

Fast multipole methods for the evaluation of layer potentials with locally-corrected quadratures



Leslie Greengard^{a,b}, Michael O'Neil^a, Manas Rachh^{b,*}, Felipe Vico^c

^a Courant Institute, NYU, USA

^b Center for Computational Mathematics, Flatiron Institute, USA

^c Instituto de Telecomunicaciones y Aplicaciones Multimedia (ITEAM) Universidad Politecnica de Valencia, Spain

ARTICLE INFO

Article history:

Received 21 August 2020

Received in revised form 22 March 2021

Accepted 11 April 2021

Available online 16 April 2021

Keywords:

Nyström method

Helmholtz equation

Quadrature

Fast multipole method

ABSTRACT

While fast multipole methods (FMMs) are in widespread use for the rapid evaluation of potential fields governed by the Laplace, Helmholtz, Maxwell or Stokes equations, their coupling to high-order quadratures for evaluating layer potentials is still an area of active research. In three dimensions, a number of issues need to be addressed, including the specification of the surface as the union of high-order patches, the incorporation of accurate quadrature rules for integrating singular or weakly singular Green's functions on such patches, and their coupling to the oct-tree data structures on which the FMM separates near and far field interactions. Although the latter is straightforward for point distributions, the near field for a patch is determined by its physical dimensions, not the distribution of discretization points on the surface.

Here, we present a general framework for efficiently coupling locally corrected quadratures with FMMs, relying primarily on what are called generalized Gaussian quadratures rules, supplemented by adaptive integration. The approach, however, is quite general and easily applicable to other schemes, such as Quadrature by Expansion (QBX). We also introduce a number of accelerations to reduce the cost of quadrature generation itself, and present several numerical examples of acoustic scattering that demonstrate the accuracy, robustness, and computational efficiency of the scheme. On a single core of an Intel i5 2.3 GHz processor, a Fortran implementation of the scheme can generate near field quadrature corrections for between 1000 and 10,000 points per second, depending on the order of accuracy and the desired precision. A Fortran implementation of the algorithm described in this work is available at <https://gitlab.com/fastalgorithms/fmm3dbie>.

© 2021 The Authors. Published by Elsevier Inc. This is an open access article under the CC BY-NC-ND license (<http://creativecommons.org/licenses/by-nc-nd/4.0/>).

1. Introduction

Over the last two decades, fast multipole methods (FMMs) and related hierarchical fast algorithms have become widespread for computing N -body interactions in computational chemistry, astrophysics, acoustics, fluid dynamics, and electromagnetics. In the time-harmonic, acoustic setting, a typical computation of interest is the evaluation of

* Corresponding author.

E-mail addresses: greengard@cims.nyu.edu (L. Greengard), oneil@cims.nyu.edu (M. O'Neil), mrachh@flatironinstitute.org (M. Rachh), felipe.vico@gmail.com (F. Vico).

<https://doi.org/10.1016/j.jcpx.2021.100092>

2590-0552/© 2021 The Authors. Published by Elsevier Inc. This is an open access article under the CC BY-NC-ND license (<http://creativecommons.org/licenses/by-nc-nd/4.0/>).

$$F_m = \sum_{\substack{n=1 \\ n \neq m}}^N \sigma_n G_k(\mathbf{x}_m, \mathbf{x}_n) \quad (1.1)$$

where

$$G_k(\mathbf{x}, \mathbf{y}) = \frac{e^{ik|\mathbf{x}-\mathbf{y}|}}{4\pi|\mathbf{x}-\mathbf{y}|} \quad (1.2)$$

is the free-space Green's function for the Helmholtz equation

$$\Delta u + k^2 u = 0.$$

Direct calculation of Equation (1.1) requires $\mathcal{O}(N^2)$ operations, while the FMM requires $\mathcal{O}(N)$ work in the low frequency regime [1] and $\mathcal{O}(N \log N)$ work in the high frequency regime [2].

When solving boundary value problems for partial differential equations (PDEs) in three dimensions, such sums arise in the discretization of layer potentials defined on a surface Γ . These potentials take the form:

$$u(\mathbf{x}) = \int_{\Gamma} K(\mathbf{x}, \mathbf{x}') \sigma(\mathbf{x}') da(\mathbf{x}'). \quad (1.3)$$

In (1.3), $K(\mathbf{x}, \mathbf{x}')$ is a Green's function for the PDE of interest, such as (1.2) or one of its directional derivatives. As a result, the governing equation is automatically satisfied, and it remains only to enforce the desired boundary condition. With a suitable choice for the kernel $K(\mathbf{x}, \mathbf{x}')$, this often leads to a Fredholm integral equation of the form

$$\sigma(\mathbf{x}) + \int_{\Gamma} K(\mathbf{x}, \mathbf{x}') \sigma(\mathbf{x}') da(\mathbf{x}') = f(\mathbf{x}), \quad \text{for } \mathbf{x} \in \Gamma. \quad (1.4)$$

As we shall see below, this can be discretized with high-order accuracy using a suitable Nyström method [3,4]

$$\sigma_i + w_{ii}\sigma_i + \sum_{j \neq i} w_{ij} K(\mathbf{x}_i, \mathbf{x}_j) \sigma_j = f(\mathbf{x}_i). \quad (1.5)$$

Here, \mathbf{x}_i and w_{ij} are the quadrature nodes and weights, respectively, while σ_i is an approximation to the true value $\sigma(\mathbf{x}_i)$. If the quadrature weights w_{ij} did not depend on the target location, i.e. $w_{ij} = w_j$, then the above sum is a standard N -body calculation of the form (1.1).

Unfortunately, when the integral equation comes from a layer potential corresponding to an elliptic PDE, the kernel K is typically singular or weakly singular so that simple high-order rules for smooth functions fail. Assuming the surface Γ is defined as the union of many smooth patches Γ_j (each with its own parameterization), high-order quadrature schemes require an analysis of the distance of the *target* \mathbf{x}_i from each *patch*.

More precisely, for a given target location \mathbf{x}_i on the boundary, the integral in (1.3) or (1.4) can be split into three pieces: a self-interaction integral, a near field integral, and a far field integral:

$$\int_{\Gamma} K(\mathbf{x}_i, \mathbf{x}') \sigma(\mathbf{x}') da(\mathbf{x}') = \int_{\text{Self}(\mathbf{x}_i)} K(\mathbf{x}_i, \mathbf{x}') \sigma(\mathbf{x}') da(\mathbf{x}') + \int_{\text{Near}(\mathbf{x}_i)} K(\mathbf{x}_i, \mathbf{x}') \sigma(\mathbf{x}') da(\mathbf{x}') + \int_{\text{Far}(\mathbf{x}_i)} K(\mathbf{x}_i, \mathbf{x}') \sigma(\mathbf{x}') da(\mathbf{x}'). \quad (1.6)$$

This splitting into target-dependent regions is essential for maintaining high-order accuracy. The region $\text{Self}(\mathbf{x}_i)$ is simply the patch on which \mathbf{x}_i itself lies. The integral over this patch involves a singular integrand (due to the kernel K). The *Near* field is defined precisely in Section 4, but consists of patches close enough to \mathbf{x}_i such that the integrand is *nearly singular* even though it is formally smooth. The *Far* region consists of all other patches, sufficiently far from \mathbf{x}_i such that high-order quadratures for smooth functions can be applied.

Definition 1. Suppose that \mathbf{x}_i is in the far field of a patch Γ_m , and that

$$\sum_{j=1}^M w_j K(\mathbf{x}_i, \mathbf{s}_j) \sigma_j \approx \int_{\Gamma_m} K(\mathbf{x}_i, \mathbf{x}') \sigma(\mathbf{x}') da(\mathbf{x}') \quad (1.7)$$

to the desired precision. Then \mathbf{s}_j and w_j will be referred to as the *far field* quadrature nodes and weights. Note that these are independent of \mathbf{x}_i .

A related task in the solution process is evaluating the computed solution $u(\mathbf{x})$ at target locations \mathbf{x} which are off-surface, but possibly arbitrarily close to the surface. In this case as well, while the integrands are formally smooth, evaluating the integral presents a similar challenge owing to the nearly singular behavior of the integrand. The evaluation of the potential can still be split into two pieces in a similar manner: a near field integral and a far field integral

$$\int_{\Gamma} K(\mathbf{x}, \mathbf{x}') \sigma(\mathbf{x}') da(\mathbf{x}') = \int_{\text{Near}(\mathbf{x})} K(\mathbf{x}, \mathbf{x}') \sigma(\mathbf{x}') da(\mathbf{x}') + \int_{\text{Far}(\mathbf{x})} K(\mathbf{x}, \mathbf{x}') \sigma(\mathbf{x}') da(\mathbf{x}'). \quad (1.8)$$

The strategy outlined above can be used for evaluating the contributions of the near and far regions for targets off-surface as well.

While methods exist for the Self and Near calculations, the use of a fast algorithm such as the FMM requires coupling these somewhat complicated quadrature schemes to the Cartesian oct-tree data structures that divide up space into a hierarchy of regular, adaptively refined cubes. Unfortunately, the surface patches Γ_m of the domain boundary Γ may be of vastly different sizes and do not, in general, conform to a spatial subdivision strategy based on the density of quadrature nodes as points in \mathbb{R}^3 . That is, many patches (curvilinear triangles) Γ_j are likely to cross leaf node boundaries in the oct-tree structure. If this were not the case, then far field interactions within the FMM could be handled by computing multipole expansions from entire patches, followed by direct calculation of the Self and Near quadratures (analogous to the direct, near neighbor calculations in a point-based FMM). Thus, one of the issues we address here concerns modifications of the FMM so that speed is conserved for far field interactions, but in a manner where the overhead for near field quadrature corrections is modest, even when the surface patches are nonuniform.

Furthermore, while we will restrict our attention here to Nyström-style discretizations, the same concerns must be addressed for collocation and Galerkin-type methods. Similar issues arise when coupling adaptive mesh refinement (AMR) data structures to complicated geometries using Cartesian cut-cell methods [5–7].

In the present paper, we develop an efficient algorithm which allows for the straightforward coupling of adaptive FMM data structures with locally corrected quadrature schemes. Our goal is to achieve $\mathcal{O}(N)$ or $\mathcal{O}(N \log N)$ performance for surfaces with $\mathcal{O}(N)$ discretizations points, depending on whether one is in the low or high frequency regime, respectively. Moreover, we would like the constant implicit in this notation to be as close to the performance of point-based FMMs as possible. We will concentrate on the use of generalized Gaussian quadrature rules [8–11] for the self interactions on curvilinear triangles, and adaptive integration for the Near region (nearly singular interactions). It is, perhaps, surprising that adaptive integration on surface patches can be competitive with other schemes such as Quadrature By Expansion (QBX), singularity subtraction, or coordinate transformations [12–19]. The key is that we have developed a careful, precision and geometry-dependent hierarchy of interpolators on each patch, after which the adaptive integration step is inexpensive when amortized over all relevant targets. As a side-effect, our scheme also provides rapid access to entries of the fully discretized system matrix which is essential for fast direct solvers.

Remark 1. Generalized Gaussian quadrature was already coupled to an FMM in [8], but the question of how to design a robust algorithm that is insensitive to large variation in triangle dimensions was not directly addressed. In some sense, the present paper is devoted to two separate issues raised in [8]: the first is to accelerate adaptive quadrature itself, and the second is to describe an FMM implementation that works for multi-scale discretizations.

Remark 2. It is worth noting that most locally corrected quadrature schemes, such as Duffy transformations [20], are designed for a target that is mapped to the origin of a local coordinate system or the vertex of a triangular patch and, hence, are suitable only for self interactions as defined above. Near interactions are not addressed. An exception is Quadrature by Expansion (QBX) which provides a systematic, uniform procedure for computing layer potentials using only smooth quadratures and extrapolation [21,16,17]. These have been successfully coupled to FMMs in [17,18]. Many aspects of the FMM modifications described here can be used in conjunction with QBX instead. Another exception is Erichsen-Sauter rules [15], which do include schemes for adjacent panels, but appear to be best suited for modest accuracy.

The paper is organized as follows: in Section 2, some basic facts regarding polynomial approximation and integration on triangles are presented. In Section 3, we describe the classical boundary integral equation for acoustic scattering from a *sound-soft* boundary, governed by the Helmholtz equation, as well as discretization and integration methods for curvilinear surfaces. Section 4 provides the algorithmic details involved in locally corrected quadrature schemes. The coupling of these quadrature schemes to FMMs is presented in Section 5, and numerical examples demonstrating the performance of the scheme are presented in Section 6. Finally, in Section 7, we discuss avenues for further research, and the application of our scheme to fast direct solvers.

2. Interpolation and integration on triangles

For the sake of simplicity, we assume that we are given a surface triangulation represented as a collection of charts $\mathbf{X}^j = \mathbf{X}^j(u, v)$, which map the standard right triangle

$$T_0 = \{(u, v) : u \geq 0, v \geq 0, u + v \leq 1\} \subset \mathbb{R}^2 \quad (2.1)$$

to the surface patch Γ_j . All discretization and integration is done over T_0 , incorporating the mapping function \mathbf{X}^j and its derivatives as needed.

In this section, we summarize the basic polynomial interpolation and integration rules we will use for smooth functions $f : T_0 \rightarrow \mathbb{R}$. A useful spectral basis is given by the orthogonal polynomials on T_0 , known as Koornwinder polynomials [22]. They are described analytically by the formula:

$$K_{nm}(u, v) = c_{nm} (1-v)^m P_{n-m}^{(0, 2m+1)}(1-2v) P_m\left(\frac{2u+v-1}{1-v}\right), \quad m \leq n, \quad (2.2)$$

where $P_n^{(a,b)}$ is the Jacobi polynomial of degree n with parameters (a, b) , P_m is the Legendre polynomial of degree m , and c_{nm} is a normalization constant such that

$$\int_{T_0} |K_{nm}(u, v)|^2 du dv = 1. \quad (2.3)$$

For convenience, our definition is slightly different from that in [22].

It is easy to see that there are $n_p = p(p+1)/2$ Koornwinder polynomials of total degree less than p . By a straightforward change of variables, and using the orthogonality relationships for Legendre and Jacobi polynomials, it is easy to show that

$$\int_{T_0} K_{nm}(u, v) K_{n'm'}(u, v) du dv = 0, \quad \text{for } n \neq n' \text{ and } m \neq m'. \quad (2.4)$$

The Koornwinder polynomials form a complete basis for $L^2(T_0)$, and can easily be used to approximate smooth functions on T_0 to arbitrary precision.

2.1. Polynomial approximation and integration

As in standard spectral approximation methods for functions defined on intervals or tensor products of intervals [23], smooth functions f defined on T_0 can be interpolated, approximated, and integrated using a Koornwinder polynomial basis.

To this end, suppose that f is defined by a p th-order Koornwinder expansion with coefficients c_{nm} :

$$f(u, v) = \sum_{n=0}^{p-1} \sum_{m=0}^n c_{nm} K_{nm}(u, v). \quad (2.5)$$

Then, the square matrix \mathbf{U} that maps the n_p coefficients in the above expansion to values of f at a selection of n_p interpolation nodes, denoted by $(u_j, v_j) \subseteq T_0$, has elements

$$U_{nm,j} = K_{nm}(u_j, v_j). \quad (2.6)$$

Let the matrix $\mathbf{V} = \mathbf{U}^{-1}$. Then \mathbf{V} maps values of f at the interpolation nodes (u_j, v_j) to coefficients in a Koornwinder polynomial expansion.

Suppose now that $f : T_0 \rightarrow \mathbb{R}$ is an arbitrary smooth function (not necessarily a polynomial) and let the values of f at the interpolation points be denoted by $f_i = f(u_i, v_i)$. Then, a p th-order approximation to f is given by

$$f(u, v) \approx \sum_{n=0}^{p-1} \sum_{m=0}^n c_{nm} K_{nm}(u, v), \quad (2.7)$$

where

$$c_{nm} = \sum_{j=1}^{n_p} V_{(nm),j} f_j. \quad (2.8)$$

We define the *conditioning of the interpolation procedure* as the condition number of the matrices \mathbf{U} or \mathbf{V} . Much like interpolation operators on the interval, these matrices are not well-conditioned for arbitrary selections of nodes (u_j, v_j) , as we will briefly discuss in the next two sections.

An n -point quadrature rule for computing the integral of a function f on T_0 is a collection of nodes and weights, (u_i, v_i) , w_i , $i = 1, 2, \dots, n$, such that

$$\int_{T_0} f(u, v) du dv = \int_0^1 \int_0^{1-u} f(u, v) du dv \approx \sum_{i=1}^n w_i f_i, \tag{2.9}$$

where $f_i = f(u_i, v_i)$ is the value of f at the i th quadrature node. The accuracy of such a quadrature rule is very dependent on the choice of the quadrature nodes; if the rule is to be exact for a selection of n functions, then the values of w_i are determined wholly by the selection of (u_i, v_i) . Since the node selection provides additional degrees of freedom, Gaussian-type quadrature rules are possible. On the triangle, a quadrature rule would be perfectly Gaussian if it integrated $3n$ functions exactly, since there are $3n$ parameters (the two coordinates of the nodes, and the weights). In one dimension, it is well-known that choosing the nodes as the roots of a suitable orthogonal polynomial leads to a perfect Gaussian rule [23]. In two dimensions, such perfect rules do not exist, but approximately Gaussian quadrature rules can be constructed.

2.2. High-order quadrature rules on the simplex T_0

First described in 2010 [24], what we will refer to as Xiao-Gimbutas quadratures are a set of Gaussian-like rules obtained through the solution of a nonlinear least-squares problem using Newton's method. The resulting nodes are contained in the interior of T_0 , and all the weights are positive. Various kinds of symmetry can also be specified. For a given $p > 0$, these rules are designed to use the minimum number of nodes with positive weights so that the resulting quadrature rule is exact for all polynomials of total degree $< p$. As noted above, there are $n_p = p(p + 1)/2$ such polynomials. While not perfect Gaussian rules, the Xiao-Gimbutas quadratures achieve remarkably high-order. A rule with 48 weights and nodes, for example, is exact for polynomials of degree 16, of which there are $n_{16} = 136$. The rule has only $3 \times 48 = 144$ free parameters.

For the sake of convenience we would like the quadrature nodes to serve as interpolation/approximation nodes as well. There are, however, far fewer Xiao-Gimbutas nodes than functions we would like to interpolate (namely n_p). Instead of using an even higher order Xiao-Gimbutas rule, with at least n_p nodes, we choose an alternative quadrature scheme, introduced by Vioreanu and Rokhlin in 2014 [25]. The Vioreanu-Rokhlin nodes of order p , obtained via a similar optimization procedure, are a collection of n_p nodes which can be used simultaneously for high-order polynomial interpolation, approximation, and integration on T_0 . We refer the reader to [25] for a thorough discussion. For our purposes, it suffices to note that the interpolation operators computed using these nodes are extremely well-conditioned.

As a quadrature rule, the nodes and weights are Gaussian-like; they have positive weights and integrate more polynomials than there are nodes in the quadrature. For example, the Vioreanu-Rokhlin rule that interpolates polynomials of degree $p = 16$, with $n_{16} = 136$ nodes, integrates all 378 polynomials of degree up to $p' = 27$. A perfect Gaussian rule would integrate $3n_p = 408$ functions exactly. The relationship between the order of the interpolation scheme p and the order of the quadrature p' is somewhat complicated, and obtained empirically [25].

3. Acoustic scattering from a sound-soft boundary

Let Ω be a bounded region in \mathbb{R}^3 , with smooth boundary $\partial\Omega = \Gamma$. Given a function f defined on Γ , a function u defined in $\mathbb{R}^3 \setminus \Omega$ is said to satisfy the exterior Dirichlet problem for the Helmholtz equation if

$$\begin{aligned} (\Delta + k^2)u &= 0 && \text{in } \mathbb{R}^3 \setminus \Omega, \\ u &= f && \text{on } \Gamma, \end{aligned} \tag{3.1}$$

$$\lim_{r \rightarrow \infty} r \left(\frac{\partial u}{\partial r} - iku \right) = 0.$$

In acoustics, Dirichlet problems such as this arise when $\partial\Omega$ is *sound-soft* and $f = -u^{in}$, where u^{in} is an impinging acoustic wave. A standard approach for solving the Dirichlet problem is to let

$$u = \mathcal{D}_k[\sigma] - ik \mathcal{S}_k[\sigma], \tag{3.2}$$

where σ is an unknown density function defined on Γ . Here \mathcal{S} and \mathcal{D} are the single layer and double layer operators, respectively, given by

$$\mathcal{S}_k[\sigma](\mathbf{x}) = \int_{\Gamma} G_k(\mathbf{x}, \mathbf{y}) \sigma(\mathbf{y}) da(\mathbf{y}), \tag{3.3}$$

$$\mathcal{D}_k[\sigma](\mathbf{x}) = \int_{\Gamma} (\mathbf{n}(\mathbf{y}) \cdot \nabla_{\mathbf{y}} G_k(\mathbf{x}, \mathbf{y})) \sigma(\mathbf{y}) da(\mathbf{y}), \tag{3.4}$$

where $\mathbf{n}(\mathbf{y})$ is the outward normal at $\mathbf{y} \in \Gamma$, and $G_k(\mathbf{x}, \mathbf{y})$ is given by (1.2). The representation (3.2) automatically satisfies the Helmholtz equation and the radiation condition in (3.1). Imposing the boundary condition, and using standard jump relations for layer potentials [26], we obtain the following second-kind integral equation along Γ for the density σ :

$$\frac{1}{2}\sigma(\mathbf{x}) + \mathcal{D}_k[\sigma](\mathbf{x}) - ik\mathcal{S}_k[\sigma](\mathbf{x}) = f(\mathbf{x}), \quad \mathbf{x} \in \Gamma. \quad (3.5)$$

This involves a slight abuse of notation: for $\mathbf{x} \in \Gamma$, $\mathcal{D}_k[\sigma](\mathbf{x})$ should be evaluated in the principal value sense.

When solving (3.5), the accurate evaluation of the layer potentials $\mathcal{S}_k[\sigma]$, $\mathcal{D}_k[\sigma]$ on Γ is essential for either direct or iterative solvers. We will focus here on the evaluation of the single layer potential $\mathcal{S}[\sigma]$, assuming σ is known. Only minor modifications are needed to address the double layer potential, as well as other scalar or vector-valued layer potentials that arise in electrostatics, elastostatics, viscous flow, or electromagnetics.

3.1. Surface parameterizations

While some simple boundaries (such as spheres, ellipsoids and tori) can be described by global parameterizations, in general it is necessary to describe a complicated surface Γ as a collection of surface patches, each of which is referred to as a *chart*. The collection of charts whose union defines Γ will be referred to as an *atlas*.

More precisely, we assume that the surface is the disjoint union of patches Γ_j

$$\Gamma = \cup_{j=1}^{N_{\text{patches}}} \Gamma_j, \quad (3.6)$$

and that the patch Γ_j is parameterized by a non-degenerate chart $\mathbf{X}^j : T_0 \rightarrow \Gamma_j$, where T_0 is the standard simplex (2.1). Given these charts \mathbf{X}^j , a local coordinate system can be defined on patch Γ_j by taking its partial derivatives. For this, we define

$$\mathbf{X}_u^j \equiv \frac{\partial \mathbf{X}^j}{\partial u}, \quad \mathbf{X}_v^j \equiv \frac{\partial \mathbf{X}^j}{\partial v}, \quad \mathbf{n}^j \equiv \mathbf{X}_u^j \times \mathbf{X}_v^j. \quad (3.7)$$

Finally, we assume that the triplet $\mathbf{X}_u^j, \mathbf{X}_v^j, \mathbf{n}^j$ forms a right-handed coordinate system with \mathbf{n}^j pointing into the unbounded region $\mathbb{R}^3 \setminus \Omega$. In general, these vectors are neither orthogonal nor of unit length. The area element on the patch Γ_j is determined by the Jacobian J^j ,

$$\begin{aligned} da(\mathbf{X}^j) &= \left| \mathbf{X}_u^j \times \mathbf{X}_v^j \right| du dv \\ &= J^j(u, v) du dv. \end{aligned} \quad (3.8)$$

3.2. Discretization and integration

If f is a function defined on Γ , then its integral can be decomposed as a sum over patches:

$$\begin{aligned} \int_{\Gamma} f(\mathbf{x}) da(\mathbf{x}) &= \sum_{j=1}^{N_{\text{patches}}} \int_{\Gamma_j} f(\mathbf{x}) da(\mathbf{x}) \\ &= \sum_{j=1}^{N_{\text{patches}}} \int_{T_0} f(\mathbf{X}^j(u, v)) J^j(u, v) dudv \\ &= \sum_{j=1}^{N_{\text{patches}}} \int_{u=0}^1 \int_{v=0}^{1-u} f(\mathbf{X}^j(u, v)) J^j(u, v) dudv, \end{aligned} \quad (3.9)$$

where the Jacobian is given in (3.8).

If, in addition, f is smooth, then each of the integrals on T_0 in (3.9) can be approximated using Xiao-Gimbutas or Vioreanu-Rokhlin quadrature rules, as discussed in Section 2. Using the latter, we have

$$\int_{\Gamma} f(\mathbf{x}) da(\mathbf{x}) \approx \sum_{j=1}^{N_{\text{patches}}} \sum_{\ell=1}^{n_p} w_{\ell} f(\mathbf{X}^j(u_{\ell}, v_{\ell})) J^j(u_{\ell}, v_{\ell}), \quad (3.10)$$

where n_p is the number of nodes in the quadrature, which varies depending on the desired order of accuracy.

For a patch Γ_j and a target \mathbf{x} , however, the integrand appearing in $\mathcal{S}_k[\sigma](\mathbf{x})$ is only smooth when \mathbf{x} is in the far field. When \mathbf{x} is either on Γ_j or nearby, we will need to modify our quadrature approach, as described at the outset. Furthermore,

in practice, we will only be given approximations of the charts \mathbf{X}^j and the function f (or the density σ) on each patch to finite order.

In what follows, we define a p th-order approximation as one for which the truncation error is $\mathcal{O}(h^p)$, where h is, say, the diameter of the patch Γ_j . For a scalar function f , it can be approximated to p th-order in $L^2(T_0)$ using Koornwinder polynomials as

$$f(u, v) \approx \sum_{n+m < p} f_{nm} K_{nm}(u, v). \tag{3.11}$$

The coefficients f_{nm} can be computed using the values-to-coefficients matrix \mathbf{V} , as described in Section 2.1. The charts \mathbf{X}^j will generally be approximated using a vector version of the above formula:

$$\begin{aligned} \mathbf{X}^j(u, v) &\approx \sum_{n+m < p} \begin{pmatrix} x_{nm}^j \\ y_{nm}^j \\ z_{nm}^j \end{pmatrix} K_{nm}(u, v) \\ &= \sum_{n+m < p} \mathbf{x}_{nm}^j K_{nm}(u, v) \end{aligned} \tag{3.12}$$

It is important to note that even if the charts \mathbf{X}^j are approximated to accuracy ε , i.e.

$$\mathbf{X}^j(u, v) = \sum_{n+m < p} \mathbf{x}_{nm}^j K_{nm}(u, v) + \mathcal{O}(\varepsilon), \tag{3.13}$$

it does not follow that the Jacobian J^j will be evaluated to precision ε as well. The function J^j is non-linear and usually requires a higher-order approximation than the individual components of the chart itself. This cannot be avoided unless analytic derivative information is provided for each patch Γ_j . This affects the accuracy of numerical approximations to surface integrals for a fixed set of patches (but not the asymptotic convergence rate).

Remark 3. In the following, we use the same order of discretization for representing the layer potential densities σ and the surface information, i.e. the charts \mathbf{X}^j and their derivatives \mathbf{X}_u^j , and \mathbf{X}_v^j . While this choice is made for convenience of notation and software implementation, our approach to evaluating layer potentials extends in a straightforward manner to the case where different orders of discretization are used for representing the surface information and layer potentials densities.

4. Locally corrected quadratures

For a target location $\mathbf{x} \in \Gamma_j$, let us first consider the *self interaction*

$$\mathcal{S}_{\text{Self}}[\sigma](\mathbf{x}) = \int_{\Gamma_j} G_k(\mathbf{x}, \mathbf{y}) \sigma(\mathbf{y}) da(\mathbf{y}). \tag{4.1}$$

In [8,9], the authors designed quadrature rules for exactly this purpose, under the assumption that σ and J^j are well-approximated by polynomials (and therefore representable by Koornwinder expansions). The quadrature schemes in these papers involve a rather intricate set of transformations but yield a set of precomputed tables which, when composed with the chart \mathbf{X}^j , yield the desired high-order accuracy. Briefly, for any $\mathbf{x} \in \Gamma_j$, and all σ of the form $\sigma(u, v) = \sum_{nm} c_{nm} K_{nm}(u, v)$, there exist $N(\mathbf{x})$ nodes $(u_{\mathbf{x},\ell}, v_{\mathbf{x},\ell}) \in T_0$ and associated quadrature weights $w_{\mathbf{x},\ell}$, such that

$$\begin{aligned} \left| \int_{u=0}^1 \int_{v=0}^{1-u} G_k(\mathbf{x}, \mathbf{X}^j(u, v)) \sigma(u, v) J^j(u, v) du dv \right. \\ \left. - \sum_{\ell=1}^{N(\mathbf{x})} G_k(\mathbf{x}, \mathbf{X}^j(u_{\mathbf{x},\ell}, v_{\mathbf{x},\ell})) \sigma(u_{\mathbf{x},\ell}, v_{\mathbf{x},\ell}) J^j(u_{\mathbf{x},\ell}, v_{\mathbf{x},\ell}) w_{\mathbf{x},\ell} \right| \leq \varepsilon \cdot \|\sigma\|_{\mathbb{L}^2(\Gamma_j)}. \end{aligned} \tag{4.2}$$

As mentioned in the introduction, in the original paper [8], which was focused on quadrature design, a simple coupling to FMMs was mentioned that relied on the underlying discretization being uniformly high-order. Near field interactions were done using *on-the-fly* adaptive integration. In their subsequent paper [9], this type of expensive adaptive integration was used for all non-self interactions (i.e. no FMM-type acceleration was used at all). Such an approach cannot be directly accelerated with standard FMMs since the effective quadrature weights are functions of both the source and target locations.




Fig. 1. The smallest sphere containing surface patch Γ_j centered at \mathbf{c}_j and the near field region $N_\eta(\Gamma_j)$. $\eta > 1$ is a free parameter whose selection is based on the order of accuracy of the far field quadrature.

Recall that for $\mathbf{x} \in \mathbb{R} \setminus \Gamma$, we split the single layer potential $\mathcal{S}_k[\sigma](\mathbf{x})$ into two pieces:

$$\begin{aligned} \mathcal{S}_k[\sigma](\mathbf{x}) &= \int_{\Gamma} G_k(\mathbf{x}, \mathbf{y}) \sigma(\mathbf{y}) da(\mathbf{y}) \\ &= \sum_{\ell=1}^{N_{\text{patches}}} \int_{\Gamma_\ell} G_k(\mathbf{x}, \mathbf{y}) \sigma(\mathbf{y}) da(\mathbf{y}) = \mathcal{S}_{\text{Near}}[\sigma](\mathbf{x}) + \mathcal{S}_{\text{Far}}[\sigma](\mathbf{x}), \end{aligned} \quad (4.3)$$

and when \mathbf{x} lies on the boundary, say on patch Γ_j , then the single layer potential $\mathcal{S}[\sigma](\mathbf{x})$ is split into three pieces:

$$\mathcal{S}_k[\sigma](\mathbf{x}) = \mathcal{S}_{\text{Self}}[\sigma](\mathbf{x}) + \mathcal{S}_{\text{Near}}[\sigma](\mathbf{x}) + \mathcal{S}_{\text{Far}}[\sigma](\mathbf{x}), \quad (4.4)$$

where $\mathcal{S}_{\text{Self}}[\sigma](\mathbf{x})$ is defined in (4.1), and the near and far regions associated with a target and the corresponding definitions of $\mathcal{S}_{\text{Near}}[\sigma](\mathbf{x})$ and $\mathcal{S}_{\text{Far}}[\sigma](\mathbf{x})$ are described below.

For this, it turns out to be easier to first take the point of view of a patch rather than a target. Let \mathbf{c}_j denote the centroid of the patch Γ_j ,

$$\begin{aligned} \mathbf{c}_j &= \int_{\Gamma_j} \mathbf{x} da(\mathbf{x}) \\ &= \int_{T_0} \mathbf{X}^j(u, v) du dv, \end{aligned} \quad (4.5)$$

and let

$$R_j = \min_{R>0} \{R \mid \Gamma_j \subset B_R(\mathbf{c}_j)\}, \quad (4.6)$$

where $B_R(\mathbf{c}_j)$ is the ball of radius R centered at \mathbf{c}_j . That is, $B_{R_j}(\mathbf{c}_j)$ is the ball of minimal radius containing the patch Γ_j . Letting $\eta > 1$ be a free parameter for the moment, we define the near field of the patch Γ_j , denoted by $N_\eta(\Gamma_j)$, to be the set of points that do not lie on Γ_j but are within the ball $B_{\eta R_j}(\mathbf{c}_j)$ (see Fig. 1). Thus,

$$N_\eta(\Gamma_j) = \{\mathbf{x} \in \mathbb{R}^3 \setminus \Gamma_j \mid d(\mathbf{c}_j, \mathbf{x}) \leq \eta R_j\}. \quad (4.7)$$

Given the collection of near field regions of the form $N_\eta(\Gamma_j)$, let $T_\eta(\mathbf{x})$ denote the dual list: that is, the collection of patches Γ_j for which the point $\mathbf{x} \in N_\eta(\Gamma_j)$,

$$T_\eta(\mathbf{x}) = \{\Gamma_j \mid \mathbf{x} \in N_\eta(\Gamma_j)\}. \quad (4.8)$$

Similarly, for $\mathbf{x} \in \mathbb{R}^3 \setminus \Gamma$, we denote the far field of \mathbf{x} by

$$F_\eta(\mathbf{x}) = \{\Gamma_j \mid \mathbf{x} \notin N_\eta(\Gamma_j)\}. \quad (4.9)$$

When $\mathbf{x} \in \Gamma_i$ is a boundary point, we let

$$F_\eta(\mathbf{x}) = \{\Gamma_j, j \neq i \mid \mathbf{x} \notin N_\eta(\Gamma_j)\}. \quad (4.10)$$

Then referring to the near-far split of the layer potential for targets off and on-surface described in Equation (4.3), and Equation (4.4), respectively, the near and far part of the layer potentials $\mathcal{S}_{\text{Near}}$, and \mathcal{S}_{Far} are given by

$$\mathcal{S}_{\text{Near}}[\sigma](\mathbf{x}) = \sum_{\Gamma_\ell \in T_\eta(\mathbf{x})} \int_{\Gamma_\ell} G_k(\mathbf{x}, \mathbf{y}) \sigma(\mathbf{y}) da(\mathbf{y}) \quad (4.11)$$

and

$$\mathcal{S}_{\text{Far}}[\sigma](\mathbf{x}) = \sum_{\Gamma_\ell \in F_\eta(\mathbf{x})} \int_{\Gamma_\ell} G_k(\mathbf{x}, \mathbf{y}) \sigma(\mathbf{y}) da(\mathbf{y}). \quad (4.12)$$

As noted in the beginning of the section, $\mathcal{S}_{\text{Self}}[\sigma](\mathbf{x})$ can be computed using the generalized Gaussian quadratures of [8,9]. By virtue of their separation from the source patches, all of the integrands in \mathcal{S}_{Far} are smooth and can be computed using either Vioreanu-Rokhlin or Xiao-Gimbutas quadrature rules, with weights that are independent of the target location \mathbf{x} . The accuracy of these rules, which is affected by the free parameter η , is discussed in Section 4.1.

It remains only to develop an efficient scheme for evaluating the integrals which define $\mathcal{S}_{\text{Near}}[\sigma]$. At present, there do not exist quadrature rules that are capable of simultaneously accounting for the singularity in the Green's function and the local geometric variation in an efficient manner. The approach developed below involves a judicious combination of precomputation and a greedy, adaptive algorithm applied, for every target point \mathbf{x} , to each $\Gamma_\ell \in T_\eta(\mathbf{x})$. Once the near field quadratures have been computed, they can be saved using only $O(N)$ storage. When solving an integral equation iteratively, this can be used to accelerate subsequent applications of the integral operator.

Remark 4. The evaluation of near field quadratures does *not* affect the overall complexity of computing layer potentials, assuming η is not too large. This follows from the fact that, for each target \mathbf{x} , there are only $O(1)$ patches contained in $T_\eta(\mathbf{x})$. Thus, the cost of all near field contributions is of the order $O(N)$. Since there can be several patches in the near field, however, this computation tends to be the rate limiting step in the overall quadrature generation procedure.

Remark 5. Without entering into a detailed literature review, it should be noted that coordinate transformation methods such as those in [12,14,15,19] can also be used for computing near field interactions for surface targets. However, these methods don't apply easily to off-surface evaluation. An alternative to our procedure is Quadrature by Expansion (QBX) [21,16–18], which handles singular and nearly-singular integrals in a unified manner and (like the method of this paper) works both on and off surface. There are distinct trade-offs to be made in QBX-based schemes and the scheme presented here. In the end, the best method will be determined by accuracy, efficiency and ease of use. At present, we have found the adaptive quadrature approach to be the most robust and fastest in terms of overall performance.

4.1. Selecting the near field cutoff

In this section, we discuss the choice of the parameter η , which defines the near field for each patch (see Fig. 1). Once η is fixed, accuracy considerations will determine whether the interpolation nodes used on each patch are sufficient for accurate calculation of the far field \mathcal{S}_{Far} , or whether we will need to increase the order of the far field quadrature.

As η increases, the near field for each patch obviously grows, so that the number of targets for which we will apply specialized quadrature increases. Since we would like to store these near field quadratures for the purpose of repeated application of the integral operator, both the storage and CPU requirements also grow accordingly. On the other hand, as η decreases, the integrand in \mathcal{S}_{Far} becomes less smooth, and the number of quadrature nodes needed to achieve the desired precision in the far field will grow. If it exceeds the number of original nodes n_p to achieve p th-order convergence, we will have to *oversample* the layer potential density σ . That is, we will have to interpolate σ to a larger number of quadrature nodes. This increases the computational cost of evaluating the far field via the FMM. Balancing the cost of the near field and far field interactions sets the optimal value for η .

Based on extensive numerical experiments, we have found that for the highest order methods ($p > 8$), $\eta = 1.25$ works well. For orders of accuracy $4 < p \leq 8$, we recommend $\eta = 2$, and for the lowest orders of accuracy $p \leq 4$, we recommend $\eta = 2.75$.

4.2. Oversampling via p -refinement

For a patch Γ_j , suppose that an order q Vioreanu-Rokhlin quadrature rule is the smallest order quadrature rule which accurately computes the contribution of Γ_j to all targets $\mathbf{x} \notin N_\eta(\Gamma_j)$ to the desired precision ε . Then, the oversampling factor for Γ_j is defined as the ratio of $q(q+1)/(p(p+1))$. For a given η , rather than trying to estimate the oversampling factor

needed for S_{Far} analytically, we compute the far field quadrature order q needed for a specified precision numerically. For each patch Γ_j , we first identify the 10 farthest targets in $N_\eta(\Gamma_j)$. If the cardinality of $N_\eta(\Gamma_j)$ is less than 20, we choose the farthest $|N_\eta(\Gamma_j)|/2$ targets from the list, and append 15 randomly chosen targets on the boundary of the sphere $\partial B_{\eta R_j}(\mathbf{c}_j)$. We denote this set of targets by $F(\Gamma_j)$. For the nm -th Koornwinder polynomial K_{nm} , let

$$I_{nm}^j(\mathbf{x}) = \int_{T_0} G_k(\mathbf{x}, \mathbf{X}^j(u, v)) K_{nm}(u, v) J^j(u, v) du dv, \quad (4.13)$$

and let $\tilde{I}_{nm,q}^j(\mathbf{x})$ denote the approximation to the integral computed using the q th-order Vioreanu-Rokhlin quadrature. Then, the far field order q_j for patch Γ_j is chosen according to the following criterion: q_j is the smallest q such that all of the integrals $\tilde{I}_{nm,q}^j(\mathbf{x})$, for $0 \leq m \leq n \leq p$ and $\mathbf{x} \in F(\Gamma_j)$, agree to a prescribed tolerance $\varepsilon d_j / \|\mathbf{V}\|$ with the corresponding integrals obtained from using a $(q+1)$ th-order Vioreanu-Rokhlin quadrature. Here, d_j is given by

$$d_j = \min_{\ell=1 \dots n_p} \sqrt{J^j(u_\ell, v_\ell) w_\ell}, \quad (4.14)$$

where u_ℓ, v_ℓ are the order p Vioreanu-Rokhlin nodes and w_ℓ are the corresponding weights and $\|\mathbf{V}\|$ is the operator norm of the values to interpolation matrix \mathbf{V} defined in Section 2.1. That is to say,

$$q_j = \min_q \text{ such that } \max_{\mathbf{x} \in F(\Gamma_j)} \sqrt{\sum_{n+m < p} \left| \tilde{I}_{nm,q}^j(\mathbf{x}) - \tilde{I}_{nm,q+1}^j(\mathbf{x}) \right|^2} \leq \varepsilon \frac{d_j}{\|\mathbf{V}\|}. \quad (4.15)$$

This seemingly arbitrary choice of scaling ε allows us to obtain an estimate for the relative error of the contribution of Γ_j to the layer potential in an L^2 sense, and will be clarified in the error analysis at the end of the section. It should be noted that d_j scales proportionally to a linear dimension of Γ_j (for example, like R_j).

On the other hand, if $\mathbf{x} \in \mathbb{R}^3 \setminus N_\eta(\Gamma_j)$, the kernel $G_k(\mathbf{x}, \mathbf{X}^j(u, v))$ in the integrand of $I_{nm}^j(\mathbf{x})$ is smoother than the corresponding kernel for $\mathbf{x} \in F(\Gamma_j)$. Thus, the above result also implies that for all $\mathbf{x} \in \mathbb{R}^3 \setminus \Gamma_j$,

$$\sqrt{\sum_{n+m < p} \left| \tilde{I}_{nm,q}^j(\mathbf{x}) - I_{nm}^j(\mathbf{x}) \right|^2} \leq \varepsilon \frac{d_j}{\|\mathbf{V}\|}. \quad (4.16)$$

However, for analyzing the error in evaluating the layer potential, we wish to obtain an estimate for the contribution of a discretized patch Γ_j to the layer potential $S_k[\sigma](\mathbf{x})$ denoted by $L_j(\mathbf{x})$,

$$L_j(\mathbf{x}) = \int_{u=0}^1 \int_{v=0}^{1-u} G_k(\mathbf{x}, \mathbf{X}^j(u, v)) \sigma(u, v) J^j(u, v) du dv. \quad (4.17)$$

Since the density σ is known on patch Γ_j through its samples σ_ℓ^j located at the Vioreanu-Rokhlin nodes, we have that

$$\sigma(u, v) = \sum_{n+m < p} s_{nm}^j K_{nm}(u, v), \quad \text{where} \quad s_{nm}^j = \sum_{\ell=1}^{n_p} V_{(nm),\ell} \sigma_\ell^j, \quad (4.18)$$

with \mathbf{V} the values-to-coefficients matrix in (2.8). Inserting the above expression into (4.17) we have:

$$\begin{aligned} L_j(\mathbf{x}) &= \sum_{n+m < p} s_{nm}^j \int_{u=0}^1 \int_{v=0}^{1-u} G_k(\mathbf{x}, \mathbf{X}^j(u, v)) K_{nm}(u, v) J^j(u, v) du dv \\ &= \sum_{n+m < p} s_{nm}^j I_{nm}^j(\mathbf{x}) \\ &= \sum_{\ell=1}^{n_p} \left(\sum_{n+m < p} V_{(nm),\ell} I_{nm}^j(\mathbf{x}) \right) \sigma_\ell^j. \end{aligned} \quad (4.19)$$

Suppose next that $\tilde{L}_{j,q_j}(\mathbf{x})$ denotes an approximation to $L_j(\mathbf{x})$, where each of the integrals I_{nm}^j are replaced by $\tilde{I}_{nm,q_j}^j(\mathbf{x})$. Let $w_{j,\ell} = \sqrt{J^j(u_\ell, v_\ell) w_\ell}$, and let \mathbf{W} be the diagonal matrix whose entries are $w_{j,\ell}$, $\ell = 1, 2, \dots, n_p$, and let $e_{nm} = I_{nm}^j(\mathbf{x}) - \tilde{I}_{nm,q_j}^j(\mathbf{x})$, $n+m < p$. Then for all $\mathbf{x} \in \mathbb{R}^3 \setminus N_\eta(\Gamma_j)$, it follows that

$$\begin{aligned}
 |L_j(\mathbf{x}) - \tilde{L}_{j,q_j}(\mathbf{x})| &= \mathbf{e}^T \mathbf{V} \mathbf{W}^{-1} \mathbf{W} \begin{bmatrix} \sigma_1^j \\ \vdots \\ \sigma_{n_p}^j \end{bmatrix} \\
 &\leq \frac{\varepsilon d_j}{\|\mathbf{V}\|} \cdot \|\mathbf{V}\| \cdot \|\mathbf{W}^{-1}\| \sqrt{\sum_{\ell=1}^{n_p} |\sigma_\ell^j|^2 w_{j,\ell}^2} \\
 &\leq \varepsilon \left(\|\sigma\|_{\mathbb{L}^2(\Gamma_j)} + O(\varepsilon) \right).
 \end{aligned} \tag{4.20}$$

The last inequality follows from the fact that $\|\mathbf{W}^{-1}\| = 1/d_j$, and that

$$\sum_{\ell=1}^{n_p} |\sigma_\ell^j|^2 w_{j,\ell}^2 = \sum_{\ell=1}^{n_p} |\sigma_\ell^j|^2 w_\ell J^j(u_\ell, v_\ell) = \int_{\Gamma_j} |\sigma(\mathbf{y})|^2 da(\mathbf{y}) + O(\varepsilon). \tag{4.21}$$

Remark 6. The same procedure as described above directly applies to the double layer potential with kernel $K(\mathbf{x}, \mathbf{y}) = \mathbf{n}(\mathbf{y}) \cdot \nabla_{\mathbf{y}} G_k(\mathbf{x}, \mathbf{y})$. For the normal derivative of the single layer potential, with kernel $K(\mathbf{x}, \mathbf{y}) = \mathbf{n}(\mathbf{x}) \cdot \nabla_{\mathbf{x}} G_k(\mathbf{x}, \mathbf{y})$, the procedure above can't be applied since $\mathbf{n}(\mathbf{x})$ isn't well-defined at off surface target points. However, the operator is simply the adjoint of the double layer, and therefore we use the same q_j as estimated for that case.

Remark 7. The q_j computed in equation (4.15) will of course depend on the kernel G_k , and any normalization factors. The oversampling factors were computed based on the Green's function $G_k(r) = e^{ikr}/(4\pi r)$.

Remark 8. A simple calculation shows that

$$L_{j,q_j} = \sum_{\ell=1}^{n_{q_j}} G_k(\mathbf{x}, \mathbf{X}^j(u_\ell, v_\ell)) J^j(u_\ell, v_\ell) w_\ell \tilde{\sigma}_\ell^j, \tag{4.22}$$

where u_ℓ, v_ℓ now are the order q_j Vioreanu-Rokhlin nodes on T_0 , w_ℓ the corresponding quadrature weights, and $\tilde{\sigma}_\ell^j$ is the interpolated density obtained by evaluating

$$\tilde{\sigma}_\ell^j = \sum_{n+m < p} s_{nm}^j K_{nm}(u_\ell, v_\ell), \tag{4.23}$$

where s_{nm}^j is defined in Equation (4.18). Since the same nodes $\mathbf{X}^j(u_\ell, v_\ell)$ on Γ_j can be used for all $\mathbf{x} \in \mathbb{R}^3 \setminus N_\eta(\Gamma_j)$, the far part of the layer potential evaluation can be trivially coupled to fast multipole methods.

4.3. Near field quadrature

Finally, we turn our attention to the evaluation of $S_{\text{Near}}[\sigma](\mathbf{x})$, for which the integrands are nearly-singular and we wish to develop a *high performance* variant of adaptive integration. Let us consider a patch $\Gamma_j \in T_\eta(\mathbf{x})$, and the integral $L_j(\mathbf{x})$ defined in Equation (4.17), which is also a near field integral for all $\mathbf{x} \in N_\eta(\Gamma_j)$.

It follows from Equation (4.19) that

$$\begin{aligned}
 L_j(\mathbf{x}) &= \sum_{\ell=1}^{n_p} \left(\sum_{n+m < p} V_{(nm),\ell} I_{nm}^j(\mathbf{x}) \right) \sigma_\ell^j \\
 &= \sum_{\ell=1}^{n_p} a_\ell^j(\mathbf{x}) \sigma_\ell^j.
 \end{aligned} \tag{4.24}$$

The numbers $a_\ell^j(\mathbf{x})$ are the matrix entries which map the function values σ_ℓ^j on patch Γ_j to the induced near field potential at location \mathbf{x} . If we approximate each $I_{nm}^j(\mathbf{x})$ by $\tilde{I}_{nm}^j(\mathbf{x})$ to precision $\frac{\varepsilon d_j}{\|\mathbf{V}\|}$, then using the same error analysis as in Section 4.2, we have that

$$\left| \sum_{\ell=1}^{n_p} a_\ell^j(\mathbf{x}) \sigma_\ell^j - \int_{u=0}^1 \int_{v=0}^{1-u} G_k(\mathbf{x}, \mathbf{X}^j(u, v)) \sigma(u, v) J^j(u, v) du dv \right| \leq \varepsilon \|\sigma\|_{\mathbb{L}^2(\Gamma_j)}. \tag{4.25}$$




Fig. 2. Adaptive integration grids used for the red target (left), blue target (center). The black grid (right) is the common set of triangles in both of the grids for which the function values of \mathbf{X}^j , K_{nm} , and J^j are reused.

We compute $\tilde{I}_{nm}^j(\mathbf{x})$ by adaptive integration on T_0 . That is, for precision ε , we compute the integral on T_0 using q th-order Vioreanu-Rokhlin nodes, and compare it to the integral obtained by

1. marking the midpoint of each edge of T_0 ,
2. subdividing T_0 into 4 smaller right triangles, which we will call its *descendants*, and
3. using q th order Vioreanu-Rokhlin nodes on each descendant.

The subdivision process is repeated until, for each triangle T , its contribution to the total integral agrees with the contribution computed using its descendants with an error less than $\varepsilon \cdot |T|/|T_0|$.

Done naively, this adaptive integration process dominates the cost of quadrature generation because of the large number of targets in $N_\eta(\Gamma_j)$. Note, however, that as we vary n, m , for a fixed target \mathbf{x} , the integrand of $I_{nm}^j(\mathbf{x})$ includes the same kernel values $G(\mathbf{x}, \mathbf{X}^j(u, v))$. Moreover, the adaptive grids generated for different targets have significant commonality. Thus, we can reuse the function values of $\mathbf{X}^j(u, v)$, $K_{nm}(u, v)$ and $J^j(u, v)$ if they have already been computed on any descendant triangle (see Fig. 2). The resulting scheme incurs very little increase in storage requirements; this significantly improves the overall performance.

Remark 9. Adaptive integration often results in much greater accuracy than requested. With this in mind, we set ε in the termination criterion to be somewhat larger than the precision requested. Our choice is based on extensive numerical experimentation and the full set of parameters used in our implementation is available at <https://gitlab.com/fastalgorithms/fmm3dbie>.

Remark 10. To further improve the performance of computing $I_{nm}^j(\mathbf{x})$, we make use of *two* parameters: η and $\eta_1 < \eta$. We only use adaptive integration for the nearest targets, inside $N_{\eta_1}(\Gamma_j)$. For targets $\mathbf{x} \in N_\eta(\Gamma_j) \setminus N_{\eta_1}(\Gamma_j)$, we use a single over-sampled quadrature without any adaptivity. This is slightly more expensive in terms of function evaluations, but eliminates the branching queries of adaptive quadrature and allows the use of highly optimized linear algebra libraries. From extensive numerical experiments, we have found that $\eta_1 = 1.25$ provides a significant speedup.

4.4. Error analysis

There are two sources of error in the computation of $S_k[\sigma]$: the discretization error due to discretizing the charts \mathbf{X}^j using order p Vioreanu-Rokhlin nodes, and the quadrature error due to using different quadrature rules for evaluating integrals over the discretized patches. As shown in [3], the discretization error can be bounded by

$$\left| S_k[\sigma](\mathbf{x}) - \sum_{j=1}^{N_{\text{patches}}} \int_{u=0}^1 \int_{v=0}^{1-u} G_k(\mathbf{x}, \mathbf{X}^j(u, v)) \sigma(u, v) J^j(u, v) du dv \right| = \left| S_k[\sigma](\mathbf{x}) - \sum_{j=1}^{N_{\text{patches}}} L_j(\mathbf{x}) \right| \leq Ch^p, \quad (4.26)$$

where $h = \max_j R_j$, and some domain dependent constant C . For a given $\mathbf{x} \in \Gamma_m$, our method approximates the layer potential as

$$I(\mathbf{x}) = \sum_{\ell=1}^{N(\mathbf{x})} G_k(\mathbf{x}, \mathbf{X}^m(u_{\mathbf{x},\ell}, v_{\mathbf{x},\ell})) \sigma(u_{\mathbf{x},\ell}, v_{\mathbf{x},\ell}) J^m(u_{\mathbf{x},\ell}, v_{\mathbf{x},\ell}) w_{\mathbf{x},\ell} + \sum_{\Gamma_j \in N_\eta(\mathbf{x})} \sum_{\ell=1}^{n_p} a_\ell^j(\mathbf{x}) \sigma_\ell^j + \sum_{\Gamma_j \in F_\eta(\mathbf{x})} \sum_{\ell=1}^{n_{qj}} G_k(\mathbf{x}, \mathbf{X}^j(u_{\ell,q_j}, v_{\ell,q_j})) J^j(u_{\ell,q_j}, v_{\ell,q_j}) w_{\ell,q_j} \tilde{\sigma}_\ell^j, \quad (4.27)$$

which are approximations to $\mathcal{S}_{\text{Self}}$, $\mathcal{S}_{\text{Near}}$, and \mathcal{S}_{Far} respectively. Here $(u_{\mathbf{x},\ell}, v_{\mathbf{x},\ell}), w_{\mathbf{x},\ell}$ are the auxiliary nodes on the self patch, $a_\ell^j(\mathbf{x})$ are the near quadrature corrections computed via adaptive integration, $\tilde{\sigma}_\ell^j$ is the oversampled density, and $(u_{\ell,q}, v_{\ell,q}), w_{\ell,q}$, $\ell = 1, 2, \dots, n_q$ are the order q Vioreanu-Rokhlin nodes on T_0 . Using the estimates in Equations (4.2), (4.20) and (4.25), combined with Equation (4.26), we conclude that

$$|S_k[\sigma](\mathbf{x}) - I(\mathbf{x})| \leq Ch^p + \varepsilon \|\sigma\|_{\mathbb{L}^2(\Gamma)}. \quad (4.28)$$

Remark 11. As shown in [3], the discretization error for evaluating $\mathcal{D}_k[\sigma]$ is $O(h^{p-1})$. The quadrature error analysis remains the same and we can evaluate $\mathcal{D}_k[\sigma](\mathbf{x})$ with accuracy $O(h^{p-1}) + \varepsilon \|\sigma\|_{\mathbb{L}^2(\Gamma)}$.

5. Coupling quadratures to FMMs

For a complete description of three-dimensional FMMs applied to sums of the form (1.1), we refer the reader to the original papers [27,2,28,29]. In order to understand the modifications needed for evaluating layer potentials, however, we will need to make reference to the adaptive oct-tree data structures on which the FMM is built. We briefly summarize that construction here.

5.1. Level-restricted, adaptive oct-trees

Suppose for the moment that we are given a collection of N points, contained in a cube C . We will superimpose on C a hierarchy of refinements as follows: the root of the tree is C itself and defined as *level 0*. Level $l+1$ is obtained from level l recursively by subdividing each cube at level l into eight equal parts, so long as the number of points in that cube at level l is greater than some specified parameter s . The eight cubes created in the above step are referred to as its children. Conversely, the box which was divided is referred to as their parent. When the refinement has terminated, C is covered by disjoint childless boxes at various levels of the hierarchy (depending on the local density of the given points). These childless boxes are referred to as leaf nodes. For any box D in the hierarchy, other boxes at the same level that touch D are called its *colleagues*. For simplicity, we assume that the oct-tree satisfies a standard restriction - namely, that two leaf nodes which share a boundary point must be no more than one refinement level apart. In creating the adaptive data structure as described above, it is very likely that the level-restriction criterion is not met. Fortunately, assuming that the tree constructed to this point has $O(N)$ leaf nodes and that its depth is of the order $O(\log N)$, it is straightforward to enforce the level-restriction in a second step requiring $O(N \log N)$ effort with only a modest amount of additional refinement [30].

5.2. Precomputation

To reiterate, on input, we assume we are given a surface Γ consisting of (curvilinear) triangles Γ_j ,

$$\Gamma = \bigcup_{j=1}^{N_{\text{patches}}} \Gamma_j, \quad (5.1)$$

each given to the desired order of accuracy p . Each Γ_j is then discretized using n_p points which are the images under the map $\mathbf{X}^j : T_0 \rightarrow \mathbb{R}^3$ of the Vioreanu-Rokhlin nodes on the standard simplex T_0 . We will refer to these as the *discretization nodes*, on which we assume that samples of the density σ are known. The total number of such points is $N = N_{\text{patches}} \times n_p$. As above, we let \mathbf{c}_j denote the centroid of the j th patch and R_j the radius of the smallest sphere centered at \mathbf{c}_j that contains Γ_j . We assume there are N_T targets, which could be either the discretization nodes themselves, a collection of off-surface points, or both.

In coupling the FMM to local quadratures, we need to determine, for each surface patch, which targets are in its near field and what order Vioreanu-Rokhlin quadratures are needed for the far field computation. Both are controlled by the parameter η , as discussed in Section 4.1. The default value for η is 2.75, 2, or 1.25 depending on whether the desired order of accuracy is $p \leq 4$, $4 < p \leq 8$ or $p > 8$, respectively. The first step is to build an adaptive oct-tree based on the patch centroids $\{\mathbf{c}_j\}$ and the target locations, with one minor modification. That modification is to prevent triangle centroids associated with large triangles from propagating to fine levels during the tree construction. For this, suppose \mathbf{c}_j is in some box, denoted $D(\mathbf{c}_j)$ at level l , and let d denote the linear dimension of $D(\mathbf{c}_j)$. If $2\eta R_j > d$, then we leave the centroid associated with $D(\mathbf{c}_j)$, while allowing smaller triangles and/or targets to be associated with the children. We will say that Γ_j is *tethered* at level l .

The near field for each patch is now easy to determine. For each triangle Γ_j , let $D(\mathbf{c}_j)$ denote the box to which the triangle centroid is associated - either a leaf node or the box at a coarser level l if is tethered there. Clearly, if $D(\mathbf{c}_j)$ is not a leaf box, then the near field region $B_{\eta R_j}(\mathbf{c}_j)$ is contained within $D(\mathbf{c}_j)$ and its colleagues. If $D(\mathbf{c}_j)$ is a leaf box, then the near field region $B_{\eta R_j}(\mathbf{c}_j)$ is contained within the union of $D(\mathbf{c}_j)$, its colleagues, and leaf boxes which are larger in size than $D(\mathbf{c}_j)$ and which share a boundary with $D(\mathbf{c}_j)$. Scanning those colleagues, all targets \mathbf{x} that do not lie on Γ_j itself and satisfy the criterion

$$|\mathbf{x} - \mathbf{c}_j| < \eta R_j$$

are assigned to the *near field list* for Γ_j . One can then compute the near field quadratures using the method of Section 4.3 for each point in the target list. This requires storing a matrix of dimension $N_{near}(j) \times n_p$, where $N_{near}(j)$ is the size of the target list. We will denote this matrix by \mathcal{N}_j .

Assuming one wishes to evaluate the layer potential on surface, we also need to compute the self interactions for each triangle using the generalized Gaussian quadrature scheme of [8,9], as described in Section 4. This requires storing an $n_p \times n_p$ matrix for each patch, which we will denote by \mathcal{S}_j .

Once the near field work has been carried out, the far field quadrature order q_j is determined, as described in Section 4.2. One can then interpolate from the n_p discretization nodes on Γ_j to the n_{q_j} quadrature nodes on Γ_j using the Koornwinder basis for interpolation. We will denote by N_{over} the total number of oversampled points used: $N_{over} = \sum_{i=1}^{N_{patches}} n_{q_i}$.

5.3. Fast evaluation of layer potentials

The simplest FMM-based scheme for evaluating a layer potential is to call the point-based FMM in the form (1.1), with N_{over} sources and N_T targets. For every target, if it is in the near field of patch Γ_j , one subtracts the contribution made in the naive, point-based FMM calculation from the n_q oversampled points on that patch. The potential at the target can then be incremented by the appropriate, near field quadrature-corrected interactions, using the stored matrix \mathcal{N}_j . If the target is on surface (one of the discretization nodes on Γ_j itself), the correct self interaction is obtained from the precomputed matrix \mathcal{S}_j .

We refer to the algorithm above as the *subtract-and-add* method. It has the drawback that it could suffer from catastrophic cancellation for dense discretizations with highly adaptive oct-trees since the near field point contributions within the naive FMM call are spurious and could be much greater in magnitude than the correct contributions. (In practice, we have not detected any such loss of accuracy, at least for single or double layer potentials.)

For readers familiar with the FMM, it is clear that one could avoid the need to compute and then subtract spurious contributions, by disabling the direct (near neighbor) interaction step in the FMM. When looping over the leaf nodes, for each source-target pair, one can first determine whether the source is on a patch for which the target is in the far field. If it is, carry out the direct interaction. If it is in the near field, omit the direct interaction. When the FMM step is completed, the subsequent processing takes place as before, but there is no need to subtract any spurious contributions.

It is, perhaps, surprising that the *subtract-and-add* method is faster in our current implementation, even though more flops are executed. This is largely because of the logical overhead and the bottlenecks introduced in loop unrolling and other compiler-level code optimizations.

Remark 12. The adaptive oct-tree used in the point FMM is different from the one used for determining the near field of the patches. The latter is constructed based on centroid and target locations, while the point FMM oct-tree is constructed based on *oversampled* source and target locations. Thus, different termination criteria can be chosen for the construction of these oct-trees in order to optimize the performance of the separate tasks.

Since the additional processing required for evaluating layer potentials is decoupled from the algorithm used for accelerating the far field interactions, one could use any fast hierarchical algorithm like the FMM, an FFT-based scheme like fast Ewald summation, or a multigrid-type PDE solver.

6. Numerical examples

In this section, we illustrate the performance of our approach. For Examples 6.1, 6.2, and 6.3, we consider a twisted torus as the geometry (typical of stellarator design in plasma physics applications). The boundary Γ is parameterized by $\mathbf{X} : [0, 2\pi]^2 \rightarrow \Gamma$ with

$$\mathbf{X}(u, v) = \sum_{i=-1}^2 \sum_{j=-1}^1 \delta_{i,j} \begin{bmatrix} \cos v \cos((1-i)u + jv) \\ \sin v \cos((1-i)u + jv) \\ \sin((1-i)u + jv) \end{bmatrix}, \quad (6.1)$$

where the non-zero coefficients are $\delta_{-1,-1} = 0.17$, $\delta_{-1,0} = 0.11$, $\delta_{0,0} = 1$, $\delta_{1,0} = 4.5$, $\delta_{2,0} = -0.25$, $\delta_{0,1} = 0.07$, and $\delta_{2,1} = -0.45$. (See Fig. 3.)

The code was implemented in Fortran and compiled using the GNU Fortran 10.2.0 compiler. We use the point-based FMMs from the FMM3D package (<https://github.com/flatironinstitute/FMM3D>). All CPU timings in these examples were obtained on a laptop using a single core of an Intel i5 2.3 GHz processor.

We use the following metrics to demonstrate the performance of our approach. As above, for discretization order p , we let $n_p = p(p+1)/2$, and we let q_j denote the far-field quadrature order for Γ_j . The user-specified precision is denoted by ε . Recall that the total number of discretization points on the boundary is denoted by $N = N_{patches} \cdot n_p$ and that the total




Fig. 3. The boundary of a stellarator like geometry. The surface is colored using its z-coordinate.

Table 1

Memory requirements, oversampling requirements, precomputation speed and layer potential evaluation speed as a function of order of accuracy p and precision ε .

(a) Oversampling parameter α					(b) Memory requirements m per point				
$p \backslash \varepsilon$	$5 \cdot 10^{-3}$	$5 \cdot 10^{-4}$	$5 \cdot 10^{-7}$	$5 \cdot 10^{-10}$	$p \backslash \varepsilon$	$5 \cdot 10^{-3}$	$5 \cdot 10^{-4}$	$5 \cdot 10^{-7}$	$5 \cdot 10^{-10}$
2	1	1	3.52	9.35	2	127	127	127	127
3	0.5	1	2.5	5.19	3	258	258	258	258
4	0.6	0.824	2.8	4.5	4	209	209	209	209
6	0.476	0.714	1.84	2.94	6	440	440	440	440
8	0.777	0.923	2.24	3.9	8	286	286	286	286

(c) Precomputation speed S_{init}					(d) Layer potential speed S_{LP}				
$p \backslash \varepsilon$	$5 \cdot 10^{-3}$	$5 \cdot 10^{-4}$	$5 \cdot 10^{-7}$	$5 \cdot 10^{-10}$	$p \backslash \varepsilon$	$5 \cdot 10^{-3}$	$5 \cdot 10^{-4}$	$5 \cdot 10^{-7}$	$5 \cdot 10^{-10}$
2	8460	6480	2970	1350	2	30100	24500	6480	1450
3	8130	5880	2460	1150	3	35800	18600	6870	2170
4	8230	6440	2820	1360	4	27500	20200	5320	2350
6	2890	2380	1250	734	6	24300	19900	6290	3390
8	2080	1750	945	581	8	20500	16000	5960	2140

number of oversampled nodes is denoted by $N_{\text{over}} = \sum_{j=1}^{N_{\text{patches}}} n_{q_j}$. We define the oversampling parameter by $\alpha = N_{\text{over}}/N$. The memory requirements per discretization node for storing all interactions in S_{Near} are given by

$$m = \frac{n_p \left(\sum_{j=1}^{N_{\text{patches}}} N_{\text{near}}(j) + n_p \right)}{N}. \quad (6.2)$$

This accounts for both off-surface targets and on surface evaluation.

Let T_{init} denote the time required to precompute all near field quadrature corrections and let T_{LP} denote the time for evaluating the layer potential given the precomputed near field quadratures. Then, the quantities $S_{\text{init}} = N/T_{\text{init}}$, and $S_{\text{LP}} = N/T_{\text{LP}}$, are the speeds of the corresponding steps, measured in points processed per second.

One feature of the surface triangulation that has some influence on speed is the aspect ratio of the patches. Letting σ_1, σ_2 be the eigenvalues of the first fundamental form of Γ_j , we define its aspect ratio by

$$a_j = \sqrt{\frac{\int_{\Gamma_j} \left(\frac{\sigma_1}{\sigma_2} \right)^2 da(\mathbf{y})}{\int_{\Gamma_j} da(\mathbf{y})}}. \quad (6.3)$$


We let $a_{\text{max}} = \max_j a_j$ and $a_{\text{avg}} = \sum_j a_j / N_{\text{patches}}$, the maximum aspect ratio and the average aspect ratio over all triangles, respectively.

6.1. Memory and oversampling requirements


To illustrate the performance of our method as a function of the order of accuracy p and the requested precision ε , we consider the evaluation of the single layer potential $\mathcal{S}[\sigma]$ with frequency $k = 1$ on the stellarator geometry discretized with $N_{\text{patches}} = 2400$ (the diameter of the stellarator with $k = 1$ is approximately 1.7 wavelengths). In Table 1, we tabulate the memory requirements per point m , the oversampling α , and the speeds S_{init} and S_{LP} , as we vary p and ε . The scheme

Table 2
Performance as a function of average aspect ratio a_{avg} .

a_{avg}	1.61	2.33	4.66	9.32	14
N_{patches}	2400	2304	2450	2304	2400
α	2.8	2.8	2.8	2.8	2.8
m	209	291	585	1220	1860
S_{init}	2890	2160	1100	519	341
S_{quad}	5390	4940	3920	2890	2270



(a) Relative L^2 error in Green’s identity, denoted by ε_g .



(b) Relative L^∞ error in solution to integral equation, denoted by ε_a .

Fig. 4. Relative errors of layer potential evaluations. In both figures, the dashed colored lines are reference curves for h^{p-1} with the corresponding p . The dashed black line is a reference line for the specified tolerance ε .

behaves as expected: for fixed p , as $\varepsilon \rightarrow 0$, α increases while S_{init} and S_{LP} decrease. The oversampling parameter α depends on both p and ε , as discussed in Section 4.2.

6.2. Effect of aspect ratio

To investigate the effect of triangle quality on the performance of our method, we vary the average aspect ratio of the discretization. The task at hand is again to compute $\mathcal{S}[\sigma]$ with $k = 1$ on the stellarator, while varying the triangulation without a significant change in the total number of patches. In Table 2, we tabulate a_{avg} , N_{patches} , α , m , S_{init} , and S_{LP} for $p = 4$ and $\varepsilon = 5 \cdot 10^{-7}$. We note that (except for the oversampling parameter), the performance of the approach deteriorates as the average aspect ratio of the discretization is increased, especially in the precomputation phase.

6.3. Order of convergence

To demonstrate the accuracy of our approach, we consider two tests. First, we verify Green’s identity along the surface:

$$\frac{u}{2} = \mathcal{S}_k \left[\frac{\partial u}{\partial n} \right] - \mathcal{D}_k[u], \quad (6.4)$$

where u is the solution to the Helmholtz equation in the interior of the domain Ω generated by a point source located in the exterior. The second test is to use the combined field representation (3.2) to solve the Dirichlet problem for an unknown density σ . With the right-hand side in the corresponding integral equation (3.5) taken to be a known solution u , σ satisfies

$$\frac{\sigma}{2} + \mathcal{D}_k[\sigma] - ik\mathcal{S}_k[\sigma] = u, \quad \text{on } \Gamma. \quad (6.5)$$

We can then check that u is correctly reproduced at any point in the exterior. For both of these tests, we discretize the stellarator using $N_{\text{patches}} = 600, 2400$, and 9600 , and compute the layer potentials with a tolerance of $\varepsilon = 5 \times 10^{-7}$ for both the quadratures and the FMM.

In Fig. 4, we plot the relative L^2 error in Green’s identity ε_g (left), and the relative L^∞ error at a point in the interior ε_a (right) as we vary the order of discretization. Note that in both tests, the errors decrease at the rate h^{p-1} until the tolerance ε is reached. This is consistent with the analysis in [3].

6.4. Large-scale examples

We demonstrate the performance of our solver on several large-scale problems. We first solve for the electrostatic field induced by an interdigitated capacitor, followed by Dirichlet and Neumann boundary value problems governed by the Helmholtz equation in the exterior of an aircraft. Our last example involves scattering in a medium with multiple sound speeds, modeled after a Fresnel lens. The results in this section were obtained using an Intel Xeon Gold 6128 Desktop with 24 cores.

Interdigitated capacitor

A challenging problem in electrostatics is the calculation of the capacitance of a configuration of two perfect compact conductors with complicated contours, which may also be close to touching. (See Fig. 5.) The capacitance is defined as the ratio $C = Q/V$, where V is the potential difference between the conductors, Q is the total charge held on one conductor and $-Q$ is the total charge held on the other. In simulations, C can be computed in two ways. First, one can solve the Dirichlet problem for the electrostatic potential u , with $u = 0$ on one conductor and $u = 1$ on the other. From the computed solution, the total charge can be obtained via the integral [31]

$$Q = \int_{\Gamma} \frac{\partial u}{\partial n}(\mathbf{x}') da(\mathbf{x}'). \quad (6.6)$$

The capacitance is then $C = Q/1 = Q$.

A second (equivalent) approach, which we will take here, is to place a net charge $q_1 = -1$ on one conductor Ω_1 with boundary Γ_1 and a net charge of $q_2 = 1$ on the other conductor Ω_2 with boundary Γ_2 . One can then determine the corresponding potential difference by solving the following boundary value problem for the potential u in the domain E exterior to Ω_1 and Ω_2 , i.e. the domain $E = \mathbb{R}^3 \setminus (\Omega_1 \cup \Omega_2)$:

$$\begin{aligned} \Delta u &= 0, & \mathbf{x} \in E, \\ u &= V_i, & \mathbf{x} \text{ on } \Gamma_i, \\ - \int_{\Gamma_j} \frac{\partial u}{\partial n} da &= q_i, \\ u &\rightarrow 0 & \text{as } |\mathbf{x}| \rightarrow \infty. \end{aligned} \quad (6.7)$$

Here, the constants V_1, V_2 are unknowns as well as the potential u . The *elastance* of the system is then given by $P = (V_2 - V_1)/(q_2 - q_1) = (V_2 - V_1)/2$. It is the inverse of the corresponding capacitance $C = 1/P = 2/(V_2 - V_1)$. This formulation (which can involve more than two conductors) is generally referred to as the *elastance* problem (see [32] and the references therein). PDEs of this type where Dirichlet data is specified up to an unknown constant are sometimes called modified Dirichlet problems [33].

We represent the solution u using the combined field representation,

$$u = S_0[\rho] + D_0[\rho], \quad (6.8)$$

where ρ is an unknown density. Imposing the boundary conditions on Γ_i , ρ must satisfy the integral equation

$$\begin{aligned} \rho/2 + D_0[\rho] + S_0[\rho] &= V_i, & \mathbf{x} \text{ on } \Gamma_i, \\ \int_{\Gamma_i} \rho &= -q_i, & i = 1, 2. \end{aligned} \quad (6.9)$$

We discretize the surface Γ with $N_{\text{patches}} = 29,888$ and $p = 4$, and then solve the resulting linear system of size $N = 298,882$ using GMRES. We set the quadrature tolerance $\varepsilon = 5 \times 10^{-7}$. For this setup, $T_{\text{init}} = 6.7$ s, $\alpha = 2.98$, $m = 154.2$. GMRES converged to a relative residual of 5×10^{-7} in 25 iterations, and the solution was obtained in 128.4 s. The reference capacitance for the system was computed by refining each patch until it had converged to 5 significant digits, given by 2237.1. The relative error in the computed capacitance was 2.2×10^{-4} . In Fig. 5, we plot the computational mesh and the solution ρ on the surface of the conductors.

Scattering from an airplane (sound-soft)

In this section we demonstrate the performance of our method on a moderate frequency acoustic scattering problem. The model airplane is 49.3 wavelengths long, with a wingspan of 49.2 wavelengths and a vertical height of 13.7 wavelengths, which we assume has a sound-soft boundary, satisfying Dirichlet boundary conditions (see section 3). The plane also has several multiscale features: 2 antennae on the top of the fuselage, and 1 *control unit* on the bottom of the fuselage. (See




Fig. 5. (left) A 4th-order computational mesh of the boundary, and (right) the solution ρ to Equation (6.9).

Fig. 7.) The plane is discretized with $N_{\text{patches}} = 125,344$, and $p = 4$ resulting in $N = 1,253,440$ discretization points. The ratio of the largest to the smallest patch size, measured by the enclosing sphere radius R_j in Equation (4.6), is 483.9. In Fig. 6, we plot a histogram of the patch sizes R_j , and the aspect ratios of the patches on the plane. The worst case patch has an aspect ratio of 35, but only 322 patches out of the total 125,344 have an aspect ratio of greater than 10.

In order to have an analytic reference solution, we assume the Dirichlet boundary data $u|_{\Gamma}$ for the governing exterior Helmholtz equation is generated using a collection of 123 interior sources, 19 of which are in the tail. Using a combined field representation

$$u = \mathcal{D}_k[\sigma] - ik\mathcal{S}_k[\sigma],$$

imposing the Dirichlet condition yields the combined field integral equation for the unknown density σ :

$$\frac{\sigma}{2} + \mathcal{D}_k[\sigma] - ik\mathcal{S}_k[\sigma] = u|_{\Gamma}. \quad (6.10)$$

For a quadrature tolerance of 5×10^{-7} , $T_{\text{init}} = 75.86$ s, the oversampling factor $\alpha = 4.51$, and the memory cost per discretization point is $m = 149.6$. GMRES converged to a relative residual of 5×10^{-7} in 59 iterations, and the solution was obtained in 4,694 s. We plot the solution on a 301×301 lattice of targets on a slice which cuts through the wing edge, whose normal is given by $(0, 0, 1)$. For targets in the interior of the airplane, we set the error to 1×10^{-6} since the computed solution there is not meaningful. The layer potential evaluation at all targets required only $T_{\text{init}} + T_{\text{quad}} = 104.37$ s. In Fig. 7, we plot the density σ on the airplane surface and the relative error on the slice with 301×301 targets. The maximum relative error at all targets is 5×10^{-4} .

Scattering from an airplane (sound-hard)

In this section we solve the Helmholtz equation in the exterior of the plane, assuming Neumann boundary conditions instead. These arise in the modeling of sound-hard scatterers [26]. We will use the following regularized combined field integral representation of [34] for the solution:

$$u = \mathcal{S}_k[\sigma] + i\alpha\mathcal{D}_k[\mathcal{S}_{i|k|}[\sigma]]. \quad (6.11)$$

Applying the boundary condition $\partial u/\partial n = g$ along Γ leads to the second-kind integral equation:

$$-\frac{\sigma}{2} + \mathcal{S}'_k[\sigma] + i\alpha\mathcal{D}'_k[\mathcal{S}_{i|k|}[\sigma]] = g. \quad (6.12)$$

Using a well-known Calderón identity for the operator $\mathcal{D}'_{i|k|}\mathcal{S}_{i|k|}$ [35], this equation can be re-written so as to avoid the application of the hypersingular operator \mathcal{D}'_k :

$$-\left(\frac{2+i\alpha}{4}\right)\sigma + \mathcal{S}'_k[\sigma] + i\alpha\left(\mathcal{D}'_k - \mathcal{D}'_{i|k|}\right)[\mathcal{S}_{i|k|}[\sigma]] + i\alpha\mathcal{S}'_{i|k|}{}^2[\sigma] = g. \quad (6.13)$$

A number of different possibilities for the regularizing operator are available depending on the frequency range (see [36, 37]). Examining the above integral equation, it is clear that a total of four separate FMM calls and four local quadrature corrections will be needed.

In order to have an analytic reference solution, we generate the Neumann boundary data $u|_{\Gamma}$ for the governing exterior Helmholtz equation by using the same 123 interior sources as in the Dirichlet problem. For a quadrature tolerance of 5×10^{-7} , $T_{\text{init}} = 387.4$ s, the oversampling factor $\alpha = 4.51$, and the memory cost per discretization point is $m = 598.3$.




Fig. 6. Histogram of size of patches R_j (left), and aspect ratio a_j right.




Fig. 7. The solution σ to Equation (6.10), and the relative error in the solution at a grid of 301×301 targets on a horizontal slice intersecting the wing edge. Zoomed in views of one antenna (top right) and the control unit (bottom left) indicate the extent of the fine multiscale features.

GMRES converged to a relative residual of 5×10^{-7} in 35 iterations, and the solution was obtained in 7617 s. We plot the solution on the 301×301 lattice of targets used for the Dirichlet problem. As before, for targets in the interior of the airplane, we set the error to 1×10^{-6} since the computed solution there is not meaningful. The layer potential evaluation at all targets required only $T_{\text{init}} + T_{\text{quad}} = 159.6$ s. In Fig. 8, we plot the induced density σ on the airplane surface and the relative error on the slice with 301×301 targets. The maximum relative error at all targets is 5×10^{-4} .

Scattering through a Fresnel lens (multiple sound speeds)

In this section we solve the Helmholtz transmission problem, i.e. scattering through media with various sound speeds, in a Fresnel lens geometry (see Fig. 9). The Helmholtz parameter for the interior region (the lens) is $k = \omega\sqrt{\epsilon\mu}$ and for the exterior region (free space) is $k_0 = \omega\sqrt{\epsilon_0\mu_0}$. We assume that the known incoming field u^{inc} exists only in the exterior, so




Fig. 8. The solution σ to integral equation (6.13), and the relative error in the solution to the PDE at a grid of 301×301 targets on a horizontal slice intersecting the wing edge. Zoomed in views of one antenna (top right) and the control unit (bottom left) indicate the extent of the fine multiscale features.

that the total field in the exterior is given by $u_t = u_0 + u^{\text{inc}}$ where u_0 is the scattered field. In the interior, the total field is merely $u_t = u$, with u the scattered field.

In the piecewise constant sound speed setup, we enforce the following transmission conditions across interfaces:

$$\begin{aligned} u_0 - u &= -u^{\text{inc}}|_{\Gamma}, \\ \frac{1}{\epsilon_0} \frac{\partial u_0}{\partial n} - \frac{1}{\epsilon_1} \frac{\partial u}{\partial n} &= -\frac{1}{\epsilon_0} \frac{\partial u^{\text{inc}}}{\partial n} \Big|_{\Gamma}. \end{aligned} \quad (6.14)$$

The scattered field in the exterior region, u_0 , and in the interior region, u , are represented as, respectively [26]:

$$\begin{aligned} u_0 &= \epsilon_0 \mathcal{D}_{k_0}[\rho] + \epsilon_0^2 \mathcal{S}_{k_0}[\sigma], \\ u &= \epsilon \mathcal{D}_k[\rho] + \epsilon^2 \mathcal{S}_k[\sigma]. \end{aligned} \quad (6.15)$$

This leads to the system of boundary integral equations along Γ

$$\begin{aligned} \left(\frac{\epsilon_0 + \epsilon}{2}\right) \rho + (\epsilon_0 \mathcal{D}_{k_0} - \epsilon \mathcal{D}_k)[\rho] + (\epsilon_0^2 \mathcal{S}_{k_0} - \epsilon^2 \mathcal{S}_k)[\sigma] &= -u^{\text{inc}}, \\ \left(\frac{\epsilon_0 + \epsilon}{2}\right) \sigma - (\mathcal{D}'_{k_0} - \mathcal{D}'_k)[\rho] - (\epsilon_0 \mathcal{S}'_{k_0} - \epsilon \mathcal{S}'_k)[\sigma] &= -\frac{1}{\epsilon_0} \frac{\partial u^{\text{inc}}}{\partial n}. \end{aligned} \quad (6.16)$$

In the following example, the Fresnel lens has $\epsilon = 2$, $\mu = 1$ and as usual, in free-space $\epsilon_0 = \mu_0 = 1$. The angular frequency is set to be $\omega = 1 + \sqrt{2}$, and the annular step size in the lens equals 1. Relative to the exterior wavenumber, the Fresnel lens is 19.11 wavelengths in diameter and has a height of 0.84 wavelengths. The geometry was designed in GiD [38], and a 4th-order curvilinear mesh was constructed using the method described in [39]. The mesh consists of 62,792 curvilinear triangles, each discretized to 5th order yielding a total of 941,880 discretization points. See Fig. 9.

We solve the transmission problem in response to an incoming plane wave $u^{\text{inc}} = e^{ik_0 z}$. For a quadrature tolerance of 5×10^{-7} , $T_{\text{init}} = 238.2$ s, the oversampling factor $\alpha = 3.07$, and the memory cost per discretization point is $m = 762.0$. GMRES converged to a relative residual of 5×10^{-7} in 294 iterations, and the solution was obtained in 23029 s. In Fig. 9 we plot the absolute value of the total field $|u_t| = |u_0 + u^{\text{inc}}|$ on a 1000×1000 lattice of targets in the yz plane in the exterior region. The layer potential evaluation at all targets required only $T_{\text{init}} + T_{\text{quad}} = 81.1$ s. We also plot the real part of the density ρ on the surface of the lens.

The accuracy of the solution is estimated by solving a transmission problem whose boundary data is computed using known solutions to the Helmholtz equation in the interior and exterior (the interior Helmholtz solution is the potential due




Fig. 9. On the left, we illustrate the triangulation of the lens surface. On the right is a plot of the real part of the solution ρ to Equation (6.16) on the surface of the lens and the absolute value of the total field above the lens where the central focusing of the beam is clearly visible.

to a point source in the exterior and the exterior Helmholtz solution is the potential due to a point source in the interior). The maximum relative error for the computed solution at the same target grid as above is 8.4×10^{-5} .

7. Conclusions

In this paper, we have presented a robust, high-order accurate method for the evaluation of layer potentials on surfaces in three dimensions. While our examples have focused on the Laplace and Helmholtz equations, the underlying methodology extends naturally to other problems in mathematical physics where the governing Green's function is singular – thereby requiring specialized quadrature schemes – but compatible with fast multipole acceleration in the far field.

To determine the highest performance scheme, we implemented generalized Gaussian quadrature [8–11], QBX [21,16–18], and coordinate transformation schemes similar to [12,14,19]. After various code optimizations (at least for surfaces defined as collections of curved, triangular patches), we found that generalized Gaussian quadrature with careful reuse of precomputed, hierarchical, adaptive interpolation tables was most efficient, as illustrated in the preceding section. It may be that an even better local quadrature scheme emerges in the future (in particular, extensions of the idea presented in [40] look quite promising). In that case, as discussed in Section 5, coupling to the FMM will be fundamentally unchanged.

A useful feature of locally corrected quadrature rules (of any kind) is that the procedure is trivial to parallelize by assigning a different patch to each computational thread. Thus, acceleration on multi-core or high performance platforms is straightforward.

Finally, although we limited ourselves here to evaluating layer potentials and solving integral equations iteratively, we note that our quadrature generation scheme and oct-tree data structure are compatible for use with fast direct solvers [10,41–48]. These require access to small blocks of the system matrix in an effort to find a compressed representation of the inverse.

There are still a number of open questions that remain to be addressed, including the development of rules for surfaces with edges and corners and the coupling of layer potential codes with volume integral codes to solve inhomogeneous or variable-coefficient partial differential equations using integral equation methods. These are all ongoing areas of research.

CRedit authorship contribution statement

All authors worked equally on all aspects of the paper.

Declaration of competing interest

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

Acknowledgements

We would like to thank Alex Barnett and Lise-Marie Imbert-Gérard for many useful discussions, and Jim Bremer and Zydrunas Gimbutas for sharing several useful quadrature codes. We also gratefully acknowledge the support of the NVIDIA

Corporation with the donation of a Quadro P6000, used for some of the visualizations presented in this research. L. Greengard was supported in part by the Office of Naval Research under award number #N00014-18-1-2307. M. O'Neil was supported in part by the Office of Naval Research under award numbers #N00014-17-1-2059, #N00014-17-1-2451, and #N00014-18-1-2307. F. Vico was supported in part by the Office of Naval Research under award number #N00014-18-2307, the Generalitat Valenciana under award number AICO/2019/018, and by the Spanish Ministry of Science and Innovation (Ministerio Ciencia e Innovación) under award number PID2019-107885GB-C32. We would also like to thank the anonymous referees for many helpful comments that led to a much-improved manuscript.

References

- [1] L. Greengard, J. Huang, V. Rokhlin, S. Wandzura, Accelerating fast multipole methods for the Helmholtz equation at low frequencies, *IEEE Comput. Sci. Eng.* 5 (3) (1998) 32–38.
- [2] H. Cheng, W.Y. Crutchfield, Z. Gimbutas, L. Greengard, J.F. Ethridge, J. Huang, V. Rokhlin, N. Yarvin, J. Zhao, A wideband fast multipole method for the Helmholtz equation in three dimensions, *J. Comput. Phys.* 216 (2006) 300–325, <https://doi.org/10.1016/j.jcp.2005.12.001>.
- [3] K.E. Atkinson, D. Chien, Piecewise polynomial collocation for boundary integral equations, *SIAM J. Sci. Comput.* 16 (1995) 651–681.
- [4] K.E. Atkinson, *The Numerical Solution of Integral Equations of the Second Kind*, Cambridge University Press, New York, NY, 1997.
- [5] M.J. Aftosis, M.J. Berger, J.E. Melton, Robust and efficient Cartesian mesh generation for component-based geometry, *AIAA J.* 36 (1998) 952–960.
- [6] J. Bell, M. Berger, J. Saltzman, M. Welcome, Three-dimensional adaptive mesh refinement for hyperbolic conservation laws, *SIAM J. Sci. Comput.* 15 (1994) 127–138.
- [7] H. Johansson, P. Colella, A Cartesian grid embedded boundary method for Poisson's equation on irregular domains, *J. Comput. Phys.* 147 (2) (1998) 60–85.
- [8] J. Bremer, Z. Gimbutas, A Nyström method for weakly singular integral operators on surfaces, *J. Comput. Phys.* 231 (2012) 4885–4903.
- [9] J. Bremer, Z. Gimbutas, On the numerical evaluation of singular integrals of scattering theory, *J. Comput. Phys.* 251 (2013) 327–343.
- [10] J. Bremer, A. Gillman, P.-G. Martinsson, A high-order accelerated direct solver for integral equations on curved surfaces, *BIT Numer. Math.* 55 (2015) 367–397, <https://doi.org/10.1007/s10543-014-0508-y>.
- [11] J. Bremer, Z. Gimbutas, V. Rokhlin, A nonlinear optimization procedure for generalized Gaussian quadratures, *SIAM J. Sci. Comput.* 32 (4) (2010) 1761–1788.
- [12] O.P. Bruno, L.A. Kunyansky, A fast, high-order algorithm for the solution of surface scattering problems: basic implementation, tests, and applications, *J. Comput. Phys.* 169 (1) (2001) 80–110.
- [13] O. Bruno, E. Garza, A Chebyshev-based rectangular-polar integral solver for scattering by geometries described by non-overlapping patches, *J. Comput. Phys.* (2020) 109740, <https://doi.org/10.1016/j.jcp.2020.109740>.
- [14] D. Malhotra, A. Cerfon, L.-M. Imbert-Gérard, M. O'Neil, Taylor states in stellarators: a fast high-order boundary integral solver, *J. Comput. Phys.* 397 (2019) 108791.
- [15] S. Erichsen, S.A. Sauter, Efficient automatic quadrature in 3-d Galerkin BEM, *Comput. Methods Appl. Mech. Eng.* 157 (1998) 215–224.
- [16] M. Siegel, A.-K. Tornberg, A local target specific quadrature by expansion method for evaluation of layer potentials in 3d, *J. Comput. Phys.* 364 (2018) 365–392.
- [17] M. Wala, A. Klöckner, A fast algorithm for quadrature by expansion in three dimensions, *J. Comput. Phys.* 388 (2018) 655–689.
- [18] M. Wala, A. Klöckner, Optimization of fast algorithms for global quadrature by expansion using target-specific expansions, *J. Comput. Phys.* 403 (2020).
- [19] L. Ying, G. Biros, D. Zorin, A high-order 3D boundary integral equation solver for elliptic PDEs in smooth domains, *J. Comput. Phys.* 219 (1) (2006) 247–275.
- [20] M.G. Duffy, Quadrature over a pyramid or cube of integrands with a singularity at a vertex, *SIAM J. Numer. Anal.* 19 (1982) 1260–1262.
- [21] A. Klöckner, A. Barnett, L. Greengard, M. O'Neil, Quadrature by expansion: a new method for the evaluation of layer potentials, *J. Comput. Phys.* 252 (2013) 332–349.
- [22] T. Koornwinder, Two-variable analogues of the classical orthogonal polynomials, in: *Theory and Application of Special Functions (Proc. Advanced Sem., Math. Res., Center, Univ. Wisconsin, Madison, Wis., 1975)*, Academic Press, New York, 1975, pp. 435–495.
- [23] L.N. Trefethen, *Approximation Theory and Approximation Practice*, SIAM, Philadelphia, PA, 2013.
- [24] H. Xiao, Z. Gimbutas, A numerical algorithm for the construction of efficient quadrature rules in two and higher dimensions, *Comput. Math. Appl.* 59 (2) (2010) 663–676.
- [25] B. Vioreanu, V. Rokhlin, Spectra of multiplication operators as a numerical tool, *SIAM J. Sci. Comput.* 36 (2014) A267–A288.
- [26] D. Colton, R. Kress, *Integral Equation Methods in Scattering Theory*, John Wiley & Sons, Inc., 1983.
- [27] H. Cheng, L. Greengard, V. Rokhlin, A fast adaptive multipole algorithm in three dimensions, *J. Comput. Phys.* 155 (2) (1999) 468–498.
- [28] L. Greengard, V. Rokhlin, A new version of the fast multipole method for the Laplace equation in three dimensions, *Acta Numer.* 6 (1997) 229–269.
- [29] L. Greengard, J. Huang, A new version of the fast multipole method for screened Coulomb interactions in three dimensions, *J. Comput. Phys.* 180 (2002) 642–658.
- [30] M. de Berg, O. Cheong, M. Kreveld, M. Overmars, *Computational Geometry: Algorithms and Applications*, Springer, 2008.
- [31] J.D. Jackson, *Classical Electrodynamics*, 3rd edition, Wiley, New York, NY, 1999.
- [32] M. Rachh, L. Greengard, Integral equation methods for elastance and mobility problems in two dimensions, *SIAM J. Numer. Anal.* 54 (5) (2016) 2889–2909.
- [33] S.G. Mikhlin, *Integral Equations: And Their Applications to Certain Problems in Mechanics, Mathematical Physics and Technology*, Elsevier, 2014.
- [34] O. Bruno, T. Elling, C. Turc, Regularized integral equations and fast high-order solvers for sound-hard acoustic scattering problems, *Int. J. Numer. Methods Eng.* 91 (2012) 1045–1072.
- [35] J.-C. Nédélec, *Acoustic and Electromagnetic Equations*, Springer, New York, NY, 2001.
- [36] D. Colton, R. Kress, *Inverse Acoustic and Electromagnetic Scattering Theory*, Springer, New York, NY, 2012.
- [37] F. Vico, L. Greengard, Z. Gimbutas, Boundary integral equation analysis on the sphere, *Numer. Math.* 128 (2014) 463–487.
- [38] International Center for Numerical Methods in Engineering (CIMNE), GiD, the Personal Pre- and Post-processor, gidhome.com, 2020.
- [39] F. Vico, L. Greengard, M. O'Neil, M. Rachh, A fast boundary integral method for high-order multiscale mesh generation, *SIAM J. Sci. Comput.* 42 (2) (2020) A1380–A1401.
- [40] B. Wu, P.-G. Martinsson, Corrected trapezoidal rules for boundary integral equations in three dimensions, [arXiv:2007.02512 \[math.NA\]](https://arxiv.org/abs/2007.02512), 2020.
- [41] L. Greengard, D. Gueyffier, P.-G. Martinsson, V. Rokhlin, Fast direct solvers for integral equations in complex three-dimensional domains, *Acta Numer.* 18 (2009) 243–275.
- [42] K.L. Ho, L. Greengard, A fast direct solver for structured linear systems by recursive skeletonization, *SIAM J. Sci. Comput.* 34 (2012) A2507–A2532.
- [43] P.-G. Martinsson, V. Rokhlin, A fast direct solver for boundary integral equations in two dimensions, *J. Comput. Phys.* 205 (2005) 1–23.

- [44] S. Börm, Directional H^2 -matrix compression for high-frequency problems, *Numer. Linear Algebra Appl.* 24 (6) (2017) e2112.
- [45] H. Guo, Y. Liu, J. Hu, E. Michielssen, A butterfly-based direct integral equation solver using hierarchical LU factorization for analyzing scattering from electrically large conducting objects, *IEEE Trans. Antennas Propag.* 65 (9) (2017) 4742–4750.
- [46] Y. Liu, H. Guo, E. Michielssen, A HSS matrix-inspired butterfly-based direct solver for analyzing scattering from two-dimensional objects, *IEEE Antennas Wirel. Propag.* 16 (2016) 1179–1183.
- [47] P. Coulier, H. Pouransari, E. Darve, The inverse fast multipole method: using a fast approximate direct solver as a preconditioner for dense linear systems, *SIAM J. Sci. Comput.* 39 (3) (2017) A761–A796.
- [48] V. Minden, K.L. Ho, A. Damle, L. Ying, A recursive skeletonization factorization based on strong admissibility, *Multiscale Model. Simul.* 15 (2) (2016) 768–796.