Gamma_6 \cup \Gamma_7,$$

as shown in Figure 2(c). Then it can be shown (see Lemma 5.1) that S_2 and S_3 are given by the formulas

$$S_2 = V_2^* \begin{bmatrix} I & S_4 \tilde{A}_{4,5} \\ S_5 \tilde{A}_{5,4} & I \end{bmatrix}^{-1} \begin{bmatrix} S_4 & 0 \\ 0 & S_5 \end{bmatrix} U_2, \quad \text{and} \quad S_3 = V_3^* \begin{bmatrix} I & S_6 \tilde{A}_{6,7} \\ S_7 \tilde{A}_{7,6} & I \end{bmatrix}^{-1} \begin{bmatrix} S_6 & 0 \\ 0 & S_7 \end{bmatrix} U_3,$$

where

$$S_4 = \hat{V}_4^* A_{4,4}^{-1} \hat{U}_4, \quad S_5 = \hat{V}_5^* A_{5,5}^{-1} \hat{U}_5, \quad S_6 = \hat{V}_6^* A_{6,6}^{-1} \hat{U}_6, \quad S_7 = \hat{V}_7^* A_{7,7}^{-1} \hat{U}_7.$$

The idea is now to continue to recursively split all patches until we get to a point where each patch holds sufficiently few discretization points that its scattering matrix can inexpensively be constructed via, e.g., Gaussian elimination.

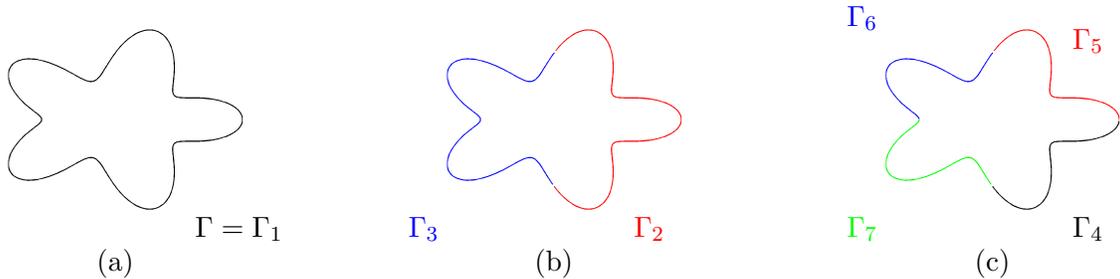


FIGURE 2. Hierarchical partitioning of the domain as described in Section 4.3. (a) The full domain $\Gamma = \Gamma_1$. (b) Partition $\Gamma_1 = \Gamma_2 \cup \Gamma_3$. (c) Partition $\Gamma_2 = \Gamma_4 \cup \Gamma_5$ and $\Gamma_3 = \Gamma_6 \cup \Gamma_7$.

5. THE HIERARCHICAL CONSTRUCTION OF A SCATTERING MATRIX

Section 4 defined the concept of a scattering matrix, and informally outlined a procedure for how to build approximate scattering matrices via a hierarchical procedure. This section provides a more rigorous description, which requires the introduction of some notational machinery. The key concept is that the matrix \mathbf{A} corresponding to the discretized operator T_Γ (defined by (4.2)) can be represented efficiently in a data sparse format that we call *Hierarchically Block Separable (HBS)*, and is closely related to the *Hierarchically Semi-Separable (HSS)* format [32, 29, 31]. Our presentation is in principle self-contained but assumes some familiarity with hierarchically rank-structured matrix formats such as HSS/HBS (or the related \mathcal{H}^2 -matrix format [2, 3]).

Remark 5.1. The method presented in Sections 4.2, 4.3, and this section use the same rank for the incoming and outgoing factorizations purely for simplicity of presentation. In practice this is not required.

5.1. Problem formulation. Let \mathbf{A} denote the $N \times N$ dense matrix obtained upon Nyström discretization (see Section 3) of the operator T_Γ as defined by (4.2). Moreover, suppose that we are given a positive tolerance ε , and two tall thin matrices $\hat{\mathbf{U}}_1$ and $\hat{\mathbf{V}}_1$ satisfying the following:

$\hat{\mathbf{U}}_1$: A well-conditioned matrix of size $N \times k_1$ whose columns span (to within precision ε) the column space of the matrix \mathbf{A}_{in} that maps the given charge distribution on the “radiation source” Γ_{in} to the collocation points on Γ . In other words, any incoming field hitting Γ can be expressed as a linear combination of the columns of $\hat{\mathbf{U}}_1$.

$\hat{\mathbf{V}}_1$: A well-conditioned matrix of size $N \times k_1$ whose columns span (to within precision ε) the row space of the matrix \mathbf{A}_{out} that maps the induced charges on the collocation points on Γ to the collocation points on the “antenna” Γ_{out} . In other words, any field generated on Γ_{out} by induced charges on Γ can be replicated by a charge distribution in the span of $\hat{\mathbf{V}}_1$.

Our objective is then to construct a highly accurate approximation to the scattering matrix

$$\mathbf{S}_1 = \hat{\mathbf{V}}_1^* \mathbf{A}^{-1} \hat{\mathbf{U}}_1. \quad (5.1)$$

5.2. Hierarchical tree. The hierarchical construction of the scattering matrix \mathbf{S}_1 is based on a binary tree structure of patches on the scattering surface Γ . We let the full domain Γ form the root of the tree, and give it the index 1, $\Gamma_1 = \Gamma$. We next split the root into two roughly equi-sized patches Γ_2 and Γ_3 so that $\Gamma_1 = \Gamma_2 \cup \Gamma_3$. The full tree is then formed by continuing to subdivide any patch that holds more than some preset fixed number of collocation points. A *leaf* is a node in the tree corresponding to a patch that never got split. For a non-leaf node τ , its *children* are the two nodes α and β such that $\Gamma_\tau = \Gamma_\alpha \cup \Gamma_\beta$, and τ is then the *parent* of α and β . Two boxes with the same parent are called *siblings*.

Let $I = I_1 = [1, 2, 3, \dots, N]$ denote the full list of indices for the collocation points on Γ , and let for any node τ the index vector $I_\tau \subseteq I$ mark the set of collocation nodes inside Γ_τ . Let n_τ denote the number of nodes in I_τ .

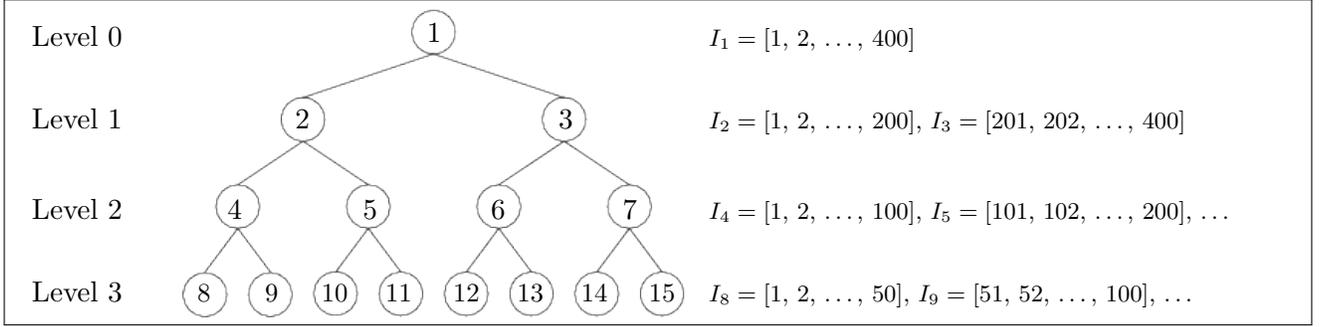


FIGURE 3. Numbering of nodes in a fully populated binary tree with $L = 3$ levels. The root is the original index vector $I = I_1 = [1, 2, \dots, 400]$.

5.3. Hierarchically Block Separable (HBS) matrices. We say that a dense $N \times N$ matrix A is HBS with respect to a given tree partitioning of its index vector $I = [1, 2, 3, \dots, N]$ if for every node τ in the tree, there exist basis matrix \hat{U}_τ and \hat{V}_τ with the following properties:

- (1) For every sibling pair $\{\alpha, \beta\}$, the corresponding off-diagonal block $A(I_\alpha, I_\beta)$ admits (up to precision ε) the factorization

$$A(I_\alpha, I_\beta) = \begin{array}{ccc} \hat{U}_\alpha & \tilde{A}_{\alpha, \beta} & \hat{V}_\beta, \\ n_\alpha \times n_\beta & n_\alpha \times k_\alpha & k_\alpha \times k_\beta & k_\beta \times n_\beta \end{array} \quad (5.2)$$

where $k_\alpha < n_\alpha$ and $k_\beta < n_\beta$.

- (2) For every sibling pair $\{\alpha, \beta\}$ with parent τ , there exist matrices U_τ and V_τ such that

$$\hat{U}_\tau = \begin{array}{cc} \left[\begin{array}{cc} \hat{U}_\alpha & 0 \\ 0 & \hat{U}_\beta \end{array} \right] & U_\tau, \\ n_\tau \times k_\tau & n_\tau \times (k_\alpha + k_\beta) & (k_\alpha + k_\beta) \times k_\tau \end{array} \quad (5.3)$$

$$\hat{V}_\tau = \begin{array}{cc} \left[\begin{array}{cc} \hat{V}_\alpha & 0 \\ 0 & \hat{V}_\beta \end{array} \right] & V_\tau. \\ n_\tau \times k_\tau & n_\tau \times (k_\alpha + k_\beta) & (k_\alpha + k_\beta) \times k_\tau \end{array} \quad (5.4)$$

For notational convenience, we set for every leaf node τ

$$U_\tau = \hat{U}_\tau, \quad \text{and} \quad V_\tau = \hat{V}_\tau.$$

Note that every basis matrix U_τ and V_τ is small. The point is that the matrix A is fully specified by giving the following for every node τ :

- (i) the two small matrices U_τ , and V_τ ,
- (ii) if τ is a leaf, the diagonal block $A(I_\tau, I_\tau)$, and
- (iii) if τ is a parent node with children $\{\alpha, \beta\}$, the sibling interaction matrices $\tilde{A}_{\alpha, \beta}$, $\tilde{A}_{\beta, \alpha}$.

*Particular attention should be paid to the fact that the long basis matrices \hat{U}_τ and \hat{V}_τ are **never** formed — these were introduced merely to facilitate the derivation of the HBS representation.* Table 4 summarizes the factors required, and Algorithm I describes how for any vector σ the matrix-vector product $u = A\sigma$ can be computed if the HBS factors of A are given.

Remark 5.2 (Meaning of basis matrices \hat{U}_τ , \hat{V}_τ , U_τ , V_τ). *Let us describe the heuristic meaning of the “tall thin” basis matrix \hat{U}_τ . Conditions (5.2) and (5.3) together imply that the columns of \hat{U}_τ span the column space of the off-diagonal block $A(I_\tau, I_\tau^c)$, as well as the columns of $A_{\text{in}}(I_\tau, \cdot)$. We can write the restriction of the local equilibrium equation on Γ_τ as*

$$A(I_\tau, I_\tau) \sigma(I_\tau) = A_{\text{in}}(I_\tau, \cdot) s_{\text{in}} - A(I_\tau, I_\tau^c) \sigma(I_\tau^c) \quad (5.5)$$

Exploiting that the columns of \hat{U}_τ span the column spaces of both matrices on the right hand side of (5.5), we can rewrite the equation as

$$A(I_\tau, I_\tau) \sigma(I_\tau) = \hat{U}_\tau (\tilde{u}_\tau - \tilde{v}_\tau), \quad (5.6)$$

	Name:	Size:	Function:
For each leaf node τ :	$\mathbf{A}(I_\tau, I_\tau)$	$n_\tau \times n_\tau$	Diagonal block.
	\mathbf{U}_τ	$n_\tau \times k_\tau$	Basis for the columns in $\mathbf{A}(I_\tau, I_\tau^c)$.
	\mathbf{V}_τ	$n_\tau \times k_\tau$	Basis for the rows in $\mathbf{A}(I_\tau^c, I_\tau)$.
For each parent node τ with children $\{\alpha, \beta\}$:	$\tilde{\mathbf{A}}_{\alpha, \beta}$	$k_\alpha \times k_\beta$	Sibling interaction matrix.
	$\tilde{\mathbf{A}}_{\beta, \alpha}$	$k_\beta \times k_\alpha$	Sibling interaction matrix.
	\mathbf{U}_τ	$(k_\alpha + k_\beta) \times k_\tau$	Basis for the (reduced) incoming fields on τ .
	\mathbf{V}_τ	$(k_\alpha + k_\beta) \times k_\tau$	Basis for the (reduced) outgoing fields from τ .

FIGURE 4. An HBS matrix \mathbf{A} associated with a tree \mathcal{T} is fully specified if the factors listed above are provided.

ALGORITHM I: HBS MATRIX-VECTOR MULTIPLY

Given a vector $\boldsymbol{\sigma}$ and a matrix \mathbf{A} in HBS format, compute $\mathbf{u} = \mathbf{A} \boldsymbol{\sigma}$.

It is assumed that the nodes are ordered so that if τ is the parent of σ , then $\tau < \sigma$.

```

for  $\tau = N_{\text{boxes}}, N_{\text{boxes}} - 1, \dots, 2$ 
  if  $\tau$  is a leaf
     $\tilde{\boldsymbol{\sigma}}_\tau = \mathbf{V}_\tau^* \boldsymbol{\sigma}(I_\tau)$ .
  else
     $\tilde{\boldsymbol{\sigma}}_\tau = \mathbf{V}_\tau^* \begin{bmatrix} \tilde{\boldsymbol{\sigma}}_\alpha \\ \tilde{\boldsymbol{\sigma}}_\beta \end{bmatrix}$ , where  $\alpha$  and  $\beta$  denote the children of  $\tau$ .
  end if
end for

 $\tilde{\mathbf{u}}_1 = 0$ 
for  $\tau = 1, 2, 3, \dots, N_{\text{boxes}}$ 
  if  $\tau$  is a parent
     $\begin{bmatrix} \tilde{\mathbf{u}}_\alpha \\ \tilde{\mathbf{u}}_\beta \end{bmatrix} = \mathbf{U}_\tau \tilde{\mathbf{u}}_\tau + \begin{bmatrix} 0 & \tilde{\mathbf{A}}_{\alpha, \beta} \\ \tilde{\mathbf{A}}_{\beta, \alpha} & 0 \end{bmatrix} \begin{bmatrix} \tilde{\boldsymbol{\sigma}}_\alpha \\ \tilde{\boldsymbol{\sigma}}_\beta \end{bmatrix}$ , where  $\alpha$  and  $\beta$  denote the children of  $\tau$ .
  else
     $\mathbf{u}(I_\tau) = \mathbf{U}_\tau \tilde{\mathbf{u}}_\tau + \mathbf{A}(I_\tau, I_\tau) \boldsymbol{\sigma}(I_\tau)$ .
  end if
end for

```

where $\tilde{\mathbf{u}}_\tau = \hat{\mathbf{U}}_\tau^\dagger \mathbf{A}_{\text{in}}(I_\tau, \cdot) \mathbf{s}_{\text{in}}$ is an efficient representation of the local incoming field, and where $\tilde{\mathbf{v}}_\tau = \hat{\mathbf{U}}_\tau^\dagger \mathbf{A}(I_\tau, I_\tau^c) \boldsymbol{\sigma}(I_\tau^c)$ is an efficient representation of the field on Γ_τ caused by charges on $\Gamma \setminus \Gamma_\tau$. The basis matrix $\hat{\mathbf{V}}_\tau$ analogously provides an efficient way to represent the outgoing fields on Γ_τ . The “small” matrices \mathbf{U}_τ and \mathbf{V}_τ then provide compact representations of $\hat{\mathbf{U}}_\tau$ and $\hat{\mathbf{V}}_\tau$. Observe for instance that

$$\hat{\mathbf{U}}_1 = \begin{bmatrix} \hat{\mathbf{U}}_2 & 0 \\ 0 & \hat{\mathbf{U}}_3 \end{bmatrix} \mathbf{U}_1 = \begin{bmatrix} \hat{\mathbf{U}}_4 & 0 & 0 & 0 \\ 0 & \hat{\mathbf{U}}_5 & 0 & 0 \\ 0 & 0 & \hat{\mathbf{U}}_6 & 0 \\ 0 & 0 & 0 & \hat{\mathbf{U}}_7 \end{bmatrix} \begin{bmatrix} \mathbf{U}_2 & 0 \\ 0 & \mathbf{U}_3 \end{bmatrix} \mathbf{U}_1.$$

These concepts are described in further detail in Appendix B.

Remark 5.3. (Exact versus numerical rank) In practical computations, the off-diagonal blocks of discretized integral operators only have low numerical (as opposed to exact) rank. However, in the present context, the singular values tend to decay exponentially fast, which means that the truncation error can often be rendered close to machine precision. We do not in this manuscript provide a rigorous error

analysis, but merely observe that across a broad range of numerical experiments, we did not see any error aggregation; instead, the local truncation errors closely matched the global error.

5.4. Hierarchical construction of a scattering matrix. For any node τ , define its associated scattering matrix via

$$\mathbf{S}_\tau = \hat{\mathbf{V}}_\tau^* (\mathbf{A}(I_\tau, I_\tau))^{-1} \hat{\mathbf{U}}_\tau. \quad (5.7)$$

The following lemma states that the scattering matrix for a parent node τ with children α and β can be formed inexpensively from the scattering matrices \mathbf{S}_α and \mathbf{S}_β , along with the sibling interaction matrices $\tilde{\mathbf{A}}_{\alpha,\beta}$ and $\tilde{\mathbf{A}}_{\beta,\alpha}$.

LEMMA 5.1. *Let τ be a node with children α and β . Then*

$$\mathbf{S}_\tau = \mathbf{V}_\tau^* \begin{bmatrix} \mathbf{I} & \mathbf{S}_\alpha \tilde{\mathbf{A}}_{\alpha,\beta} \\ \mathbf{S}_\beta \tilde{\mathbf{A}}_{\beta,\alpha} & \mathbf{I} \end{bmatrix}^{-1} \begin{bmatrix} \mathbf{S}_\alpha & \mathbf{0} \\ \mathbf{0} & \mathbf{S}_\beta \end{bmatrix} \mathbf{U}_\tau. \quad (5.8)$$

Lemma 5.1 immediately leads to an $O(N^{3/2})$ algorithm for computing the scattering matrix \mathbf{S}_1 . A pseudo-code description is given as Algorithm II.

Proof of Lemma 5.1: Using formula (5.2) for the factorization of sibling interaction matrices we get

$$\begin{aligned} \mathbf{A}(I_\tau, I_\tau) &= \begin{bmatrix} \mathbf{A}_{\alpha,\alpha} & \mathbf{A}_{\alpha,\beta} \\ \mathbf{A}_{\beta,\alpha} & \mathbf{A}_{\beta,\beta} \end{bmatrix} = \begin{bmatrix} \mathbf{A}_{\alpha,\alpha} & \mathbf{0} \\ \mathbf{0} & \mathbf{A}_{\beta,\beta} \end{bmatrix} + \begin{bmatrix} \mathbf{0} & \hat{\mathbf{U}}_\alpha \tilde{\mathbf{A}}_{\alpha,\beta} \hat{\mathbf{V}}_\beta^* \\ \hat{\mathbf{U}}_\beta \tilde{\mathbf{A}}_{\beta,\alpha} \hat{\mathbf{V}}_\alpha^* & \mathbf{0} \end{bmatrix} \\ &= \begin{bmatrix} \mathbf{A}_{\alpha,\alpha} & \mathbf{0} \\ \mathbf{0} & \mathbf{A}_{\beta,\beta} \end{bmatrix} + \begin{bmatrix} \mathbf{0} & \hat{\mathbf{U}}_\alpha \tilde{\mathbf{A}}_{\alpha,\beta} \\ \hat{\mathbf{U}}_\beta \tilde{\mathbf{A}}_{\beta,\alpha} & \mathbf{0} \end{bmatrix} \begin{bmatrix} \hat{\mathbf{V}}_\alpha^* & \mathbf{0} \\ \mathbf{0} & \hat{\mathbf{V}}_\beta^* \end{bmatrix} \\ &= \left(\mathbf{I} + \begin{bmatrix} \mathbf{0} & \hat{\mathbf{U}}_\alpha \tilde{\mathbf{A}}_{\alpha,\beta} \\ \hat{\mathbf{U}}_\beta \tilde{\mathbf{A}}_{\beta,\alpha} & \mathbf{0} \end{bmatrix} \begin{bmatrix} \hat{\mathbf{V}}_\alpha^* \mathbf{A}_{\alpha,\alpha}^{-1} & \mathbf{0} \\ \mathbf{0} & \hat{\mathbf{V}}_\beta^* \mathbf{A}_{\beta,\beta}^{-1} \end{bmatrix} \right) \begin{bmatrix} \mathbf{A}_{\alpha,\alpha} & \mathbf{0} \\ \mathbf{0} & \mathbf{A}_{\beta,\beta} \end{bmatrix} \\ &= (\mathbf{I} + \mathbf{E}\mathbf{F}^*) \begin{bmatrix} \mathbf{A}_{\alpha,\alpha} & \mathbf{0} \\ \mathbf{0} & \mathbf{A}_{\beta,\beta} \end{bmatrix}, \end{aligned} \quad (5.9)$$

where we defined

$$\mathbf{E} = \begin{bmatrix} \mathbf{0} & \hat{\mathbf{U}}_\alpha \tilde{\mathbf{A}}_{\alpha,\beta} \\ \hat{\mathbf{U}}_\beta \tilde{\mathbf{A}}_{\beta,\alpha} & \mathbf{0} \end{bmatrix} \quad \text{and} \quad \mathbf{F}^* = \begin{bmatrix} \hat{\mathbf{V}}_\alpha^* \mathbf{A}_{\alpha,\alpha}^{-1} & \mathbf{0} \\ \mathbf{0} & \hat{\mathbf{V}}_\beta^* \mathbf{A}_{\beta,\beta}^{-1} \end{bmatrix}.$$

Formula (5.9) shows that $\mathbf{A}(I_\tau, I_\tau)$ can be written as a product between a low-rank perturbation to the identity, and a block-diagonal matrix. The Woodbury formula for inversion of a low-rank perturbation of the identity states that $(\mathbf{I} + \mathbf{E}\mathbf{F}^*)^{-1} = \mathbf{I} - \mathbf{E}(\mathbf{I} + \mathbf{F}^*\mathbf{E})^{-1}\mathbf{F}^*$. To exploit this formula, we first define

$$\mathbf{Z}_\tau \stackrel{\text{def}}{=} (\mathbf{I} + \mathbf{F}^*\mathbf{E})^{-1} = \left(\mathbf{I} + \begin{bmatrix} \mathbf{0} & \hat{\mathbf{V}}_\alpha^* \mathbf{A}_{\alpha,\alpha}^{-1} \hat{\mathbf{U}}_\alpha \tilde{\mathbf{A}}_{\alpha,\beta} \\ \hat{\mathbf{V}}_\beta^* \mathbf{A}_{\beta,\beta}^{-1} \hat{\mathbf{U}}_\beta \tilde{\mathbf{A}}_{\beta,\alpha} & \mathbf{0} \end{bmatrix} \right)^{-1} = \begin{bmatrix} \mathbf{I} & \mathbf{S}_\alpha \tilde{\mathbf{A}}_{\alpha,\beta} \\ \mathbf{S}_\beta \tilde{\mathbf{A}}_{\beta,\alpha} & \mathbf{I} \end{bmatrix}^{-1}. \quad (5.10)$$

Then

$$\mathbf{A}(I_\tau, I_\tau)^{-1} = \begin{bmatrix} \mathbf{A}_{\alpha,\alpha}^{-1} & \mathbf{0} \\ \mathbf{0} & \mathbf{A}_{\beta,\beta}^{-1} \end{bmatrix} \left(\mathbf{I} - \begin{bmatrix} \mathbf{0} & \hat{\mathbf{U}}_\alpha \tilde{\mathbf{A}}_{\alpha,\beta} \\ \hat{\mathbf{U}}_\beta \tilde{\mathbf{A}}_{\beta,\alpha} & \mathbf{0} \end{bmatrix} \mathbf{Z}_\tau \begin{bmatrix} \hat{\mathbf{V}}_\alpha^* \mathbf{A}_{\alpha,\alpha}^{-1} & \mathbf{0} \\ \mathbf{0} & \hat{\mathbf{V}}_\beta^* \mathbf{A}_{\beta,\beta}^{-1} \end{bmatrix} \right). \quad (5.11)$$

Combining (5.7) with (5.11) and inserting the condition on nestedness of the basis matrices (5.3), we get

$$\begin{aligned} \mathbf{S}_\tau &= \mathbf{V}_\tau^* \begin{bmatrix} \hat{\mathbf{V}}_\alpha^* & \mathbf{0} \\ \mathbf{0} & \hat{\mathbf{V}}_\beta^* \end{bmatrix} \begin{bmatrix} \mathbf{A}_{\alpha,\alpha}^{-1} & \mathbf{0} \\ \mathbf{0} & \mathbf{A}_{\beta,\beta}^{-1} \end{bmatrix} \left(\mathbf{I} - \begin{bmatrix} \mathbf{0} & \hat{\mathbf{U}}_\alpha \tilde{\mathbf{A}}_{\alpha,\beta} \\ \hat{\mathbf{U}}_\beta \tilde{\mathbf{A}}_{\beta,\alpha} & \mathbf{0} \end{bmatrix} \mathbf{Z}_\tau \begin{bmatrix} \hat{\mathbf{V}}_\alpha^* \mathbf{A}_{\alpha,\alpha}^{-1} & \mathbf{0} \\ \mathbf{0} & \hat{\mathbf{V}}_\beta^* \mathbf{A}_{\beta,\beta}^{-1} \end{bmatrix} \right) \begin{bmatrix} \hat{\mathbf{U}}_\alpha & \mathbf{0} \\ \mathbf{0} & \hat{\mathbf{U}}_\beta \end{bmatrix} \mathbf{U}_\tau \\ &= \mathbf{V}_\tau^* \left(\mathbf{I} - \begin{bmatrix} \mathbf{0} & \mathbf{S}_\alpha \tilde{\mathbf{A}}_{\alpha,\beta} \\ \mathbf{S}_\beta \tilde{\mathbf{A}}_{\beta,\alpha} & \mathbf{0} \end{bmatrix} \mathbf{Z}_\tau \right) \begin{bmatrix} \mathbf{S}_\alpha & \mathbf{0} \\ \mathbf{0} & \mathbf{S}_\beta \end{bmatrix} \mathbf{U}_\tau \\ &= \mathbf{V}_\tau^* \left(\mathbf{Z}_\tau^{-1} - \begin{bmatrix} \mathbf{0} & \mathbf{S}_\alpha \tilde{\mathbf{A}}_{\alpha,\beta} \\ \mathbf{S}_\beta \tilde{\mathbf{A}}_{\beta,\alpha} & \mathbf{0} \end{bmatrix} \right) \mathbf{Z}_\tau \begin{bmatrix} \mathbf{S}_\alpha & \mathbf{0} \\ \mathbf{0} & \mathbf{S}_\beta \end{bmatrix} \mathbf{U}_\tau = \mathbf{V}_\tau^* \mathbf{Z}_\tau \begin{bmatrix} \mathbf{S}_\alpha & \mathbf{0} \\ \mathbf{0} & \mathbf{S}_\beta \end{bmatrix} \mathbf{U}_\tau. \end{aligned}$$

So (5.8) holds. \square

ALGORITHM II: HIERARCHICAL COMPUTATION OF SCATTERING MATRICES

```

for  $\tau = N_{\text{boxes}} : (-1) : 1$ 
  if  $\tau$  is a leaf
     $S_\tau = V_\tau^* A(I_\tau, I_\tau)^{-1} U_\tau.$ 
  else
     $S_\tau = V_\tau^* \begin{bmatrix} I & S_\alpha \tilde{A}_{\alpha,\beta} \\ S_\beta \tilde{A}_{\beta,\alpha} & I \end{bmatrix}^{-1} \begin{bmatrix} S_\alpha & 0 \\ 0 & S_\beta \end{bmatrix} U_\tau,$  where  $\alpha$  and  $\beta$  denote the children of  $\tau$ .
  end if
end for

```

Remark 5.4 (Extension to a full direct solver). *The scattering matrix formulation described in this section can be viewed as an efficient direct solver for the equation $A\sigma = \mathbf{u}$ for the particular case where the load vector \mathbf{u} belongs to the low-dimensional subspace spanned by the columns of the given matrix \hat{U}_1 . A general direct solver for applying A^{-1} to an arbitrary vector for an HBS matrix A (under some general conditions on invertibility) can be obtained by slight variations to the simplified scheme described here, see Appendix B for details.*

5.5. Efficient construction of the HBS representation and recursive skeletonization. In this section, we describe a technique for computing the HBS representation of a discretized boundary integral operator via a procedure described in [26, 17], and sometimes referred to as *recursive skeletonization* [24].

5.5.1. *The interpolative decomposition (ID).* The *interpolative decomposition* (ID) [13] is a low-rank factorization in which a matrix of rank k is expressed by using a selection of k of its columns (rows) as a basis for its column (row) space. To be precise, let X be matrix of size $m \times n$ and of rank k . Then it is possible to determine an index vector $J_{\text{colskel}} \subseteq \{1, 2, 3, \dots, n\}$ of length k such that

$$\begin{matrix} X \\ m \times n \end{matrix} = \begin{matrix} X(:, J_{\text{colskel}}) \\ m \times k \end{matrix} \begin{matrix} V^* \\ k \times n \end{matrix}, \quad (5.12)$$

where V is an $n \times k$ matrix that contains the $k \times k$ unit matrix as a sub-matrix,

$$V(J_{\text{colskel}}, :) = I.$$

Moreover, no entry of V is of magnitude greater than 1, which implies that V is well-conditioned. We call the index vector J_{colskel} the *column skeleton* of X .

Computing the optimal ID of a general matrix is hard [18], but if we slightly relax the restriction on the basis matrix V to allow its entries to be bounded by, say, 2 instead of one, then very efficient algorithms are available [13, 18, 21].

5.5.2. *Interpolative decompositions and HBS matrices.* It is highly advantageous to use the ID to factor the off-diagonal blocks in the HBS representation of a matrix. The first step in the construction is to form an extended system matrix

$$B = \begin{bmatrix} A & A_{\text{in}} \\ A_{\text{out}} & 0 \end{bmatrix}.$$

For each leaf node τ , identify an ID basis matrix U_τ and a subset $\tilde{I}_\tau^{\text{in}} \subset I_\tau$ such that

$$\begin{matrix} B(I_\tau, I_\tau^c) \\ n_\tau \times (N + N_{\text{in}} - n_\tau) \end{matrix} = \begin{matrix} U_\tau \\ n_\tau \times k_\tau \end{matrix} \begin{matrix} B(\tilde{I}_\tau^{\text{in}}, I_\tau^c) \\ k_\tau \times (N + N_{\text{in}} - n_\tau) \end{matrix}. \quad (5.13)$$

The vector $\tilde{I}_\tau^{\text{in}}$ marks a subset of the collocation points in Γ that we call the *skeleton points* of τ . These points have the property that all interactions between Γ_τ and the rest of the geometry can be done via

evaluation through these points. Next let τ be a node whose children $\{\alpha, \beta\}$ are leaves. Then we find the skeleton vector $\tilde{I}_\tau^{\text{in}}$ and the basis matrix \mathbf{U}_τ by simply factoring the matrix $\mathbf{B}([\tilde{I}_\alpha^{\text{in}}, \tilde{I}_\beta^{\text{in}}], I_\tau^{\text{c}})$ to form its ID

$$\begin{array}{ccc} \mathbf{B}([\tilde{I}_\alpha^{\text{in}}, \tilde{I}_\beta^{\text{in}}], I_\tau^{\text{c}}) & = & \mathbf{U}_\tau \quad \mathbf{B}(\tilde{I}_\tau^{\text{in}}, I_\tau^{\text{c}}). \\ (k_\alpha + k_\beta) \times (N + N_{\text{in}} - n_\tau) & & (k_\alpha + k_\beta) \times k_\tau \quad k_\tau \times (N + N_{\text{in}} - n_\tau) \end{array} \quad (5.14)$$

All remaining skeleton sets and basis matrices $\{\mathbf{U}_\tau\}_\tau$ are found by simply continuing this process up through the tree.

The process for determining the basis matrices $\{\mathbf{V}_\tau\}_\tau$ is analogous. For each leaf form the factorization

$$\begin{array}{ccc} \mathbf{B}(I_\tau^{\text{c}}, I_\tau) & = & \mathbf{B}(\tilde{I}_\tau^{\text{out}}, I_\tau^{\text{c}}) \quad \mathbf{V}_\tau^*. \\ (N + N_{\text{out}} - n_\tau) \times n_\tau & & (N + N_{\text{out}} - n_\tau) \times k_\tau \quad k_\tau \times n_\tau \end{array} \quad (5.15)$$

For a parent τ with children $\{\alpha, \beta\}$ (whose skeletons are available), factor

$$\begin{array}{ccc} \mathbf{B}(I_\tau^{\text{c}}, [\tilde{I}_\alpha^{\text{out}}, \tilde{I}_\beta^{\text{out}}]) & = & \mathbf{B}(I_\tau^{\text{c}}, \tilde{I}_\tau^{\text{out}}) \quad \mathbf{V}_\tau^*. \\ (N + N_{\text{out}} - n_\tau) \times (k_\alpha + k_\beta) & & (N + N_{\text{out}} - n_\tau) \times k_\tau \quad k_\tau \times (k_\alpha + k_\beta) \end{array} \quad (5.16)$$

Once the skeletons $\{\tilde{I}_\tau^{\text{in}}, \tilde{I}_\tau^{\text{out}}\}$ have been determined for all nodes, the sibling interaction matrices for a sibling pair $\{\alpha, \beta\}$ are given simply by

$$\tilde{\mathbf{A}}_{\alpha, \beta} = \mathbf{A}(\tilde{I}_\alpha^{\text{out}}, \tilde{I}_\beta^{\text{in}}), \quad \text{and} \quad \tilde{\mathbf{A}}_{\beta, \alpha} = \mathbf{A}(\tilde{I}_\beta^{\text{out}}, \tilde{I}_\alpha^{\text{in}}). \quad (5.17)$$

The formulas (5.17) are very useful since they make it extremely cheap to form the sibling interaction matrices — they are merely submatrices of the original system matrix!

Remark 5.5. It is often convenient to force the incoming and the outgoing skeletons to be identical, $\tilde{I}_\tau^{\text{out}} = \tilde{I}_\tau^{\text{in}}$. For instance, in order to determine a single set of skeleton nodes \tilde{I}_τ for a leaf τ , we can combine (5.13) and (5.14) to the single factorization

$$\begin{array}{ccc} \begin{bmatrix} \mathbf{B}(I_\tau^{\text{c}}, I_\tau) \\ \mathbf{B}(I_\tau, I_\tau^{\text{c}})^* \end{bmatrix} & = & \begin{bmatrix} \mathbf{B}(I_\tau^{\text{c}}, \tilde{I}_\tau) \\ \mathbf{B}(\tilde{I}_\tau, I_\tau^{\text{c}})^* \end{bmatrix} \quad \mathbf{V}_\tau^*. \\ (2N + N_{\text{out}} + N_{\text{in}} - n_\tau) \times n_\tau & & (2N + N_{\text{out}} + N_{\text{in}} - n_\tau) \times k_\tau \quad k_\tau \times n_\tau \end{array} \quad (5.18)$$

Then simply set $\mathbf{U}_\tau = \mathbf{V}_\tau$.

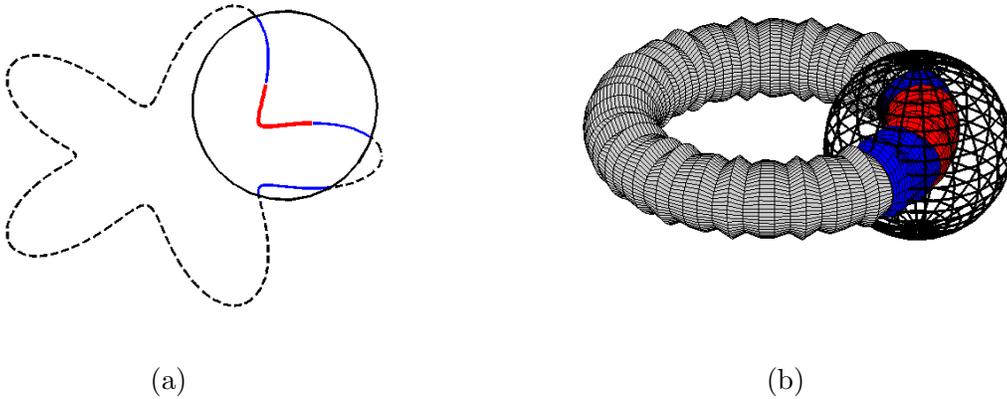


FIGURE 5. Illustration of the geometric objects Γ_τ , Γ_{near} and Γ_{proxy} used for the Green's representation technique for a two-dimensional (a) and three-dimensional surface (b). Γ_τ is given in red, Γ_{near} is given in blue, and Γ_{proxy} is given in black.

5.5.3. *The Green’s representation technique.* The compression technique described in Section 5.5.2 would be prohibitively expensive if executed as stated, since it requires the formation and factorization of very large matrices. For instance, equation (5.13) requires the factorization of a matrix of size $n_\tau \times (N + N_{\text{in}}1 - n_\tau)$. It turns out that these computations can be localized by exploiting that the column space of the matrix $\mathbf{B}(I_\tau, I_\tau^c)$ (which is what we need to span) expresses merely a set of solutions to the homogeneous Helmholtz equation. This means that we can use representation techniques from potential theory to replace all “far-field” interactions by interactions with a small proxy surface Γ_{proxy} that encloses the patch Γ_τ that we seek to compress. All discretization points inside on Γ_{proxy} but not in Γ_τ are labeled Γ_{near} . Figure 5 illustrates the geometry used in the Green’s representation technique for a two-dimensional contour and a three surface that is defined rigorously in Section 6.2. This technique is described in detail in Section 6.2 of [15]. For the experiments in this paper, it is sufficient to let Γ_{proxy} be a sphere with a radius that is twice the radius of the smallest ball containing Γ_τ . (In other words, Γ_{proxy} is the surface of the ball “ $2B$ ” defined in Section 3.3.)

6. NUMERICAL EXPERIMENTS.

We now present the results of a number of experiments conducted to measure the performance of the approach of this paper. All code was written in Fortran 77 and compiled with the Intel Fortran Compiler version 12.1. The experiments were carried out on a workstation equipped with 12 Intel Xeon processor cores running at 3.47 GHz and 192 GB of RAM. As a point of reference, we note that MATLAB requires approximately 35 seconds to invert a 8192×8192 matrix with complex-valued i.i.d. Gaussian entries on this workstation.

6.1. **Sound-soft scattering from grids of ellipsoids.** The purpose of this first set of experiments was to measure the growth in the running time of the algorithm of Section 4 as the geometry of the scatterer becomes more complicated and more nodes are required to discretize it. To that end we considered the exterior Dirichlet problem (DBVP) on a collection of domains consisting of grids of ellipsoids of various sizes and eccentricities. In each experiment, the ellipsoids were arranged on an $n \times n \times n$ grid with center points

$$(1.25 + 2.5i, 1.25 + 2.5j, 1.25 + 2.5k), \quad 1 \leq i, j, k \leq n.$$

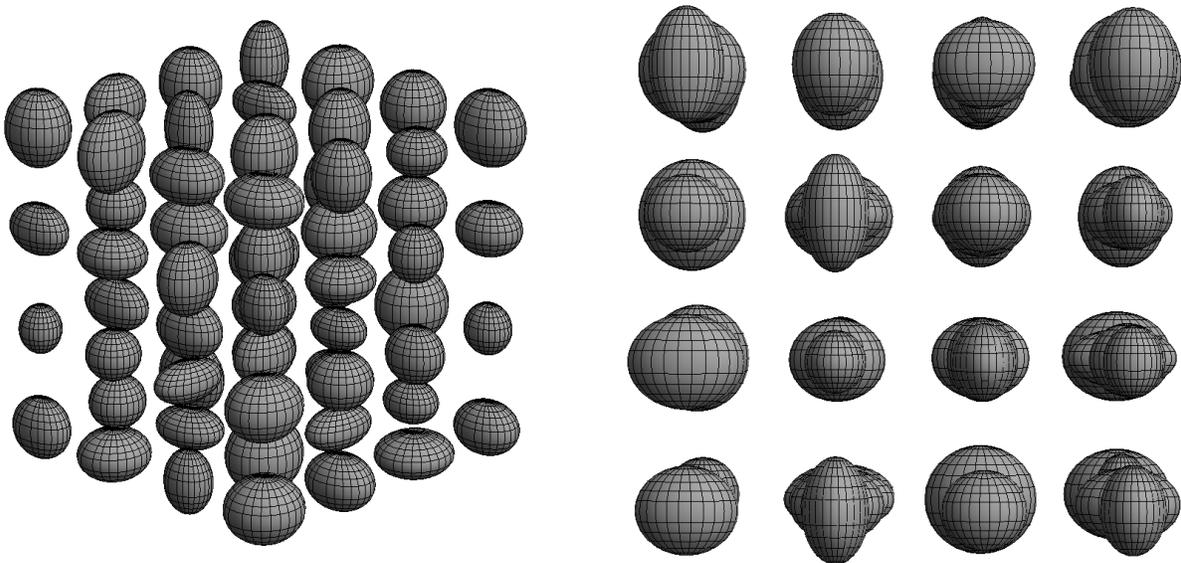


FIGURE 6. Two views of one of the “grid of ellipsoids” domains of Section 6.1. Domains of this type pose significant challenges for iterative solvers.

Grid dimensions	N	$N_{\text{out}} \times N_{\text{in}}$	T	E	Ratio	Predicted
$2 \times 2 \times 2$	12 288	410×401	$1.02 \times 10^{+1}$	3.37×10^{-04}	-	-
$3 \times 3 \times 3$	41 472	464×475	$3.43 \times 10^{+1}$	4.81×10^{-04}	3.4	6.2
$4 \times 4 \times 4$	98 304	483×532	$7.92 \times 10^{+1}$	1.57×10^{-04}	2.3	3.7
$6 \times 6 \times 6$	331 776	504×707	$2.96 \times 10^{+2}$	7.03×10^{-04}	3.7	6.2
$8 \times 8 \times 8$	786 432	513×1014	$6.70 \times 10^{+2}$	4.70×10^{-04}	2.3	3.7
$10 \times 10 \times 10$	1 536 000	518×1502	$2.46 \times 10^{+3}$	3.53×10^{-04}	3.7	2.7

TABLE 2. Results for the experiments of Section 6.1 in the case $\epsilon = 1.0 \times 10^{-3}$.

The axis lengths of the ellipsoids were determined randomly — each axis length was taken to be $0.5 + 0.5\eta$ where η is a computer-generated pseudorandom number between 0 and 1. A $4 \times 4 \times 4$ domain of this type is depicted in Figure 6. For each experiment, the wavenumber κ was chosen to be

$$\kappa = \frac{8\pi}{2.5 \times n}$$

so that each domain considered was bounded by a box whose sides were 4 wavelengths in size.

Each experiment consisted of forming a scattering matrix for the integral equation (2.3) and using it to solve an instance of the boundary value problem (DBVP). The boundary data for each problem was taken to be the restriction of the function

$$g(x) = \sum_{j=1}^N \frac{\exp(i\kappa|x - x_j|)}{|x - x_j|},$$

where N is the total number of ellipsoids and the x_1, \dots, x_N are the centers of the ellipsoids, to the boundary of the ellipsoids. Let r and x_0 denote the radius and center of the minimum bounding sphere containing the boundary of the domain, respectively. Then the surface Γ_{in} which specified incoming potentials was the union of the spheres of radius 0.1 centered at the points x_1, \dots, x_N and the sphere of radius $2r$ centered at x_0 while the surface Γ_{out} specifying outgoing potentials was taken to be the sphere of radius $2r$ centered at x_0 . The function $g(x)$ is, of course, the unique solution of the boundary value problem and we measured the accuracy of each obtained solution σ of the the integral equation by computing the relative error

$$E = \left(\iint_{\Gamma} |g(x)|^2 dx \right)^{-1/2} \left(\iint_{\Gamma} |g(x) - (S_{\kappa} + i\kappa D_{\kappa})\sigma(x)|^2 dx \right)^{1/2}, \quad (6.1)$$

where Γ is the sphere of radius 2 centered at the point $x_0 - (0, 0, -2r - 10)$. The ellipsoids were parameterized by projection onto the cube $[-1, 1]^3$ and the faces of the cube were triangulated in order to produce a discretization. Three sets of experiments were performed, each at a different level of desired precision and the ellipsoids were discretized as follows in each case:

- For precision $\epsilon = 1.0 \times 10^{-3}$, each cube face comprising the parameterization domain was divided into 8 triangles and a 32 point 6th order quadrature was applied to each triangle. This yields a total of 1536 discretization nodes per ellipsoid.
- For precision $\epsilon = 1.0 \times 10^{-6}$, each cube face was subdivided into 32 triangles and a 32 point 6th order quadrature was applied to each triangle. This resulted in 6144 discretization nodes per ellipsoid.
- For precision $\epsilon = 1.0 \times 10^{-9}$, each cube face was subdivided into 8 triangles and a 112 point 12th order quadrature was applied to each triangle. This results in 5376 discretization nodes per ellipsoid.

Tables 2, 3, and 4 present the results of the experiments with $\epsilon = 1.0 \times 10^{-3}$, 1.0×10^{-6} , 1.0×10^{-9} , respectively. In each of those tables, N refers to the number of discretization nodes on the surface; $N_{\text{out}} \times N_{\text{in}}$ denotes the dimension of the resulting scattering matrix; T is the time in seconds required to construct the scattering matrix; E is the error (6.1); Ratio refers to the ratio of running times; and Predicted is the ratio of running times one would expect for an algorithm whose cost scales as $N^{1.5}$.

Grid dimensions	N	$N_{\text{out}} \times N_{\text{in}}$	T	E	Ratio	Predicted
$2 \times 2 \times 2$	49 152	601×584	$1.61 \times 10^{+2}$	1.22×10^{-07}	-	-
$3 \times 3 \times 3$	165 888	676×677	$6.87 \times 10^{+2}$	4.92×10^{-07}	4.3	6.2
$4 \times 4 \times 4$	393 216	703×747	$1.68 \times 10^{+3}$	5.31×10^{-07}	2.4	3.6
$6 \times 6 \times 6$	1 327 104	728×925	$6.66 \times 10^{+3}$	4.60×10^{-06}	4.0	6.2
$8 \times 8 \times 8$	3 145 728	$742 \times 1 237$	$1.59 \times 10^{+4}$	2.30×10^{-07}	2.4	3.6

TABLE 3. Results for the numerical experiments of Section 6.1 in the case $\epsilon = 1.0 \times 10^{-6}$.

Grid dimensions	N	$N_{\text{out}} \times N_{\text{in}}$	T	E	Ratio	Predicted
$2 \times 2 \times 2$	43 008	997×946	$5.14 \times 10^{+2}$	8.98×10^{-09}	-	-
$3 \times 3 \times 3$	145 152	$1 363 \times 1 301$	$2.32 \times 10^{+3}$	4.97×10^{-10}	4.5	6.2
$4 \times 4 \times 4$	344 064	$1 377 \times 1 353$	$5.56 \times 10^{+3}$	1.10×10^{-10}	2.4	3.6

TABLE 4. Results for the numerical experiments of Section 6.1 in the case $\epsilon = 1.0 \times 10^{-9}$.

Remark 6.1. *The boundary value problems solved in the experiments of this section present considerable difficulties for iterative solvers. We attempted to compare the running time of the fast solver of this paper with the running time of an iterative solver. However, it took more than 1000 iterations for GMRES residuals to fall below 1.0×10^{-3} when we attempted to solve the boundary value problem (DBVP) given on the $2 \times 2 \times 2$ grid of ellipsoids with 1536 discretization nodes per ellipsoid via the combined field integral equation. Even with the aid of an extremely efficient multipole code, replicating the experiments of this section with an iterative solver appears to be prohibitively expensive.*

6.2. Sound-hard scattering from a deformed torus. In the experiments described here we considered the problem (NBVP) with Γ the surface parameterized via

$$r(s, t) = \begin{pmatrix} (2 + 0.50 \cos(s)) \sin(t) \\ (2 + 0.50 \cos(s)) \cos(t) \\ 0.50 \sin(s) (1 + 0.15 \cos(36t)) \end{pmatrix}, \quad 0 \leq s, t < 2\pi. \quad (6.2)$$

Three views of this surface, which is enclosed in the box with lower left corner $(-2.5, -2.5, -0.65)$ and upper right corner $(2.5, 2.5, 0.65)$, are shown in Figure 7.



FIGURE 7. Three views of the deformed torus of Section 6.2.

Each experiment consisted of constructing a scattering matrix for the integral equation formulation (2.4) of the problem (NBVP) and using that scattering matrix to solve an instance of the boundary value problem. The wavenumber κ was set at $4\pi/5$, which makes the domain approximately $2 \times 2 \times 0.7$ wavelengths in size, and the boundary data was taken to be the normal derivative of the function

$$g(x) = \frac{1}{4\pi} \frac{\exp(i\kappa|x - x_0|)}{|x - x_0|},$$

where x_0 is the point $(2, 0, 0)$. Since x_0 lies inside of the domain Ω , the function $g(x)$ is, of course, the unique solution of (NBVP). The surface Γ_{in} which specified incoming potentials was taken to be the union of the sphere of radius 0.1 centered at the point $(2, 0, 0)$ and the sphere of radius 6 centered at $(0, 0, 0)$

while the surface Γ_{out} specifying outgoing potentials was taken to be the sphere of radius 6 centered at $(0, 0, 0)$. The accuracy of each obtained solution σ was measured by computing the relative error

$$E = \left(\iint_S |g(x)|^2 dx \right)^{-1/2} \left(\iint_S |g(x) - S_\kappa \sigma(x)|^2 dx \right)^{1/2}, \quad (6.3)$$

where S is the sphere of radius 1 centered at the point $(-20, 0, 0)$. The order of the discretization quadrature was fixed at 8 and the requested precision for the scattering matrix was set to $\epsilon = 1.0 \times 10^{-7}$ in each experiment. The number of triangles used to decompose the parameterization domain was varied.

N_{tris}	N	T	E	$N_{\text{out}} \times N_{\text{in}}$	Ratio	Predicted
32	1 664	$7.16 \times 10^{+00}$	3.51×10^{-02}	395×447	-	-
128	6 656	$6.29 \times 10^{+01}$	4.41×10^{-03}	407×460	8.79	8
512	26 624	$2.81 \times 10^{+02}$	4.08×10^{-05}	400×456	4.47	8
2 048	106 496	$2.60 \times 10^{+03}$	7.80×10^{-07}	408×463	9.25	8
8 192	425 984	$1.47 \times 10^{+04}$	3.25×10^{-08}	409×466	5.66	8

TABLE 5. The results of the experiments of Section 6.2.

Table 5 shows the results. There, N_{tris} is the number of triangles into which the parameterization domain was partitioned; N is the total number of discretization nodes on the boundary Γ ; T is the total running time of the algorithm in seconds; E is the relative error (6.3); $N_{\text{out}} \times N_{\text{in}}$ gives the dimensions of the scattering matrix; Ratio is the quotient of the running time at the current level of discretization to the running time at the preceding level of discretization (where it is defined); and Predicted gives the expected ratio of the running times (assuming $O(N^{1.5})$ asymptotic cost).

6.3. A heart-shaped domain. In these experiments, we solved the problem (NBVP) on the domain Ω shown in Figure 8. The boundary Γ is described by the parameterization

$$r(s, t) = \begin{pmatrix} \cos(s) \cos(t) + \cos^2(s) \sin^2(t) \\ \cos(s) \sin(t) \\ \sin(s) \end{pmatrix}, \quad 0 \leq s, t < 2\pi,$$

although for the purposes of producing discretizations of the surface Γ a parameterization obtained via projection onto a cube was used rather than this function $r(s, t)$. The domain Ω is contained in the box with lower left corner $(-1, -1, -1)$ and upper right corner $(1.25, 1, 1)$.

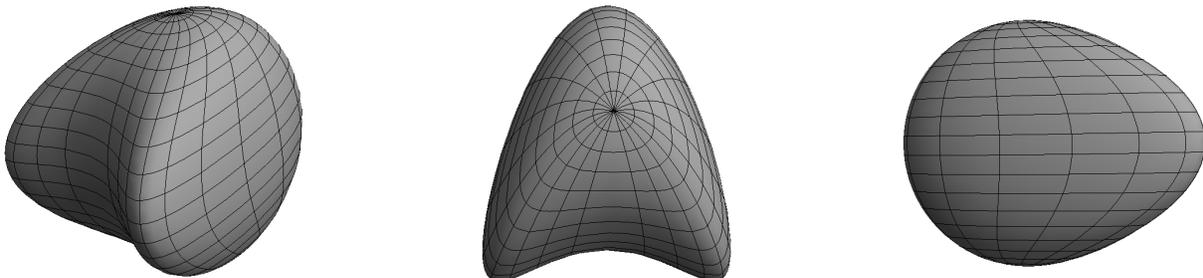


FIGURE 8. Three views of the heart shaped surface of Section 6.3.

These experiments consisted of solving the problem (NBVP) on the domain Ω at a variety of wavenumbers. In each experiment, a scattering matrix for the domain was formed and an instance of the boundary value problem was solved using this scattering matrix. The boundary data was taken to be the derivative of the restriction of the function

$$g(x) = \frac{1}{4\pi} \frac{\exp(i\kappa|x|)}{|x|}$$

to Γ with respect to the outward pointing unit normal vector on Γ . The error in the obtained charge distribution was measured by computing

$$E = \left(\iint_S |g(x)|^2 dx \right)^{-1/2} \left(\iint_S |g(x) - S_\kappa \sigma(x)|^2 dx \right)^{1/2}, \quad (6.4)$$

where S is the ball of radius 10 centered at the point $(-30, -30, -30)$ and σ is the approximate charge distribution obtained using the scattering matrix. Discretizations of Γ were obtained by triangulating the faces of the cube which constitutes the parameterization domain. An 8th order discretization quadrature with 45 points was applied to each triangle. The number of triangles was increased with wavenumber. In all cases, the desired precision for the scattering matrix was taken to be $\epsilon = 1.0 \times 10^{-3}$. Table 6 reports the number of discretization nodes N on the surface; the relative error E ; the number of triangles N_{tris} into which the parameterization domain was divided; the total running time T ; the dimensions $N_{\text{out}} \times N_{\text{in}}$ of the scattering matrix; the wavenumber κ for the problem; and the approximate diameter of the domain in wavelengths.

κ	Diameter	N_{tris}	N	E	T	$N_{\text{out}} \times N_{\text{in}}$
2π	3.46λ	48	2 496	4.20×10^{-04}	$6.84 \times 10^{+00}$	251×229
4π	6.92λ	192	9 984	1.33×10^{-04}	$6.93 \times 10^{+01}$	571×546
6π	10.39λ	768	39 936	1.88×10^{-04}	$3.74 \times 10^{+02}$	996×972
8π	13.87λ	768	39 936	3.61×10^{-04}	$4.99 \times 10^{+02}$	1491×1471
10π	17.32λ	3 072	159 744	4.96×10^{-04}	$2.83 \times 10^{+03}$	2157×2126
12π	20.79λ	3 072	159 744	1.89×10^{-04}	$3.78 \times 10^{+03}$	2851×2829
14π	24.25λ	12 288	638 976	6.75×10^{-04}	$1.81 \times 10^{+04}$	3615×3589

TABLE 6. The results of the experiments of Section 6.3.

6.4. Scattering from a deformed cube. In this next set of experiments, we considered the problem (NBVP) on a deformed cube Ω with 12 curved edges and 8 corner points. The domain Ω is pictured in Figure 9. The top face of the boundary Γ was parameterized over the region $-1 \leq s, t \leq 1$ via

$$r(s, t) = (2 + s^2 + t^2) \begin{pmatrix} s \\ t \\ 1 \end{pmatrix}$$

and the other faces were handled in a like manner. The domain Ω is contained in the box with lower left corner $(-4, -4, -4)$ and upper right corner $(4, 4, 4)$.

In these experiments, the integral equation (2.4) defined on Γ was discretized by applying a 12th order discretization quadrature with 112 nodes to increasing dense triangulations of the parameterization domain. We arranged for a denser distribution of discretization nodes near the edges and corners of the deformed cube. Figure 10 shows one of the unevenly refined meshes used in these experiments.

A scattering matrix was constructed for each discretization and an instance of the boundary value problem was solved using this scattering matrix. In each experiment, the wavenumber was taken to be $\kappa = \pi/2$, making the domain roughly 3.46 wavelengths in diameter, the desired precision for the scattering matrix was set to $\epsilon = 1.0 \times 10^{-10}$ and the boundary data was taken to be the outward normal derivative on Γ of the function

$$g(x) = \frac{1}{4\pi} \iint_{\Gamma_{\text{in}}} \frac{\exp(i\kappa|x-y|)}{|x-y|} ds(y),$$

where Γ_{in} denotes the sphere of radius 0.1 centered at the point $(0.1, 0, 0)$. The error in the obtained charge distribution σ was measured via the quantity

$$E = \left(\iint_{\Gamma_{\text{out}}} |g(x)|^2 dx \right)^{-1/2} \left(\iint_{\Gamma_{\text{out}}} |g(x) - S_\kappa \sigma(x)|^2 dx \right)^{1/2}, \quad (6.5)$$

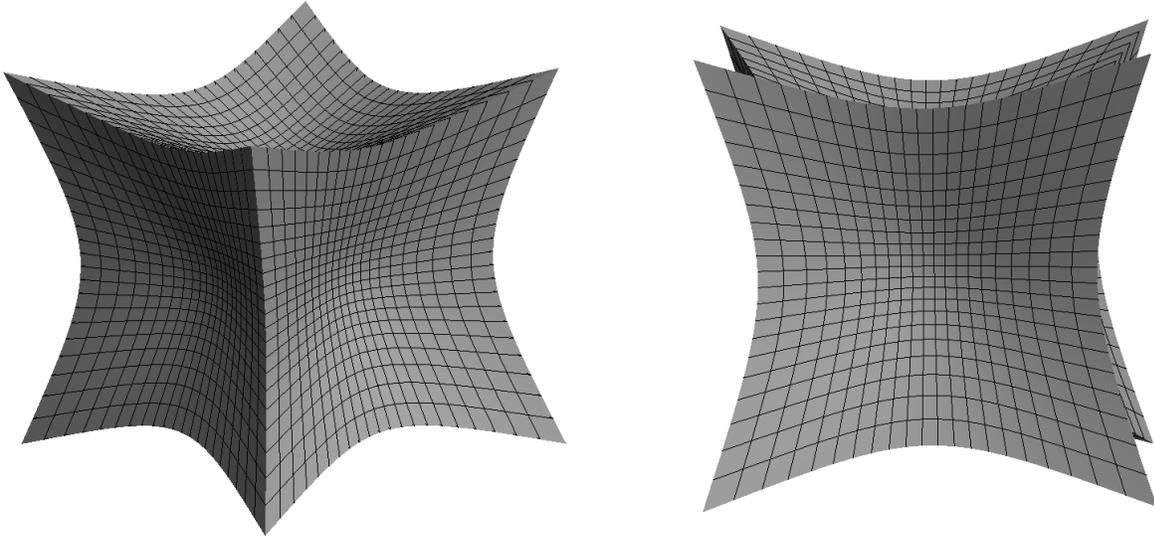


FIGURE 9. The deformed cube domain considered in Section 6.4. The domain has 12 curved edges and 8 corner points.

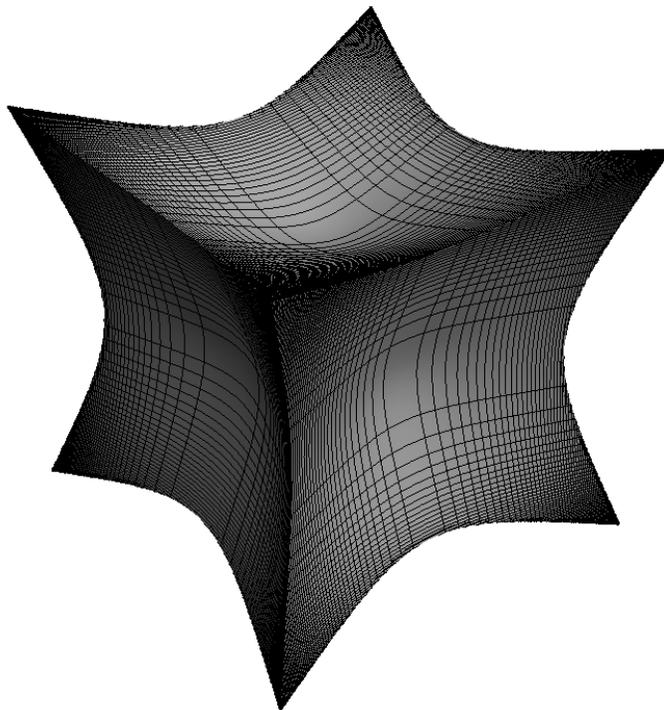


FIGURE 10. Illustration of the unevenly refined grid used to generate discretization nodes on the deformed cube surface considered in Section 6.4.

where σ is the obtained charge distribution and Γ_{out} denotes the sphere of radius 1 centered at the point $(-10, -10, 0)$. Table 7 reports the error (6.5) along with N_{tris} , N , T and $N_{\text{out}} \times N_{\text{in}}$ (whose meanings are as in the preceding sections).

N_{tris}	N	E	T	$N_{\text{out}} \times N_{\text{in}}$
192	21 504	2.60×10^{-08}	$6.11 \times 10^{+02}$	617×712
432	48 384	2.13×10^{-09}	$1.65 \times 10^{+03}$	620×694
768	86 016	3.13×10^{-10}	$3.58 \times 10^{+03}$	612×685

TABLE 7. The results of the experiments performed on the deformed cube of Section 6.4.

6.5. A Convergence Study involving a Trefoil Domain. The experiments described in the preceding sections all involved problems with artificial boundary data associated with known solutions. In this section, we describe an experiment for which we did not have access to an exact solution, and instead estimate the error by comparing against a very finely resolved computed solution.

We considered the boundary value problem (DBVP) given on the trefoil domain Ω shown in Figure 11. The parameterization of the boundary of Ω is given by

$$r(s, t) = \gamma(s) + \cos(s)\mathbf{n}(s) + \sin(-t)\mathbf{b}(s), \quad 0 \leq s, t < 2\pi,$$

where

$$\gamma(s) = \begin{pmatrix} \sin(3s) \\ \sin(s) + 2\sin(2s) \\ \cos(s) - 2\cos(2s) \end{pmatrix},$$

$\mathbf{n}(s)$ is the unit normal vector to the curve $\gamma(s)$ at s and $\mathbf{b}(s)$ is the unit binormal vector at s . The domain is contained in a box with bottom left corner at $[-1.5, -3, -3.5]$ and top right corner $[1.5, 3, 3.5]$.

Discretizations of the surface Γ were formed by partitioning the parameter space into N_{tris} triangles and utilizing an 8th order (52 point) quadrature on each triangle so that $N = 52N_{\text{tris}}$. A scattering matrix for each discretization was then formed; Γ_{in} consisted of the point $(1, 0, 0)$ and Γ_{out} was taken to be the surface of a sphere of radius 1 centered at $(30, 30, 30)$. The boundary data was generated by a point charge on Γ_{in} . As the solution to this problem is unknown, we report the L^2 convergence of the solution on the surface of Γ_{out} . Specifically, we report

$$E_N = \left(\iint_{\Gamma_{\text{out}}} |(S_\kappa + i\kappa D_\kappa)\sigma_{4N}(x) - (S_\kappa + i\kappa D_\kappa)\sigma_N(x)|^2 dx \right)^{1/2}, \quad (6.6)$$

where σ_N is the computed boundary charge distribution with N discretization points.

We considered two choices of wavenumber and tolerance: (i) the domain was approximately $1 \times 3 \times 3$ wavelengths in size and the tolerance is set to $\epsilon = 10^{-9}$; (ii) the domain was approximately $2 \times 6 \times 6$ wavelengths in size and the tolerance is set to $\epsilon = 10^{-6}$. Tables 8 and 9 report on the convergence of the solution as the number of discretization points N is increased.

N_{tris}	N	E_N	T	$N_{\text{out}} \times N_{\text{in}}$
16	832	6.73×10^{-04}	$1.17 \times 10^{+00}$	754×737
64	3 328	2.33×10^{-06}	$3.78 \times 10^{+01}$	939×910
256	13 312	2.59×10^{-08}	$3.61 \times 10^{+02}$	945×917
1 024	53 248	2.47×10^{-11}	$2.55 \times 10^{+03}$	948×918
4 096	212 992	-	$2.83 \times 10^{+04}$	949×921

TABLE 8. The results for the experiments performed on the trefoil of Section 6.5 when the domain is is approximately $1 \times 3 \times 3$ wavelengths in size and $\epsilon = 10^{-9}$.

7. ACKNOWLEDGEMENTS

J. Bremer was supported by a fellowship from the Alfred P. Sloan Foundation and by the Office of Naval Research under contract N00014-12-1-0117. P.G. Martinsson was supported by the National Science Foundation under contracts DMS0748488 and DMS0941476.

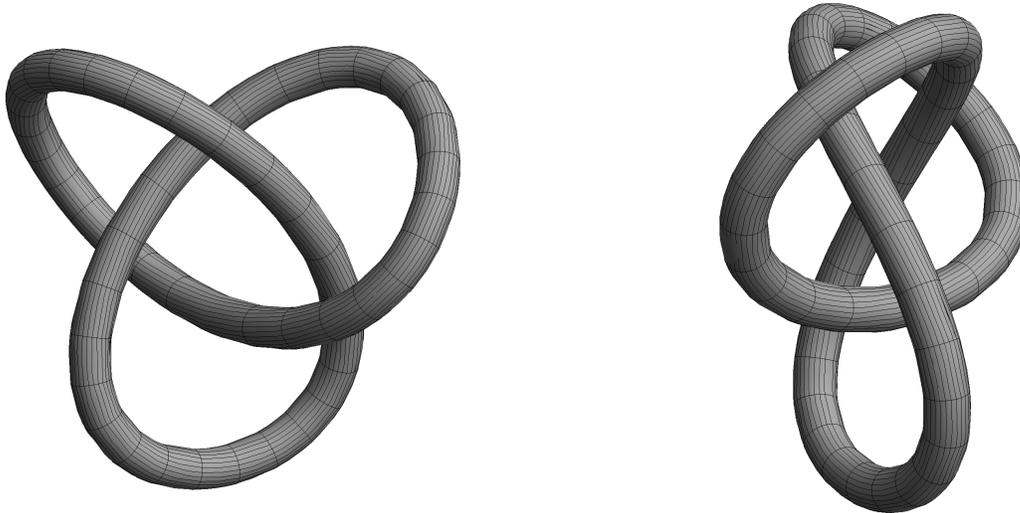


FIGURE 11. The trefoil domain Ω considered in Section 6.5.

N_{tris}	N	E_N	T	$N_{\text{out}} \times N_{\text{in}}$
16	832	2.87×10^{-02}	$1.94 \times 10^{+00}$	832×832
64	3328	1.08×10^{-03}	$7.59 \times 10^{+01}$	1771×1745
256	13312	5.57×10^{-06}	$4.12 \times 10^{+02}$	1820×1782
1024	53248	3.04×10^{-08}	$2.16 \times 10^{+03}$	1822×1788
4096	212992	-	$1.27 \times 10^{+04}$	1828×1794

TABLE 9. The results for the experiments performed on the trefoil of Section 6.5 when the domain is is approximately $2 \times 6 \times 6$ wavelengths in size and $\epsilon = 10^{-6}$.

APPENDIX A. A ROBUST INTEGRAL FORMULATION FOR SOLVING (NBVP)

Consider the task of solving the Neumann boundary value problem (NBVP) using a boundary integral equation formulation. In order to avoid the problem of spurious resonances, it would be natural to use a combined layer representation of the form

$$u(x) = D_\kappa \sigma(x) + i|\kappa|S_\kappa \sigma(x). \tag{A.1}$$

However, the BIE obtained by simply combining (A.1) with the Neumann boundary condition involves the *hypersingular* operator

$$N_\kappa \sigma(x) = \int_{\partial\Omega} (\eta_x \cdot \nabla_x (\eta_y \cdot \nabla_y G_\kappa(x, y))) \sigma(y) ds(y).$$

Two difficulties are encountered when handling integral equations involving hypersingular operators:

1. Viewed as an operator from $L^2(\Gamma)$ to $L^2(\Gamma)$, N_κ is not bounded. In consequence, if standard discretization methods are applied to integral equations involving this operator, ill-conditioned linear systems will result. This complication can be overcome by viewing N_κ as an operator between two different Sobolev spaces; however, more complicated analytic machinery would then be required.
2. The singularity in the kernel of the hypersingular operator N_κ is more severe than the singularities of the standard operators of acoustic scattering. As discussed in [8], the approach used here to evaluate singular integrals can be generalized to this case. However, considerable work is required and the resulting quadratures are somewhat less efficient than those used here.

A technique known as “regularization” offers an alternative to handling the operator N_κ [11]. For example, the representation (2.2) can be replaced with

$$u(x) = S_\kappa \sigma(x) + D_\kappa S_\kappa \sigma(x)$$

which leads to the integral equation

$$-\frac{1}{2}\sigma(x) + [N_\kappa S_\kappa]\sigma(x) + D_\kappa^* \sigma(x) = g(x). \quad (\text{A.2})$$

Using Calderón identities (which are discussed in, for instance, [28]), we can rewrite (A.2) as

$$-\frac{3}{4}\sigma(x) + D_\kappa^* \sigma(x) + (D_\kappa^*)^2 \sigma(x) = g(x), \quad (\text{A.3})$$

which does not involve hypersingular operators. While the equation (A.3) is suitable for use by iterative solvers, the composition of the operator D_κ^* makes it cumbersome for direct solvers.

To make the formulation (A.3) more suitable for direct solvers, we introduce an auxiliary variable

$$\phi(x) = D_\kappa^* \sigma(x),$$

which results to the system of integral equations

$$\begin{aligned} -\frac{3}{4}\sigma(x) + D_\kappa^* \sigma(x) + D_\kappa^* \phi(x) &= g(x) \\ -\phi(x) + D_\kappa^* \sigma(x) &= 0. \end{aligned} \quad (\text{A.4})$$

The system (A.4) is now ideal for direct solvers as it does not involve hypersingular operators and no compositions of operators appear.

APPENDIX B. DERIVATION OF SCATTERING MATRICES AND THE MERGE FORMULA FROM A FIELD FORMULATION

The derivation of the formula for merging two scattering matrices given in Section 5.4 is quite formal. In this section, we provide a different derivation that is hopefully more intuitive and gives a better idea of the physical meaning of a scattering matrix S_τ for a subpatch $\Gamma_\tau \subseteq \Gamma$. This derivation was omitted from the main text since it requires the introduction of several auxiliary fields.

Throughout this section of the appendix, we rely on the hierarchical tree partitioning of the domain Γ described in Section 5.2, and the associated HBS representation of a matrix given in Section 5.3.

B.1. Problem formulation. Our objective is to evaluate the map (4.4), which is the discretization of the scattering problem (4.1). We restate it for easy reference: Given a vector \mathbf{s}_{in} representing a source distribution on the “radiation source,” we seek to compute the field \mathbf{u}_{out} on the “antenna” given by

$$\begin{array}{ccccccc} \mathbf{u}_{\text{out}} & = & \mathbf{A}_{\text{out}} & \mathbf{A}^{-1} & \mathbf{A}_{\text{in}} & \mathbf{s}_{\text{in}}, & \\ N_{\text{out}} \times 1 & & N_{\text{out}} \times N & N \times N & N \times N_{\text{in}} & N_{\text{in}} \times 1 & \end{array} \quad (\text{B.1})$$

We define the *external incoming field* on the scattering surface Γ via

$$\mathbf{u} = \mathbf{A}_{\text{in}} \mathbf{s}_{\text{in}}.$$

The core task is then to solve the equation

$$\mathbf{A}\boldsymbol{\sigma} = \mathbf{u} \quad (\text{B.2})$$

to compute the field of induced charges $\boldsymbol{\sigma}$ on Γ . As described in Section 4.2, we do not need to solve (B.2) in full generality, but only under two simplifying constraints:

- (1) The external incoming field \mathbf{u} is restricted to a low-dimensional space spanned by the provided basis matrix $\hat{\mathbf{U}}_1$. In other words, there exists a short vector $\tilde{\mathbf{u}}_1$ such that $\mathbf{u} = \hat{\mathbf{U}}_1 \tilde{\mathbf{u}}_1$.
- (2) We only need the projection of $\boldsymbol{\sigma}$ onto a low-dimensional space spanned by the provided basis matrix $\hat{\mathbf{V}}_1$. In other words, we only need the short vector $\tilde{\boldsymbol{\sigma}}_1 = \hat{\mathbf{V}}_1^* \boldsymbol{\sigma}$.

For convenience in the hierarchical formulation of the scattering problem, we define for a subpatch Γ_τ of the scattering surface an *internal incoming field* via

$$\mathbf{v}_\tau = \mathbf{A}(I_\tau, I_\tau^c) \boldsymbol{\sigma}(I_\tau^c). \quad (\text{B.3})$$

Then \mathbf{v}_τ is the field on Γ_τ induced by charges on $\Gamma \setminus \Gamma_\tau$. For a node τ , the restriction of (B.2) to I_τ ,

$$\mathbf{A}(I_\tau, I_\tau) \boldsymbol{\sigma}(I_\tau) + \mathbf{A}(I_\tau, I_\tau^c) \boldsymbol{\sigma}(I_\tau^c) = \mathbf{u}(I_\tau) \quad (\text{B.4})$$

can then be written

$$\mathbf{A}(I_\tau, I_\tau) \boldsymbol{\sigma}(I_\tau) = \mathbf{u}(I_\tau) - \mathbf{v}_\tau. \quad (\text{B.5})$$

In words, equation (B.5) states that the induced charges $\boldsymbol{\sigma}(I_\tau)$ on Γ_τ must create a local field that precisely match the *total incoming field* $\mathbf{u}(I_\tau) - \mathbf{v}_\tau$, which is the difference between the externally imposed incoming field $\mathbf{u}(I_\tau)$ and the internal incoming field \mathbf{v}_τ which is caused by induced charges on the rest of Γ .

B.2. Short representations and the definition of the scattering matrix. By definition, the columns of the long basis matrix $\hat{\mathbf{U}}_\tau$ span both the restriction $\mathbf{u}(I_\tau)$ of the external incoming field to the patch Γ_τ , and the columns of the off-diagonal block $\mathbf{A}(I_\tau, I_\tau^c)$. As a consequence, there exist short vectors $\tilde{\mathbf{u}}_\tau$ and $\tilde{\mathbf{v}}_\tau$ such that

$$\mathbf{v}_\tau = \hat{\mathbf{U}}_\tau \tilde{\mathbf{v}}_\tau, \quad \text{and} \quad \mathbf{u}(I_\tau) = \hat{\mathbf{U}}_\tau \tilde{\mathbf{u}}_\tau. \quad (\text{B.6})$$

We say that $\tilde{\mathbf{v}}_\tau$ and $\tilde{\mathbf{u}}_\tau$ are the *compressed representations* of \mathbf{v}_τ and $\mathbf{u}(I_\tau)$, respectively. Using these representations, (B.5) can be written

$$\mathbf{A}(I_\tau, I_\tau) \boldsymbol{\sigma}(I_\tau) = \hat{\mathbf{U}}_\tau (\tilde{\mathbf{u}}_\tau - \tilde{\mathbf{v}}_\tau). \quad (\text{B.7})$$

We also define a compressed representation of the induced charges $\boldsymbol{\sigma}(I_\tau)$ on Γ_τ via

$$\tilde{\boldsymbol{\sigma}}_\tau = \hat{\mathbf{V}}_\tau^* \boldsymbol{\sigma}(I_\tau). \quad (\text{B.8})$$

Multiplying (B.7) by $\hat{\mathbf{V}}_\tau^* \mathbf{A}(I_\tau, I_\tau)^{-1}$, we now get the scattering equation

$$\tilde{\boldsymbol{\sigma}}_\tau = \mathbf{S}_\tau (\tilde{\mathbf{u}}_\tau - \tilde{\mathbf{v}}_\tau), \quad (\text{B.9})$$

where, recall from (4.6),

$$\mathbf{S}_\tau = \hat{\mathbf{V}}_\tau^* \mathbf{A}(I_\tau, I_\tau)^{-1} \hat{\mathbf{U}}_\tau. \quad (\text{B.10})$$

Equation (B.9) states that *the scattering matrix* \mathbf{S}_τ *maps the short representation* $\tilde{\mathbf{u}}_\tau - \tilde{\mathbf{v}}_\tau$ *of the total incoming field on* Γ_τ *to the short representation* $\tilde{\boldsymbol{\sigma}}_\tau$ *of the induced charges.*

B.3. Hierarchical construction of scattering matrices. For leaf nodes τ , we compute \mathbf{S}_τ directly from the definition (B.10). For parent nodes, we compute the scattering matrices via a hierarchical merging procedure; it turns out that if we know the scattering matrices \mathbf{S}_α and \mathbf{S}_β of the children $\{\alpha, \beta\}$ of a node τ , then \mathbf{S}_τ can be built inexpensively as follows: Observe that (B.7) can be written

$$\begin{bmatrix} \mathbf{A}(I_\alpha, I_\alpha) & \mathbf{A}(I_\alpha, I_\beta) \\ \mathbf{A}(I_\beta, I_\alpha) & \mathbf{A}(I_\beta, I_\beta) \end{bmatrix} \begin{bmatrix} \boldsymbol{\sigma}(I_\alpha) \\ \boldsymbol{\sigma}(I_\beta) \end{bmatrix} = \hat{\mathbf{U}}_\tau (\tilde{\mathbf{u}}_\tau - \tilde{\mathbf{v}}_\tau). \quad (\text{B.11})$$

Inserting into (B.11) the factorizations (5.2) of $\mathbf{A}(I_\alpha, I_\beta)$ and $\mathbf{A}(I_\beta, I_\alpha)$, and the hierarchical representation (5.3) of $\hat{\mathbf{U}}_\tau$, we get

$$\begin{bmatrix} \mathbf{A}(I_\alpha, I_\alpha) & \hat{\mathbf{U}}_\alpha \tilde{\mathbf{A}}_{\alpha, \beta} \hat{\mathbf{V}}_\beta^* \\ \hat{\mathbf{U}}_\beta \tilde{\mathbf{A}}_{\beta, \alpha} \hat{\mathbf{V}}_\alpha^* & \mathbf{A}(I_\beta, I_\beta) \end{bmatrix} \begin{bmatrix} \boldsymbol{\sigma}(I_\alpha) \\ \boldsymbol{\sigma}(I_\beta) \end{bmatrix} = \begin{bmatrix} \hat{\mathbf{U}}_\alpha & \mathbf{0} \\ \mathbf{0} & \hat{\mathbf{U}}_\beta \end{bmatrix} \mathbf{U}_\tau (\tilde{\mathbf{u}}_\tau - \tilde{\mathbf{v}}_\tau). \quad (\text{B.12})$$

Left-multiply (B.12) by $\begin{bmatrix} \hat{\mathbf{V}}_\alpha^* \mathbf{A}(I_\alpha, I_\alpha)^{-1} & \mathbf{0} \\ \mathbf{0} & \hat{\mathbf{V}}_\beta^* \mathbf{A}(I_\beta, I_\beta)^{-1} \end{bmatrix}$ to obtain

$$\begin{bmatrix} \mathbf{I} & \mathbf{S}_\alpha \tilde{\mathbf{A}}_{\alpha, \beta} \\ \mathbf{S}_\beta \tilde{\mathbf{A}}_{\beta, \alpha} & \mathbf{I} \end{bmatrix} \begin{bmatrix} \tilde{\boldsymbol{\sigma}}_\alpha \\ \tilde{\boldsymbol{\sigma}}_\beta \end{bmatrix} = \begin{bmatrix} \mathbf{S}_\alpha & \mathbf{0} \\ \mathbf{0} & \mathbf{S}_\beta \end{bmatrix} \mathbf{U}_\tau (\tilde{\mathbf{u}}_\tau - \tilde{\mathbf{v}}_\tau). \quad (\text{B.13})$$

Solving (B.13) for $\{\tilde{\sigma}_\alpha, \tilde{\sigma}_\beta\}$ and left multiplying the result by \mathbf{V}_τ^* , we find

$$\underbrace{\mathbf{V}_\tau^* \begin{bmatrix} \tilde{\sigma}_\alpha \\ \tilde{\sigma}_\beta \end{bmatrix}}_{=\tilde{\sigma}_\tau} = \mathbf{V}_\tau^* \underbrace{\begin{bmatrix} \mathbf{I} & \mathbf{S}_\alpha \tilde{\mathbf{A}}_{\alpha,\beta} \\ \mathbf{S}_\beta \tilde{\mathbf{A}}_{\beta,\alpha} & \mathbf{I} \end{bmatrix}^{-1} \begin{bmatrix} \mathbf{S}_\alpha & \mathbf{0} \\ \mathbf{0} & \mathbf{S}_\beta \end{bmatrix}}_{=\mathbf{S}_\tau} \mathbf{U}_\tau(\tilde{\mathbf{u}}_\tau - \tilde{\mathbf{v}}_\tau). \quad (\text{B.14})$$

Comparing equation (B.14) to (B.9), we rediscover the formula (5.8) for the merge of two scattering matrices given in Lemma 5.1.

REFERENCES

- [1] BARNETT, A., AND GREENGARD, L. A new integral representation for quasi-periodic fields and its application to two-dimensional band structure calculations. *Journal of Computational Physics* 229 (2010), 6898–6914.
- [2] BÖRM, S. *Efficient numerical methods for non-local operators*, vol. 14 of *EMS Tracts in Mathematics*. European Mathematical Society (EMS), Zürich, 2010. \mathcal{H}^2 -matrix compression, algorithms and analysis.
- [3] BÖRM, S., AND HACKBUSCH, W. Approximation of boundary element operators by adaptive \mathcal{H}^2 -matrices. In *Foundations of computational mathematics: Minneapolis, 2002*, vol. 312 of *London Math. Soc. Lecture Note Ser.* Cambridge Univ. Press, Cambridge, 2004, pp. 58–75.
- [4] BRAKHAGE, H., AND WERNER, P. über das dirichletsche ausssenraum problem für die helmholtzsche schwingungsgleichung. *Arch. Math* 16 (1965), 325–329.
- [5] BREMER, J. A fast direct solver for the integral equations of scattering theory on planar curves with corners. *Journal of Computational Physics* 231 (2012), 45–64.
- [6] BREMER, J. On the Nyström discretization of integral operators on planar domains with corners. *Applied and Computational Harmonic Analysis* 32 (2012), 45–64.
- [7] BREMER, J., AND GIMBUTAS, Z. A Nyström method for weakly singular integral operators on surfaces. *Journal of Computational Physics* 231 (2012), 4885–4903.
- [8] BREMER, J., AND GIMBUTAS, Z. On the numerical evaluation of the singular integrals of scattering theory. *Journal of Computational Physics* 251 (2013), 327–343.
- [9] BREMER, J., GIMBUTAS, Z., AND ROKHLIN, V. A nonlinear optimization procedure for generalized Gaussian quadratures. *SIAM Journal of Scientific Computing* 32 (2010), 1761–1788.
- [10] BREMER, J., ROKHLIN, V., AND SAMMIS, I. Universal quadratures for boundary integral equations on two-dimensional domains with corners. *Journal of Computational Physics* 229 (2010), 8259–8280.
- [11] BRUNO, O., ELLING, T., AND TURC, C. Well-conditioned high-order algorithms for the solution of three-dimensional surface acoustic scattering problems with Neumann boundary conditions. *Journal of Numerical Methods in Engineering* 91 (2012), 1045–1072.
- [12] CHANDRASEKARAN, S., AND GU, M. Fast and stable algorithms for banded plus semiseparable systems of linear equations. *SIAM J. Matrix Anal. Appl.* 25, 2 (2003), 373–384 (electronic).
- [13] CHENG, H., GIMBUTAS, Z., MARTINSSON, P., AND ROKHLIN, V. On the compression of low rank matrices. *SIAM Journal of Scientific Computing* 26, 4 (2005), 1389–1404.
- [14] COLTON, D., AND KRESS, R. *Inverse Acoustic and Electromagnetic Scattering Theory*, 2nd ed. Springer-Verlag, New York, 1998.
- [15] GILLMAN, A., YOUNG, P., AND MARTINSSON, P.-G. A direct solver $O(N)$ complexity for integral equations on one-dimensional domains. *Frontiers of Mathematics in China* 7 (2012), 217–247. 10.1007/s11464-012-0188-3.
- [16] GRASEDYCK, L., AND HACKBUSCH, W. Construction and arithmetics of \mathcal{H} -matrices. *Computing* 70, 4 (2003), 295–334.
- [17] GREENGARD, L., GUEYFFIER, D., MARTINSSON, P.-G., AND ROKHLIN, V. Fast direct solvers for integral equations in complex three-dimensional domains. *Acta Numer.* 18 (2009), 243–275.
- [18] GU, M., AND EISENSTAT, S. C. Efficient algorithms for computing a strong rank-revealing QR factorization. *SIAM J. Sci. Comput.* 17, 4 (1996), 848–869.
- [19] HACKBUSCH, W. *Integral Equations: theory and numerical treatment*. Birkhäuser, Berlin, 1995.
- [20] HACKBUSCH, W., KHOROMSKIJ, B., AND SAUTER, S. On \mathcal{H}^2 -matrices. In *Lectures on Applied Mathematics*. Springer Berlin, 2002, pp. 9–29.
- [21] HALKO, N., MARTINSSON, P.-G., AND TROPP, J. A. Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions. *SIAM Review* 53, 2 (2011), 217–288.
- [22] HELSING, J. A fast and stable solver for singular integral equations on piecewise smooth curves. *SIAM Journal of Scientific Computing* 33 (2011), 153–174.
- [23] HELSING, J., AND OJALA, R. Corner singularities for elliptic problems: integral equations, graded meshes, quadrature, and compressed inverse preconditioning. *Journal of Computational Physics* 227 (2008), 8820–8840.
- [24] HO, K., AND GREENGARD, L. A fast direct solver for structured linear systems by recursive skeletonization. *SIAM Journal on Scientific Computing* 34, 5 (2012), 2507–2532.
- [25] KRESS, R. *Integral Equations*. Springer-Verlag, New York, 1999.

- [26] MARTINSSON, P., AND ROKHLIN, V. A fast direct solver for boundary integral equations in two dimensions. *J. Comp. Phys.* 205, 1 (2005), 1–23.
- [27] MARTINSSON, P., AND ROKHLIN, V. A fast direct solver for boundary integral equations in two dimensions. *Journal of Computational Physics* 205 (2006).
- [28] NÉDÉLEC, J.-C. *Acoustic and Electromagnetic Equations: Integral Representations for Harmonic Functions*. Springer, New York, 2012.
- [29] SHENG, Z., DEWILDE, P., AND CHANDRASEKARAN, S. Algorithms to solve hierarchically semi-separable systems. In *System theory, the Schur algorithm and multidimensional analysis*, vol. 176 of *Oper. Theory Adv. Appl.* Birkhäuser, Basel, 2007, pp. 255–294.
- [30] TASKINEN, M., AND YLA-OIJALA, P. Current and charge integral equation formulation. *IEEE Transactions on Antennas and Propagation* 54 (2006), 58–67.
- [31] XIA, J., CHANDRASEKARAN, S., GU, M., AND LI, X. Fast algorithms for hierarchically semiseparable matrices. *Numerical Linear Algebra with Applications* 17, 6 (2010), 953–976.
- [32] XIA, J., CHANDRASEKARAN, S., GU, M., AND LI, X. S. Superfast multifrontal method for large structured linear systems of equations. *SIAM J. Matrix Anal. Appl.* 31, 3 (2009), 1382–1411.
- [33] XIA, J., CHANDRASEKARAN, S., GU, M., AND LI, X. S. Fast algorithms for hierarchically semiseparable matrices. *Numerical Linear Algebra with Applications* 17, 6 (2010), 953–976.