



A robust solver for elliptic PDEs in 3D complex geometries

Matthew J. Morse ^{a,*}, Abtin Rahimian ^b, Denis Zorin ^a



^a Courant Institute of Mathematical Sciences, New York University, New York, NY 10003, United States of America

^b Department of Computer Science, University of Colorado - Boulder, Boulder, CO 80309, United States of America

ARTICLE INFO

Article history:

Available online 11 June 2021

ABSTRACT

We develop a boundary integral equation solver for elliptic partial differential equations on complex 3D geometries. Our method is efficient, high-order accurate and robustly handles complex geometries. A key component is our singular and near-singular layer potential evaluation scheme, hedgehog: a simple extrapolation of the solution along a line to the boundary. We present a series of geometry-processing algorithms required for hedgehog to run efficiently with accuracy guarantees on arbitrary geometries and an adaptive upsampling scheme based on a iteration-free heuristic for quadrature error. We validate the accuracy and performance with a series of numerical tests and compare our approach to a competing local evaluation method.

© 2021 Elsevier Inc. All rights reserved.

1. Introduction

Linear elliptic homogeneous partial differential equations (PDEs) play an important role in modeling many physical interactions, including electrostatics, elastostatics, acoustic scattering, and viscous fluid flow. Ideas from potential theory allow us to reformulate the associated boundary value problem (BVP) as an integral equation [31]. The solution to the BVP can then be expressed as a surface convolution against the PDE's fundamental solution called a *layer potential*. Discretizing this boundary integral equation (BIE) formulation offers several advantages over commonly used PDE discretization methods such as finite element or finite volume methods.

First, the system of equations uses asymptotically fewer variables because only the boundary of the PDE's domain requires discretization. There is no need to directly discretize the domain itself, which is often time-consuming and error-prone, especially when complex or unbounded domains are involved. This makes the boundary integral formulation well-suited for electromagnetic problems [47] and indispensable for particulate flow simulations with changing, moving, or deforming geometries [49]. Second, although the algebraic system resulting from discretization of BIE's is dense, efficient methods based on the Fast Multipole Method [26] can solve it in $O(N)$ time. A suitable integral formulation can yield a well-conditioned system that can be solved using an iterative method like GMRES in relatively few iterations. Third, high-order quadrature rules can be leveraged to dramatically improve the accuracy of a given discretization size.

For elliptic problems with smooth domain boundaries, fast, high-order methods have a significant advantage over standard methods, drastically reducing the number of degrees of freedom needed to approximate a solution to a given accuracy. However, achieving this with a BIE discretization presents a significant challenge. In particular, integral equation solvers require accurate quadrature rules for *singular* integrals, as the formulation requires the solution of an integral equation involving the singular fundamental solution of the PDE. Moreover, if the solution needs to be evaluated arbitrarily close to the

* Corresponding author.

E-mail addresses: mmorse@cs.nyu.edu (M.J. Morse), arahimian@acm.org (A. Rahimian), dzorin@cs.nyu.edu (D. Zorin).

boundary, then one must numerically compute *nearly singular* integrals with high-order accuracy (e.g., [13,36,75]). Precomputing high-order singular/near-singular quadrature weights also presents a considerable problem. Such weights necessarily depend on the surface geometry, so each sample point requires a unique set of weights. Furthermore, the sampling density required for accurate singular/near-singular integration is highly dependent on the boundary geometry. For example, two nearly touching pieces of the boundary require a sampling density proportional to the distance between them. Applying such a fine discretization globally would be prohibitively expensive, highlighting the need for adaptive refinement.

1.1. Contributions

Our main contribution is a high-order, boundary integral solver for non-oscillatory elliptic PDEs, and experimental evaluation of this solver. An earlier parallel version of this method is used in [42] to simulate red blood cell flows through complex blood vessel with high numerical accuracy. More specifically, the main features of our solver include:

- **Singular and near-singular quadrature scheme.** We introduce an approximation-based singular/near-singular quadrature scheme for single- and double-layer potentials in 3D: after computing the solution at a set of nearby *check points*, placed along a line intersecting the target, we extrapolate the solution to the target point. We have named this scheme hedgehog, for reasons that are apparent from Fig. 2. In order to ensure accuracy of the scheme for complex geometries, a key component of our scheme is a set of geometric criteria for surface sampling needed for accurate integration. Our approach is motivated by the near-singular evaluation scheme of [52,75], which implements a similar scheme that includes an additional on-surface singular evaluation to allow for interpolation of the solution. We eliminate the need for explicit on-surface singular evaluation. An important consequence of this includes the use of *smooth* quadrature rules only, removing the need for an explicit singular quadrature scheme. This allows for much greater flexibility in the choice of surface representation (e.g., the representation of [75] was explicitly designed to support singular quadratures).
- **Surface representation.** Our quadrature scheme enables us to use standard Bézier patches to define the domain boundary, which simplifies the use of the solver on CAD geometry, increases the efficiency of surface evaluation and simplifies parallelization. It also allows for adaptive *quad-trees* of patches to approximate complex surfaces with nonuniform curvature distribution efficiently. Our method can be applied to other surface representations with minimal changes.
- **Refinement for geometric admissibility and quadrature accuracy.** An essential aspect of our method is a set of fast adaptive geometry refinement algorithms to ensure that the assumptions required for the validity and accuracy of hedgehog are satisfied. These conditions are similar in spirit to [54] and [68], but adapted to the geometry of our particular quadrature scheme. To guarantee quadrature accuracy of our method, we detail an adaptive *h*-refinement approach for the integral equation discretization points.

We evaluate hedgehog for a variety of problems on complex geometries to demonstrate high-order convergence and compare to [75].

1.2. Related work

We restrict our discussion to elliptic PDE solvers in 3D using boundary integral formulations. The common schemes to discretize boundary integral equations are the *Galerkin* method, the *collocation* method, and the *Nyström* method [2]. Galerkin and collocation methods are usually referred as Boundary Element Methods (BEM). BEM has been applied to a variety of problems in elastodynamics, electromagnetics and acoustics [1,16,17]. There are a variety of BEM implementations available; one that is most notable is BEM++, which includes high-order elements [57] with extensions for adaptivity added in [8,12]. In this paper, we focus on the *Nyström* discretization, in which the integral in the equation is replaced by its quadrature approximation. The *Nyström* method is simple, yet it enables very efficient methods to solve the discretized integral equation. Compared to BEM methods, *Nyström* methods tend to be more efficient, especially for changing or moving surfaces. However, *Nyström* methods are more difficult to apply to non-smooth surfaces (we do not consider high-order methods for surfaces with sharp edges and corners in this work).

The key element of *Nyström* methods for BIE equations is efficient quadrature rules for singular and near-singular integrals. In the BIE literature, such integration schemes fall into one of the several categories: *singularity cancellation*, *asymptotic correction*, *singularity subtraction*, *custom quadratures* or approximation-based quadrature schemes.

Singularity cancellation schemes apply a change of variables to remove the singularity in the layer potential, allowing for the application of standard smooth quadrature rules. The first polar change of variables was detailed in the context of acoustic scattering [13], which leveraged a partition of unity and a polar quadrature rule to remove the singularity in the integrand of layer potential. The method was extended to open surfaces in [14]. This methodology was applied to general elliptic PDEs in [75] and coupled with the kernel-independent fast multipole method [74] and a general C^∞ surface representation for complex geometries [76]. Its advantages and disadvantages compared to hedgehog are discussed in Section 6. Recently, [44] demonstrated that the choice of partition of unity function used for the change of variables has a dramatic effect on overall convergence order. The first singularity cancellation scheme in 3D on general surfaces composed of piecewise smooth triangles was presented in [10,11]. [25] introduced a change of variables method for acoustic scattering on 3D surfaces, parametrized by spherical coordinates by integrating over a rotated coordinate system that cancels out the singularity.

Asymptotic correction methods study the inaccuracies due to the singular PDE kernel with asymptotic analysis and apply a compensating correction. [9,15,63] compute the integral with a regularized kernel and add corrections for regularization and discretization for the single and double layer Laplace kernel in 3D, along with the Stokeslet and stresslet in 3D. [18] computes an asymptotic expansion of the kernel itself, which is used to remove the aliasing error incurred when applying smooth quadrature rules to near-singular layer potentials. This method is extended to 3D in [19] and a complete asymptotic analysis of the double-layer integral is performed in [37]. *Singularity subtraction methods* [33,34] explicitly subtract the singular component of the integrand analytically, which produces a smooth bounded integral that can be integrated with standard quadrature rules. However, the analytic calculations involved in these approaches are often tailored to a particular PDE and require recalculation for each new PDE of interest.

Custom quadrature rules aim to integrate a particular family of functions to high-order accuracy. This can allow for arbitrarily accurate and extremely fast singular integration methods, since the quadrature rules can be precomputed and stored [5,73].

Our method falls into the final category: *approximation-based quadrature schemes*. The first use of a local expansion to approximate a layer potential near the boundary of a 2D boundary was presented in [6]. By using a refined, or *upsampled*, global quadrature rule to accurately compute coefficients of a Taylor series, the resulting expansion serves as a reasonable approximation to the solution near the boundary where quadrature rules for smooth functions are inaccurate. This scheme was then adapted to evaluate the solution both near and on the boundary, called Quadrature by Expansion (QBX) [21,36]. The first rigorous error analysis of the truncation error of QBX was carried out in [21].

A fast implementation of QBX in 2D, along with a set of geometric constraints required for well-behaved convergence, was presented in [54]. However, the interaction of the expansions of QBX and the translation operator expansions of the FMM resulted in a loss of accuracy, which required an artificially high multipole order to compensate for this additional error. [67] addresses this shortcoming by enforcing a confinement criteria on the location of expansion disks relative to FMM tree boxes. [3] provided extremely tight error heuristics for various kernels and quadrature rules in 2D using contour integration and the asymptotic approach of [22]. [4] then leveraged these estimates in a QBX algorithm for Laplace and Helmholtz problems in 2D that adaptively selects quadrature upsampling and the expansion order for each QBX expansion. In the spirit of [74], [53] generalizes QBX to any elliptic PDE by using potential theory to form a local, least-squares solution approximation using only evaluations of the PDE's fundamental solution.

The first extension of QBX to 3D was [62], where the authors present a *local, target-specific* QBX method on spheroidal geometries. In a local QBX scheme, an upsampled accurate quadrature is used as a local correction to the expansion coefficients computed from the coarse quadrature rule over the boundary. This is in contrast with a *global* scheme, where the expansion coefficients are computed from the upsampled quadrature with no need for correction. The first local QBX scheme appears in [6] in 2D, but the notion of local FMM corrections dates back to earlier work such as [5,38]. The expansions in [62] computed in a target-specific QBX scheme can only be used to evaluate a single target point, but each expansion can be computed at a lower cost than a regular expansion valid in a disk. The net effect of both these algorithmic variations are greatly improved constants, which are required for complicated geometries in 3D. [68] extends the QBX-FMM coupling detailed in [67] to 3D surfaces, along with the geometric criteria and algorithms of [54] that guarantees accurate quadrature. [69] improves upon this by adding target-specific expansions to [68], achieving a 40% speed-up and [70] provides a thorough error analysis of the interaction between computing QBX expansions and FMM local expansions.

In addition to techniques described above, a singular quadrature scheme of [29], further extended to 2D Stokes flows in [72] and to near-singular 3D line integrals in [35], does not fit into one of the above categories. While this method performs exceptionally well in practice, it does not immediately generalize to 3D surfaces in an efficient manner.

Most techniques mentioned above assume smooth domain boundaries or use adaptive refinement to handle non-smooth features. There has been a great deal of recent work on special quadratures for regions with corners [30,55,58–61]. Although not yet generalized to 3D, this work has the potential to vastly improve the performance of 3D Nyström boundary integral methods on regions with corners and edges.

A way to avoid singular quadratures entirely is to use the *method of fundamental solutions* (MFS), which represents the solution as a sum of point charges on an equivalent surface outside of the PDE domain. MFS was successfully applied in 2D [7] and in axis-symmetric 3D problems [40]. Recently, [27] has introduced an 2D approach similar in spirit to MFS, but reformulated as a rational approximation problem. Eliminating the need for singular integration makes these methods advantageous, but placing the point charges robustly can be challenging in practice and general 3D geometries remain a challenge.

We also briefly mention the use of *isogeometric analysis* (IGA) [28] in the context of boundary integral equations. IGA aims to use the same basis functions for geometry and solution representation, in particular, similar to our work, reducing the gap between representations used in CAD, and those needed for high-order BEM. IGA has been successfully applied to singular and hypersingular boundary integral equations with a collocation discretization [65]. A Nyström IGA method coupled with a regularized quadrature scheme is detailed in [77].

The rest of the paper is organized as follows: In Section 2, we briefly summarize the problem formulation, geometry representation and discretization. In Section 3, we detail our singular evaluation scheme and with algorithms to enforce admissibility, adaptively upsample the boundary discretization, and query surface geometry to evaluate singular/near-singular integrals. In Section 4, we provide error estimates for hedgehog. In Section 5, we summarize the complexity of each of the

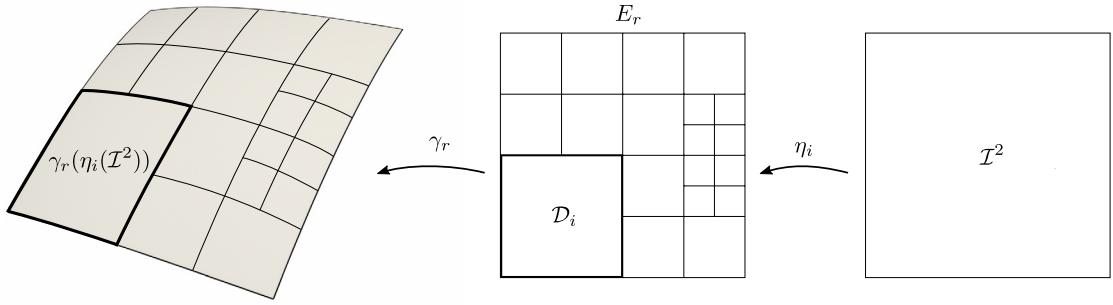


Fig. 1. PATCH QUADRISECTION. Right: the standard domain \mathcal{I}^2 of a single surface or quadrature patch. Middle: a collection of subdomains \mathcal{D}_i of E_r , produced by quadrisecion. Each \mathcal{D}_i corresponds to a map η_i such that $\mathcal{D}_i = \eta_i(\mathcal{I}^2)$; a single \mathcal{D}_i is highlighted in bold. Left: the image of E_r under the patch γ_r . The final image of each subdomain is outlined, with the image of \mathcal{D}_i in bold.

algorithms described in Section 3. In Section 6, we detail convergence tests of our singular evaluation scheme and compare against other state-of-the-art methods.

2. Formulation

2.1. Problem setup

We restrict our focus to interior Dirichlet boundary value problems of the form

$$Lu(\mathbf{x}) = 0, \quad \mathbf{x} \in \Omega, \quad (1)$$

$$u(\mathbf{x}) = f(\mathbf{x}), \quad \mathbf{x} \in \partial\Omega = \Gamma, \quad (2)$$

with multiply- or singly-connected domain Ω of arbitrary genus. Our approach applies directly to standard integral equation formulations of exterior Dirichlet and Neumann problems; we include results for an exterior Dirichlet problem in Section 6.4. Here L is a linear elliptic operator and f is at least C^k . While our method can be applied to any non-oscillatory elliptic PDE, we use the following equations in our examples:

$$Lu = \begin{cases} \Delta u & \text{Laplace} \\ \Delta u - \nabla p, \quad \nabla \cdot u = 0 & \text{Stokes} \\ \Delta u + \frac{1}{1-2\nu} \nabla \nabla \cdot u & \text{Navier (linear elasticity)} \end{cases} \quad (3)$$

We follow the approach of [75]. We can express the solution at a point $\mathbf{x} \in \Omega$ in terms of the double-layer potential

$$u(\mathbf{x}) = D[\phi](\mathbf{x}) = \int_{\Gamma} \frac{\partial G(\mathbf{x}, \mathbf{y})}{\partial \mathbf{n}(\mathbf{y})} \phi(\mathbf{y}) d\mathbf{y}_{\Gamma}, \quad (4)$$

where $G(\mathbf{x}, \mathbf{y})$ is the *fundamental solution* or *kernel* of Eq. (2), $\mathbf{n}(\mathbf{y})$ is the normal at \mathbf{y} on Γ pointing into the exterior of Ω , and ϕ is an unknown function, or *density*, defined on Γ . We list the kernels associated with the PDEs in Eq. (3) in [46, Section 1]. Using the jump relations for the interior and exterior limits of $u(\mathbf{x})$ as \mathbf{x} tends towards Γ [39,45,48,50], we know that Eq. (4) is a solution to Eq. (2) if ϕ satisfies

$$\left(\frac{1}{2}I + D + M \right) [\phi](\mathbf{x}) = f(\mathbf{x}), \quad \mathbf{x} \in \Gamma \quad (5)$$

with identity operator I . We will refer to ϕ as the *density* and $u(\mathbf{x})$ as the *potential* at \mathbf{x} . The double-layer integrals in this equation are *singular*, due to the singularity in the integrand of Eq. (4). Additionally, as \mathbf{x} approaches Γ , Eq. (4) becomes a *nearly singular* integral.

The operator M completes the rank of $\frac{1}{2}I + D$ to ensure invertibility of Eq. (5). If $\frac{1}{2}I + D$ is full-rank, $M = 0$. When $\frac{1}{2}I + D$ has a non-trivial null space, M accounts for the additional constraints to complete the rank of the left-hand side of Eq. (5). For example, for the exterior Laplace problem on ℓ multiply-connected domains, the null space of $\frac{1}{2}I + D$ has dimension ℓ [62]. The full set of cases for each kernel is considered in this work and their corresponding values of M have been detailed in [75].

2.2. Geometry representation

We assume that the smooth domain boundary Γ is given by a *quadrilateral mesh* consisting of quadrilateral faces Q_r referred to as *quads*. Each quad is associated with a parametric domain $\mathcal{I}^2 = [-1, 1]^2 = E_r$, along with embeddings $\gamma_r : E_r \rightarrow \mathbb{R}^3$ for each quad such that $Q_r = \gamma_r(E_r)$. We assume that the quad mesh is *conforming*, i.e., two non-disjoint faces either share a whole edge or a single vertex; examples of this are shown in Figs. 8 and 9. We assume that no two images $\gamma_r(E_r)$ intersect, except along the shared edge or vertex. The surface Γ is the union of patches $\cup_r \gamma_r(E_r) = \cup_r Q_r$. We also assume that Γ is sufficiently smooth to recover the solution of Eq. (2) up to the boundary [39] and is at least C^k .

To represent the surface geometry, we approximate Γ with a collection of *Bézier patches*, given by a linear combination of tensor-product Bernstein polynomials

$$\mathbf{P}_i(s, t) = \sum_{\ell=0}^n \sum_{m=0}^n \mathbf{a}_{\ell m}^{(i)} B_\ell^n(s) B_m^n(t), \quad (6)$$

where $B_\ell^n(t) = \binom{n}{\ell} t^{n-\ell} (1-t)^\ell$ for each ℓ are the n -th degree Bernstein polynomials, i denotes the index of a patch in the collection and $\mathbf{a}_{\ell m}^{(i)} \in \mathbb{R}^3$. Each patch \mathbf{P} is a vector function from \mathcal{I}^2 to \mathbb{R}^3 , so $s, t \in [-1, 1]$. We will refer to this approximation of Γ as $\hat{\Gamma}$.

The domain E_r of each embedding function γ_r is adaptively refined using *quadrisection*, i.e., splitting a square domain into four square subdomains of equal size. Quadrisection induces a *quadtree* structure on each E_r . The root of the quadtree is the original domain \mathcal{I}^2 and each node of the tree is related by a single quadrisection of a subdomain of E_r . The leaves of the quadtree form a collection of subdomains \mathcal{D}_i whose union equals E_r , as shown in Fig. 1-middle. Given an indexing scheme of all \mathcal{D}_i 's over all E_r 's, we define the function $r(i)$ that maps the leaf node index i to its root node index r in the quadtree forest, indicating that $\mathcal{D}_i \subset E_r$. For each r , E_r can have a distinct sequence of associated quadrisects and therefore a distinct quadtree structure. We refer to the process of *refinement* or *refining a patch \mathbf{P}* as the construction of such quadtrees for each E_r subject to some set of criteria.

On each \mathcal{D}_i at the quadtree leaves, we define a Bézier patch and reparametrize each patch over \mathcal{I}^2 by defining the affine map $\eta_i : \mathcal{I}^2 \rightarrow E_{r(i)}$ such that $\eta_i(\mathcal{I}^2) = \mathcal{D}_i \subseteq E_{r(i)}$. It follows that the set of subdomains $\{\eta_i(\mathcal{I}^2) | r(i) = \kappa\}$ form a cover of E_κ and $\{\gamma_\kappa(\eta_i(\mathcal{I}^2)) | r(i) = \kappa\}$ likewise covers $\gamma_\kappa(E_\kappa)$. We summarize this setup in Fig. 1; examples of surfaces of this form can be seen in Figs. 8, 9, 12 and 13.

2.3. Problem discretization

We use two collections of patches in the form described above: $\mathcal{P}_{\text{coarse}}$ and $\mathcal{P}_{\text{fine}}$. The patches in $\mathcal{P}_{\text{coarse}}$, called *surface patches*, determine $\hat{\Gamma}$ from Γ and the set of patches $\mathcal{P}_{\text{fine}}$, called *quadrature patches*, are obtained by further quadrisection of the surface patches in $\mathcal{P}_{\text{coarse}}$. The geometry of $\hat{\Gamma}$ is not changed by this additional refinement of $\mathcal{P}_{\text{coarse}}$, but the total number of subdomains $E_{r(i)}$ is increased. We will detail the geometric criteria that $\mathcal{P}_{\text{coarse}}$ and $\mathcal{P}_{\text{fine}}$ must satisfy in Section 3.2. Discretizing $\hat{\Gamma}$ with a quadrature rule based on $\mathcal{P}_{\text{fine}}$ results in a denser sampling of $\hat{\Gamma}$ than a similar discretization of $\mathcal{P}_{\text{coarse}}$. We will refer to $\mathcal{P}_{\text{coarse}}$ as the *coarse discretization* of $\hat{\Gamma}$ and $\mathcal{P}_{\text{fine}}$ as the *upsampled* or *fine discretization* of $\hat{\Gamma}$.

We index the patches in $\mathbf{P}_i \in \mathcal{P}_{\text{coarse}}$ by $i = 1, \dots, N$; we can then rewrite Eq. (4) as a sum of integrals over surface patches:

$$u(\mathbf{x}) = \sum_{i=1}^N \int_{\mathbf{P}_i} \frac{\partial G(\mathbf{x}, \mathbf{y})}{\partial \mathbf{n}(\mathbf{y})} \phi(\mathbf{y}) d\mathbf{y}_{\mathbf{P}_i}. \quad (7)$$

We discretize functions defined on $\hat{\Gamma}$, such as Eq. (7), at q -node composite tensor-product Clenshaw-Curtis quadrature points on \mathcal{I}^2 of patches in $\mathcal{P}_{\text{coarse}}$. We refer to these points and weights on a single patch \mathbf{P}_i as x_j and w_j^{CC} respectively, for $j = 1 \dots q^2$. The quadrature point \mathbf{y}_{ij} from \mathbf{P}_i is defined as $\mathbf{y}_{ij} = \mathbf{P}_i(\eta_i(x_j))$. We assume that the boundary condition f is given by a black-box evaluator on \mathbb{R}^3 that can be used to obtain values at \mathbf{y}_{ij} . For clarity, we reindex the surface points by a global index $I = 1, \dots, q^2 N$. We discretize the double layer integral Eq. (7) on $\mathcal{P}_{\text{coarse}}$ to approximate the solution $u(\mathbf{x})$:

$$u(\mathbf{x}, \mathcal{P}_{\text{coarse}}) \approx \hat{u}(\mathbf{x}, \mathcal{P}_{\text{coarse}}) = \sum_{i=1}^N \sum_{j=1}^{q^2} \frac{\partial G(\mathbf{x}, \mathbf{y}_{ij})}{\partial \mathbf{n}(\mathbf{y}_{ij})} \phi_{ij} \sqrt{g_{ij}} w_j^{\text{CC}} = \sum_{I=1}^{q^2 N} \frac{\partial G(\mathbf{x}, \mathbf{y}_I)}{\partial \mathbf{n}(\mathbf{y}_I)} \phi_I \hat{w}_I \quad (8)$$

with g_{ij} being the determinant of the metric tensor of \mathbf{P}_i at x_j and $\hat{w}_{i+q^2+j} = \sqrt{g_{ij}} w_j^{\text{CC}}$. In other words, $\hat{u}(\mathbf{x}, \mathcal{P}_{\text{coarse}}) = \hat{D}[\phi](\mathbf{x})$, where $\hat{D}[\phi](\mathbf{x}) \approx D[\phi](\mathbf{x})$.

We can also discretize functions with tensor-product Clenshaw-Curtis nodes on the domains of patches in $\mathcal{P}_{\text{fine}}$. The values of functions on $\mathcal{P}_{\text{fine}}$ are *interpolated* from their values on the quadrature nodes of $\mathcal{P}_{\text{coarse}}$ rather than being computed directly on $\mathcal{P}_{\text{fine}}$. We call this interpolation from $\mathcal{P}_{\text{coarse}}$ to $\mathcal{P}_{\text{fine}}$ *upsampling*. We denote the quadrature nodes and weights

on $\mathcal{P}_{\text{fine}}$ by \tilde{x}_j and \tilde{w}_j with a similar global index J and refer to them as the *upsampled* nodes and weights. Identical formulas are used for computing quadrature on $\mathcal{P}_{\text{fine}}$ with the nodes and weights \tilde{x}_j , \tilde{w}_j on $\mathcal{P}_{\text{fine}}$, denoted $u(\mathbf{x}, \mathcal{P}_{\text{fine}})$ and $\hat{u}(\mathbf{x}, \mathcal{P}_{\text{fine}})$, respectively.

In the next section, we describe the algorithm to compute an accurate approximation to the singular/near-singular double-layer integral in Eq. (4), using a quadrature rule for smooth functions (Eq. (8)) as a building block. This algorithm allows us to compute the matrix-vector products $A\phi$, for a vector of values ϕ defined at the quadrature points \mathbf{y}_I , where A is the discrete operator obtained from the left-hand side of Eq. (5) after approximating $D[\phi](\mathbf{y})$ with the singular integration scheme. As a result, we can solve the linear system using GMRES, which only requires a matrix-vector product

$$A\phi = f, \quad (9)$$

where f is the boundary condition sampled at the points \mathbf{y}_I . The evaluation of these integrals is accelerated in a standard manner using the fast multipole method (FMM) [26,43,74].

3. Algorithms

We now detail a set of algorithms to solve the integral equation in Eq. (5) and evaluate the solution via the double layer integral in Eq. (4) at a given target point $\mathbf{x} \in \Omega$. As described in the previous section, both solving Eq. (5) and evaluating Eq. (4) require accurate evaluation of singular/near-singular integrals of functions defined on the surface $\hat{\Gamma}$. We first outline our unified singular/near-singular integration scheme, hedgehog, its relation to existing approximation-based quadrature methods and geometric problems that can impede accurate solution evaluation. We then describe two geometry preprocessing algorithms, *admissibility refinement* and *adaptive upsampling*, that address these issues to obtain the sets of patches $\mathcal{P}_{\text{coarse}}$ and $\mathcal{P}_{\text{fine}}$ used by hedgehog.

3.1. Singular and near-singular evaluation

We begin with an outline of the algorithm. For a point $\mathbf{s}_x \in \hat{\Gamma}$ on a patch \mathbf{P} from $\mathcal{P}_{\text{coarse}}$ that is closest to \mathbf{x} , we first upsample the density ϕ from $\mathcal{P}_{\text{coarse}}$ to $\mathcal{P}_{\text{fine}}$ and compute the solution at a set of points \mathbf{c}_s , $s = 1, \dots, p$ called *check points*, sampled along the surface normal at \mathbf{s}_x away from $\hat{\Gamma}$. We use Eq. (8) to approximate the solution at the check points. We then extrapolate the solution to \mathbf{x} .

For a given surface or quadrature patch $\mathbf{P} : \mathcal{I}^2 \rightarrow \mathbb{R}^3$, we define the *characteristic length* $L(\mathbf{P})$ as the square root of the surface area of \mathbf{P} , i.e., $L(\mathbf{P}) = \sqrt{\int_{\mathbf{P}} d\mathbf{y}_{\mathbf{P}}}$. We use $L = L(\mathbf{P})$ or L_y for $\mathbf{y} \in \mathbf{P}(D)$ to denote the characteristic length when \mathbf{P} is clear from context. For a point $\mathbf{x} \in \Omega$, we assume that there is a single closest point $\mathbf{s}_x \in \hat{\Gamma}$ to \mathbf{x} ; all points to which the algorithm is applied will have this property by construction. Note that $\mathbf{n}(\mathbf{s}_x)$, the vector normal to $\hat{\Gamma}$ at \mathbf{s}_x , is chosen to point outside of Ω .

We define three zones in Ω for which Eq. (4) is evaluated differently in terms of Eq. (8) and the desired solution accuracy ϵ_{target} . The *far field* $\Omega_F = \{\mathbf{x} \in \Omega \mid \|u(\mathbf{x}) - \hat{u}(\mathbf{x}; \mathcal{P}_{\text{coarse}})\|_2 \leq \epsilon_{\text{target}}\}$, where the quadrature rule corresponding to $\mathcal{P}_{\text{coarse}}$ is sufficiently accurate, and the *intermediate field* $\Omega_I = \{\mathbf{x} \in \Omega \mid \|u(\mathbf{x}) - \hat{u}(\mathbf{x}; \mathcal{P}_{\text{fine}})\|_2 \leq \epsilon_{\text{target}}\}$, where quadrature over $\mathcal{P}_{\text{fine}}$ is sufficiently accurate. The remainder of Ω is the *near field* $\Omega_N = \Omega \setminus \Omega_I$.

Non-singular integration. To compute the solution at points \mathbf{x} in Ω_F , Eq. (8) is accurate to ϵ_{target} , so we can simply compute $\hat{u}(\mathbf{x}, \mathcal{P}_{\text{coarse}})$ directly. Similarly for points in $\Omega_I \setminus \Omega_F$, we know by definition that $\hat{u}(\mathbf{x}, \mathcal{P}_{\text{fine}})$ is sufficiently accurate, so it can also be applied directly.

Singular/near-singular integration algorithm. For the remaining points in Ω_N , we need an alternative means of evaluating the solution. In the spirit of the near-singular evaluation method of [75], we construct a set of *check points* $\mathbf{c}_0, \dots, \mathbf{c}_p$ in Ω_I along a line intersecting \mathbf{x} to approximate the solution near \mathbf{x} . However, instead of interpolating the solution as in [75], we instead extrapolate the solution from the check points to \mathbf{x} . We define two distances relative to \mathbf{s}_x : $R(\mathbf{s}_x) = bL_{\mathbf{s}_x} = \|\mathbf{c}_0 - \mathbf{s}_x\|_2$, the distance from the first check point \mathbf{c}_0 to $\hat{\Gamma}$, and $r(\mathbf{s}_x) = aL_{\mathbf{s}_x} = \|\mathbf{c}_i - \mathbf{c}_{i+1}\|_2$, the distance between consecutive check points. We assume $0 < a, b < 1$.

The overall algorithm for the unified singular/near-singular evaluation scheme is as follows. A schematic for hedgehog is depicted in Fig. 2.

1. Find the closest point \mathbf{s}_x on $\hat{\Gamma}$ to \mathbf{x} .
2. Given values a and b , generate check points $C = \{\mathbf{c}_0, \dots, \mathbf{c}_p\}$

$$\mathbf{c}_s = \mathbf{s}_x - (R(\mathbf{s}_x) + sr(\mathbf{s}_x))\mathbf{n}(\mathbf{s}_x), \quad s = 0, \dots, p \quad (10)$$

The center of mass of these check points $\hat{\mathbf{c}}$ is called the *check center* for \mathbf{x} . Note that $\mathcal{P}_{\text{fine}}$ must satisfy the condition that \mathbf{c}_s are in Ω_I for a given choice of a and b .

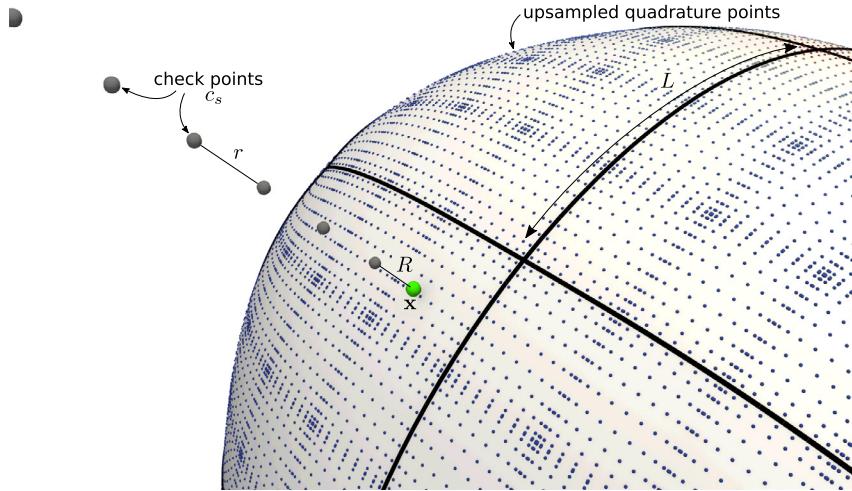


Fig. 2. SCHEMATIC OF SINGULAR/NEAR-SINGULAR EVALUATION. A small piece of a boundary $\hat{\Gamma}$ is shown, along with the set of patches $\mathcal{P}_{\text{coarse}}$ (patch boundaries are drawn in black). The target point \mathbf{x} , in this case on $\hat{\Gamma}$, is shown in green. The solution is evaluated at the check points \mathbf{c}_s (gray points off-surface) using the fine discretization $\mathcal{P}_{\text{fine}}$ (small dots on-surface). The distance from the first check point \mathbf{c}_0 to $\hat{\Gamma}$ is R and the distance between consecutive check points \mathbf{c}_i and \mathbf{c}_{i+1} is r . In this example, $\mathcal{P}_{\text{fine}}$ is computed from $\mathcal{P}_{\text{coarse}}$ with two levels of uniform quadrissection, producing 16 times more patches. The patch length L is roughly proportional to the average edge length of the patch.

3. Upsample ϕ . We interpolate the density values ϕ_I at x_I on patches in $\mathcal{P}_{\text{coarse}}$ to quadrature points \tilde{x}_J on patches in $\mathcal{P}_{\text{fine}}$ with global indices I and J on $\mathcal{P}_{\text{coarse}}$ and $\mathcal{P}_{\text{fine}}$ respectively. If a patch \mathbf{P}_i in $\mathcal{P}_{\text{coarse}}$ is split into m_i patches in $\mathcal{P}_{\text{fine}}$, we are interpolating from q^2 points to $m_i q^2$ points.
4. Evaluate the potential at check points via smooth quadrature with the upsampled density, i.e. evaluate $\hat{u}(\mathbf{c}_s) = \hat{u}(\mathbf{c}_s, \mathcal{P}_{\text{fine}})$ for $s = 0, \dots, p$.
5. Compute a Lagrange interpolant \tilde{u} through the check points $\mathbf{c}_0, \dots, \mathbf{c}_p$ and values $\hat{u}(\mathbf{c}_0), \dots, \hat{u}(\mathbf{c}_p)$ and evaluate at the interpolant at \mathbf{x} :

$$\tilde{u}(\mathbf{x}) = \sum_{s=0}^p \hat{u}(\mathbf{c}_s) \ell_s(t_{\mathbf{x}}), \quad (11)$$

where $\ell_s(\mathbf{x})$ is the s th Lagrange basis function through the points $\mathbf{c}_0, \dots, \mathbf{c}_p$, and $t_{\mathbf{x}} \in \mathbb{R}$ is such that $\mathbf{x} = \mathbf{s}_{\mathbf{x}} - t_{\mathbf{x}} \mathbf{n}(\mathbf{s}_{\mathbf{x}})$ (see Fig. 6 for a schematic of the check points). Since \mathbf{x} lies between \mathbf{c}_0 and $\hat{\Gamma}$, we are extrapolating when computing $\tilde{u}(\mathbf{x})$.

Ill-conditioning of the discrete integral operator. This evaluation scheme can be used directly to extrapolate all the way to the surface and obtain the values of the singular integral in Eq. (5). However, in practice, due to a distorted eigenspectrum of this approximate operator, GMRES tends to stagnate at a level of error corresponding to the accuracy of hedgehog when it is used to compute the matrix-vector product. This is a well-known phenomenon of approximation-based singular quadrature schemes; [36, Section 3.5][53, Section 4.2] present a more detailed study. To address this, we average the interior and exterior limits of the solution at the quadrature nodes, computed via hedgehog, to compute the on-surface potential and add $\frac{1}{2}I$ to produce the interior limit. This shifts the clustering of eigenvalues from around zero to around $\frac{1}{2}$, which is ideal from the perspective of GMRES. We call this *two-sided* hedgehog, while the standard version described above is called *one-sided* hedgehog. We observe stable and consistent convergence of GMRES when two-sided hedgehog is used to evaluate the matrix-vector multiply to solve Eq. (9). In light of this, we always use two-sided hedgehog within GMRES and set the stopping tolerance for GMRES to $\epsilon_{\text{GMRES}} = 10^{-12}$, regardless of the geometry, boundary condition or quadrature order.

3.2. Geometric criteria for accurate quadrature

The accuracy of the method outlined above is controlled by two competing error terms: *quadrature error* incurred from approximating the layer potential Eq. (4) with Eq. (8) in Step 4 and *extrapolation error* due to approximating the singular integral with an extrapolated value in Step 5. Both errors are determined by the location of check points relative to the patches in $\mathcal{P}_{\text{coarse}}$ and $\mathcal{P}_{\text{fine}}$ (see Heuristic 4.1 and Theorem 4.2).

In Fig. 3, we show three examples of different choices of check point locations to evaluate the potential at a point with hedgehog. In Fig. 3-left, \mathbf{c}_0 is placed close to the target point, while in Fig. 3-middle, \mathbf{c}_0 is far from the target point, but

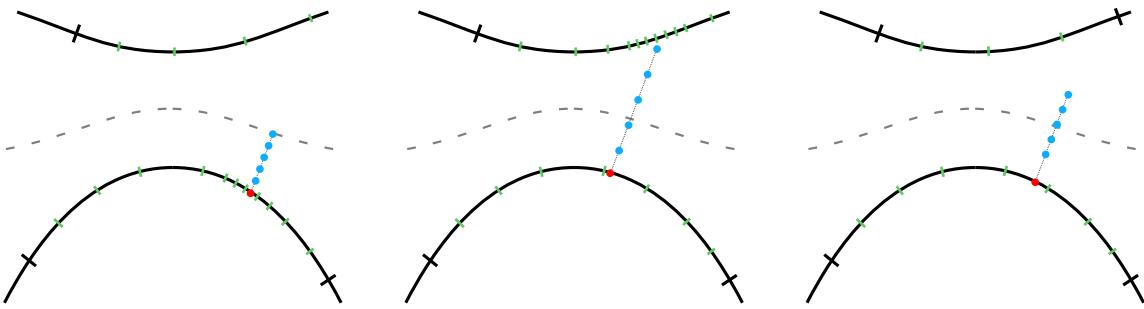


Fig. 3. POSSIBLE CHECK POINT CONFIGURATIONS. A 2D example depicting three choices of a and b in Eq. (10). Shown is the boundary $\hat{\Gamma}$, with black tick marks denoting patch boundaries of $\mathcal{P}_{\text{coarse}}$, green tick marks denoting patch boundaries of $\mathcal{P}_{\text{fine}}$, the target point (red dots), its check points (blue dots) along the normal closest to the target point, and the medial axis of $\hat{\Gamma}$ (gray dotted line). Large (left) and small (middle) values of a and b can cause clustering of check points near to $\hat{\Gamma}$, which requires large amounts of upsampling to compute the potential accurately. Using the medial axis as a heuristic to for admissibility (right), we can minimize the amount of adaptive upsampling required.

\mathbf{c}_p is close to a non-local piece of $\hat{\Gamma}$. Both cases will require excessive refinement of $\mathcal{P}_{\text{coarse}}$ in order to resolve Eq. (8) accurately with $\mathcal{P}_{\text{fine}}$. On the other hand, in Fig. 3-right, we can either perform one refinement step on $\mathcal{P}_{\text{coarse}}$ or adjust a and b , which will result in fewer patches in $\mathcal{P}_{\text{fine}}$, and therefore provide a faster integral evaluation, while maintaining accuracy.

In an attempt to strike this balance between speed and accuracy, we need certain constraints on the geometry of $\hat{\Gamma}$ to ensure the efficient and accurate application of hedgehog, which we impose on the patch sets $\mathcal{P}_{\text{coarse}}$ and $\mathcal{P}_{\text{fine}}$. We will first outline our constraints on the quadrature patch sets $\mathcal{P}_{\text{coarse}}$ and $\mathcal{P}_{\text{fine}}$ which allow for accurate evaluation with hedgehog.

3.2.1. Admissibility criteria

A set of patches \mathcal{P} is *admissible* if the following statements are satisfied on each quadrature patch in \mathcal{P} :

- 1 The error of a surface patch \mathbf{P}_i approximating an embedding γ_r is below some absolute target accuracy ϵ_g
- 2 The interpolation error of the boundary condition f is below some absolute target accuracy ϵ_f
- 3 For each check center $\hat{\mathbf{c}}_j$ corresponding to the quadrature point \mathbf{y}_j on the surface, the closest point on $\hat{\Gamma}$ to $\hat{\mathbf{c}}_j$ is \mathbf{y}_j .

Criterion 1 is required to ensure that $\hat{\Gamma}$ approximates Γ with sufficient accuracy to solve the integral equation. We discuss how to choose ϵ_g in [46, Section 6]; for the tests in this paper, we simply choose $\epsilon_g < \epsilon_{\text{target}}$. Criterion 2 guarantees that f can be represented at least as accurately as the desired solution accuracy. We therefore similarly choose $\epsilon_f < \epsilon_{\text{target}}$. Criterion 3 balances the competing geometric constraints of cost and accuracy by flexibly placing check points as far as possible from $\hat{\Gamma}$ without causing too much upsampling on other patches. If a check point \mathbf{c} constructed from a surface patch \mathbf{P} is too close to another surface patch \mathbf{P}' , Criterion 3 will indicate that \mathbf{P} is inadmissible. If \mathbf{P} is subdivided into its children, new check points \mathbf{c}' generated from these children of \mathbf{P} will be closer to \mathbf{P} and further from \mathbf{P}' . Since check points are placed at distances proportional to $L(\mathbf{P})$, repeated refinement of \mathbf{P} will eventually satisfy Criterion 3.

3.2.2. Upsampling criteria

Once we have a set of admissible surface patches satisfying Criteria 1 to 3, we need to determine the upsampled quadrature patches $\mathcal{P}_{\text{fine}}$ that ensure that the check points generated from $\mathcal{P}_{\text{coarse}}$ are in Ω_I , i.e., $\|u(\mathbf{c}) - \hat{u}(\mathbf{c}, \mathcal{P}_{\text{fine}})\| < \epsilon_{\text{target}}$. To achieve this, we need a criterion to determine which patches are “too close” to a given check point for the error to be below ϵ_{target} . We make the following assumption about the accuracy of our smooth quadrature rule: *Eq. (8) is accurate to ϵ_{target} at points further than $L(\mathbf{P})$ from \mathbf{P} , for $\epsilon_{\text{target}} > 10^{-12}$.* This is motivated by [3,6], which demonstrate the rapid convergence of the layer potential quadrature error with respect to $\|\mathbf{x} - \mathbf{s}_x\|_2$. For sufficiently high quadrature orders, such as $q = 20$, this assumption seems to hold in practice. We say that a point \mathbf{x} is *near* to \mathbf{P} if the distance from \mathbf{x} to \mathbf{P} is less than $L(\mathbf{P})$; otherwise, \mathbf{x} is *far* from \mathbf{P} . We would like all check points required for the singular/near-singular evaluation of the discretization of Eq. (4) using hedgehog to be far from all patches in $\mathcal{P}_{\text{fine}}$. If this is satisfied, then we know that the Clenshaw-Curtis quadrature rule will be accurate to 10^{-12} at each check point.

3.3. Refinement algorithm preliminaries

Computing the distance from a check point to a given patch is a fundamental step in verifying the constraints on $\mathcal{P}_{\text{coarse}}$ and $\mathcal{P}_{\text{fine}}$ from Sections 3.2.1 and 3.2.2. Before detailing our refinement algorithms to enforce these criteria, we introduce several geometric algorithms and data structures that will be used to compute the closest point on piecewise polynomial surfaces.

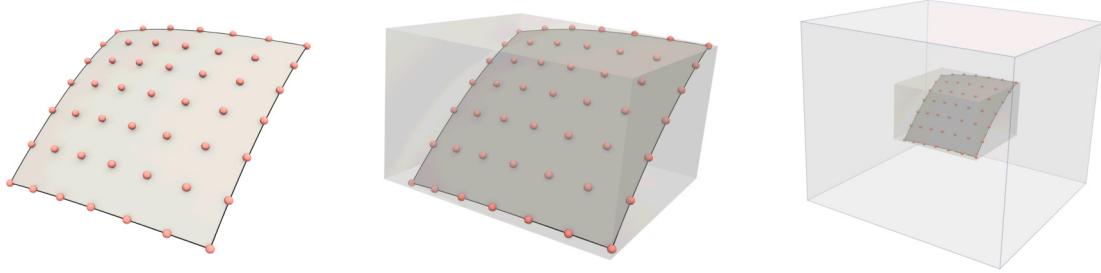


Fig. 4. RELATIONSHIP BETWEEN CONTROL POINTS AND BOUNDING BOXES. Left: a patch in the tensor product Bézier basis, with control points ($\mathbf{a}_{\ell m}$'s from Eq. (6)) plotted. The convex hull of the control points of a patch is guaranteed to contain the patch. Center: The patch bounding box, computed from the control points. Right: The near-zone bounding box of the patch from Section 3.5 computed by inflating the bounding box by $L(\mathbf{P})$.

3.3.1. AABB trees

In order to implement our algorithms to enforce admissibility efficiently, we use a fast spatial data structure to find the patches that are close to a query point \mathbf{x} . In [54,68], the quadtree and octree within an FMM is extended to support the geometric queries needed for a fast QBX algorithm. In this work, we use an axis-aligned bounding box (AABB) tree, which is a type of bounding volume hierarchy [56], implemented in `geogram` [41]. An AABB is a tree with nodes corresponding to bounding boxes and leaves corresponding to bounding boxes containing single objects. A bounding box B_0 is a child of another box B_1 if $B_0 \subset B_1$; the root node is a bounding box of the entire domain of interest. Operations supported by AABB trees include: (i) finding all bounding boxes containing a query point, (ii) finding all bounding boxes that intersect another query box, (iii) finding the closest triangle to a query point (because triangles have trivial bounding boxes). By decoupling geometric queries from fast summation, the individual algorithms can be more thoroughly optimized, in exchange for the additional memory overhead of maintaining two distinct data structures. The query algorithm presented in [42] likely has better parallel scalability, but AABB trees are faster for small to medium problem sizes on a single machine due to less redundant computation.

To define an AABB tree for our patch-based surface $\hat{\Gamma}$, we make use of the following fact: the control points of a Bézier surface ($\mathbf{a}_{\ell m}$'s from Eq. (6)) form a convex hull around the surface that they define [24]. As a result, we can compute a bounding box of a surface or quadrature patch \mathbf{P} directly from the Bézier coefficients simply by computing the maximum and minimum values of each component of the $\mathbf{a}_{\ell m}$'s, as shown in Fig. 4-middle. This bounding box can then be inserted into the AABB tree as a proxy for a surface or quadrature patch.

3.3.2. Computing the closest point to a patch

To find a candidate closest patch \mathbf{P}_{i_0} to \mathbf{x} , we construct a fine triangle mesh and bounding boxes of each patch in $\mathcal{P}_{\text{coarse}}$ and insert them into an AABB tree. We can query the AABB tree for the nearest triangle to \mathbf{x} with the AABB tree, which corresponds to \mathbf{P}_{i_0} . We then compute the accurate true distance d_{i_0} to \mathbf{P}_{i_0} using a constrained Newton method, presented in detail in [46, Section 2].

However, there may be other patches whose distance to \mathbf{x} is less than d_{i_0} , as shown in Fig. 5. To handle this case, we then query the AABB tree for all patches $\mathbf{P}_{i_1}, \dots, \mathbf{P}_{i_k}$ that are distance at most d_{i_0} from \mathbf{x} . This is achieved by forming a query box centered at \mathbf{x} with edge length $2d_{i_0}$ and querying the AABB tree for all intersection bounding boxes. The precise distance is then computed for each patch $\mathbf{P}_{i_1}, \dots, \mathbf{P}_{i_k}$ with [46, Section 2] and the smallest distance is chosen. We summarize this process in Algorithm 1.

Algorithm 1: Compute the closest point to \mathbf{x} .

Data: A set of quadrature patches \mathcal{P} , a query point \mathbf{x} , Newton method tolerance ϵ_{opt}
Result: The closest point $\mathbf{s}_{\mathbf{x}}$ on \mathcal{P} to \mathbf{x}

- 1 Construct an AABB tree T_T from a fine triangle mesh of the quadrature patches of \mathcal{P}
- 2 Construct an AABB tree T_B from bounding boxes of quadrature patches in \mathcal{P} .
- 3 $\tau_0 = \text{closest triangle to } \mathbf{x} \text{ computed with } T_T$
- 4 $\mathbf{P}_{i_0} = \text{patch corresponding to } \tau_0$
- 5 Find the closest point $\mathbf{s}_{\mathbf{x},0}$ on \mathbf{P}_{i_0} to \mathbf{x} with [46, Section 2].
- 6 $d_{i_0} = \|\mathbf{x} - \mathbf{s}_{\mathbf{x},0}\|_2$
- 7 $B_{d_{i_0}}(\mathbf{x}) = \text{a box centered at } \mathbf{x} \text{ with edge length } 2d_{i_0}$
- 8 Find the boxes B_{i_1}, \dots, B_{i_k} in T_B that intersect $B_{d_{i_0}}(\mathbf{x})$
- 9 **for** $B_{i_j} \in B_{i_1}, \dots, B_{i_k}$ **do**
- 10 $\mathbf{P}_{i_j} = \text{quadrature patch corresponding to } B_{i_j}$
- 11 Find the closest point $\mathbf{s}_{\mathbf{x},j}$ on \mathbf{P}_{i_j} to \mathbf{x} with [46, Section 2] to precision ϵ_{opt} .
- 12 $d_{i_j} = \|\mathbf{x} - \mathbf{s}_{\mathbf{x},j}\|_2$
- 13 $j^* = \text{argmin}_j \{d_{i_j}\}$
- 14 **return** $\mathbf{s}_{\mathbf{x},j^*}$

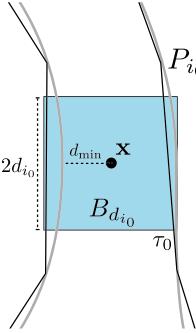


Fig. 5. A 2D SCHEMATIC OF NEAR-PATCH CANDIDATE SELECTION. A visual depiction of the quantities defined in lines 3-7 of Algorithm 1 (shown here in 2D for simplicity), with notation matching Algorithm 4. The triangle-mesh proxy is drawn in as black lines and patches are drawn as gray curves. We have found an initial closest triangle τ_0 to \mathbf{x} corresponding to patch \mathbf{P}_{i_0} and computed $d(\mathbf{x}, \mathbf{P}_{i_0}) = d_{i_0}$. We then query the AABB tree for all patches that intersect box $B_{d_{i_0}}$ with edge length $2d_{i_0}$, shown in blue. There is clearly a patch that is closer to \mathbf{x} than \mathbf{P}_{i_0} that will be returned from the query, which will be distance d_{\min} from \mathbf{x} .

3.4. Admissibility algorithm

Our algorithm to enforce Criteria 1 to 3 proceeds as follows:

- To enforce Criterion 1, we adaptively fit a set of surface patches to the embeddings γ_r representing Γ . We construct a bidegree (n, n) piecewise polynomial least-squares approximation \mathbf{P}_i in the form of Eq. (6) to γ_r on \mathbb{I}^2 . If \mathbf{P}_i 's domain \mathcal{D}_i is obtained by refinement of E_r , we fit $\mathbf{P}_i \circ \eta_i$ to γ_r on \mathbb{I}^2 , using $4n \times 4n$ samples on \mathbb{I}^2 . If the pointwise error of \mathbf{P}_i and its partial derivatives is greater than ϵ_g , then it is quadrisected and the process is repeated.
- Once the embeddings are resolved, we resolve f on each surface patch produced from the previous step in a similar fashion to enforce Criterion 2. However, rather than a least-squares approximation in this stage, we use piecewise polynomial interpolation.
- To enforce Criterion 3, we construct the set of check centers $\hat{\mathbf{c}}_l$ which correspond to the check points required to evaluate the solution at the quadrature nodes \mathbf{y}_l . For each check center $\hat{\mathbf{c}}_l$, we find the closest point $\mathbf{s}_{\hat{\mathbf{c}}_l} \in \hat{\Gamma}$. If $\|\mathbf{s}_{\hat{\mathbf{c}}_l} - \mathbf{y}_l\| \geq \epsilon_{\text{opt}}$, we split the quadrature patch \mathbf{P} containing \mathbf{y}_l . The tolerance ϵ_{opt} is used in the Newton's method in [46, Section 2]; we usually choose $\epsilon_{\text{opt}} = 10^{-14}$. Since $d(\hat{\mathbf{c}}_l, \hat{\Gamma})$ is proportional to $L_{\mathbf{y}_l}$, the new centers $\hat{\mathbf{c}}_l$ for the refined patches will be closer to the surface. We use Algorithm 1 to compute $\mathbf{s}_{\hat{\mathbf{c}}_l}$. However, in the case of check points, we can skip lines 1-6 to compute d_{i_0} , since $\hat{\mathbf{c}}_l$ is $R + r(p+1)/2$ away from $\mathbf{y}_l \in \mathbf{P}(D)$ by construction. We can apply lines 7-14 of Algorithm 1 with $d_{i_0} = R + r(p+1)/2$ to compute $\mathbf{s}_{\hat{\mathbf{c}}_l}$.

We summarize the algorithm to enforce Criterion 3 in Algorithm 2. At each refinement iteration, the offending patches are decreased by quadrisection, which reduces the distance from the quadrature point \mathbf{y}_l to its checkpoints. This eventually satisfies Criterion 3 and the algorithm terminates.

3.5. Adaptive upsampling algorithm

Before detailing our upsampling algorithm to satisfy the criteria outlined in Section 3.2.2, we must define the notion of a *near-zone bounding box* of a quadrature patch \mathbf{P} , denoted $B_{\text{near}}(\mathbf{P})$. The near-zone bounding box of \mathbf{P} is computed as described in Section 3.3.1, but then is inflated by $2L(\mathbf{P})$, as shown in Fig. 4-right. This inflation guarantees that any point \mathbf{x} that is near \mathbf{P} is contained in $B_{\text{near}}(\mathbf{P})$ and, for an admissible set of quadrature patches $\mathcal{P}_{\text{coarse}}$, that any $\mathbf{x} \in \Omega_N$ must be contained in some quadrature patch's near-zone bounding box. This means that by forming $B_{\text{near}}(\mathbf{P})$ for each quadrature patch in $\mathcal{P}_{\text{fine}}$, a check point is in Ω_l if it is not contained in any near-zone bounding boxes.

To compute the upsampled patch set from $\mathcal{P}_{\text{coarse}}$, we initially set $\mathcal{P}_{\text{fine}} = \mathcal{P}_{\text{coarse}}$, compute the near-zone bounding boxes of each patch in $\mathcal{P}_{\text{fine}}$ and insert them into an AABB tree. We also construct the set of check points C required to evaluate our discretized layer-potential with hedgehog (Section 3.1). For each check point $\mathbf{c} \in C$, we query the AABB tree for all near-zone bounding boxes that contain \mathbf{c} . If there are no such boxes, we know \mathbf{c} is far from all quadrature patches and can continue. If, however, there are near-zone bounding boxes B_{i_0}, \dots, B_{i_k} containing \mathbf{c} , we compute the distances d_{i_k} from \mathbf{c} to $\mathbf{P}_{i_1}, \dots, \mathbf{P}_{i_k}$ using [46, Section 2]. If $d_{i_k} < L(\mathbf{P}_{i_k})$, we replace \mathbf{P}_{i_k} in $\mathcal{P}_{\text{fine}}$ with its four children produced by quadrisection.

To improve the performance of this refinement procedure, we allow for the option to skip the Newton method in Algorithm 1 and immediately refine all patches $\mathbf{P}_{i_0}, \dots, \mathbf{P}_{i_k}$. This is advantageous in the early iterations of the algorithm, when most check points are near to patches by design. We allow for a parameter n_{skip} to indicate the number of iterations to skip the Newton optimization and trigger refinement immediately. We typically set $n_{\text{skip}} = 2$. We summarize our algorithm in Algorithm 3.

Algorithm 2: Enforce admissibility Criterion 3 on a set of quadrature patches.

Data: A set of quadrature patches \mathcal{P} , optimization tolerance ϵ_{opt}
Result: An admissible set of quadrature patches \mathcal{P}

- 1 $\mathcal{P} = \mathcal{P}_{\text{coarse}}$
- 2 Mark all patches in \mathcal{P} as inadmissible.
- 3 **while** any patch in \mathcal{P} is inadmissible **do**
- 4 Construct an AABB tree T as described in Section 3.3.2 from \mathcal{P}
- 5 **for** $\mathbf{P} \in \mathcal{P}$ **do**
- 6 **if** \mathbf{P} is inadmissible **then**
- 7 Construct a set of check centers $\mathbf{C}_{\mathbf{P}}$ for each $\mathbf{y}_j \in \mathbf{P}(D)$
- 8 **for** $\hat{\mathbf{c}} \in \mathbf{C}_{\mathbf{P}}$ **do**
- 9 $d_{i_0} = R + r(p+1)/2$
- 10 Compute $\mathbf{s}_{\hat{\mathbf{c}}}$ with lines 7-14 of Algorithm 1 with precision ϵ_{opt} and d_{i_0} .
- 11 **if** $\|\mathbf{s}_{\hat{\mathbf{c}}} - \mathbf{y}_j\|_2 < \epsilon_{\text{opt}}$ **then**
- 12 Mark \mathbf{P} as admissible.
- 13 **else**
- 14 Mark \mathbf{P} as inadmissible.
- 15 **break** // only need one bad check center to mark \mathbf{P} for refinement
- 16 **for** $\mathbf{P} \in \mathcal{P}$ **do**
- 17 **if** \mathbf{P} is inadmissible **then**
- 18 Split \mathbf{P} into its four child patches, mark each as inadmissible, and replace \mathbf{P} with its children in \mathcal{P} .
- 19 **return** \mathcal{P}

Algorithm 3: Adaptively upsample to accurately evaluate Eq. (8) at check points.

Data: An admissible patch set \mathcal{P} , number of iterations n_{skip} before using [46, Section 2]
Result: An upsampled set of quadrature patches

- 1 Compute inflated near-zone bounding boxes B_1, \dots, B_N of each $\mathbf{P} \in \mathcal{P}$.
- 2 Construct an AABB tree T from the near-zone bounding boxes.
- 3 Construct all check points C required to evaluate the Eq. (5) on \mathcal{P} .
- 4 $\mathcal{P}_{\text{fine}} = \mathcal{P}$
- 5 Mark all check points in C as near.
- 6 $i = 0$
- 7 **while** any $\mathbf{c} \in C$ is marked near **do**
- 8 **for** $\mathbf{c} \in C$ **do**
- 9 **if** \mathbf{c} is marked near **then**
- 10 Query T for all bounding boxes B_{i_1}, \dots, B_{i_k} containing \mathbf{c} .
- 11 $\mathbf{P}_{i_1}, \dots, \mathbf{P}_{i_k}$ = patches corresponding to boxes B_{i_1}, \dots, B_{i_k}
- 12 Mark \mathbf{c} as far
- 13 **for** $\mathbf{P} \in \mathbf{P}_{i_1}, \dots, \mathbf{P}_{i_k}$ **do**
- 14 **if** $i > n_{\text{skip}}$ **then**
- 15 Find the closest point $\mathbf{s}_{\mathbf{c}}$ on \mathbf{P} to \mathbf{c} with Algorithm 1.
- 16 **if** $\|\mathbf{s}_{\mathbf{c}} - \mathbf{c}\|_2 < L(\mathbf{P})$ **then**
- 17 Split \mathbf{P} and replace it in $\mathcal{P}_{\text{fine}}$ with its children.
- 18 Mark \mathbf{c} as near
- 19 **else**
- 20 Split \mathbf{P} and replace it in $\mathcal{P}_{\text{fine}}$ with its children.
- 21 Mark \mathbf{c} as near
- 22 $i = i + 1$

3.6. Marking target points for evaluation

Once we have solved Eq. (9) for ϕ on $\hat{\Gamma}$, we need the ability to evaluate Eq. (4) at an arbitrary set of points in the domain. For a target point \mathbf{x} , in order apply the algorithm in Section 3.1, we need to determine whether or not $\mathbf{x} \in \Omega$ and, if so, whether \mathbf{x} is in Ω_N , Ω_I or Ω_F . Both of these questions can be answered by computing the closest point $\mathbf{s}_{\mathbf{x}}$ on $\hat{\Gamma}$ to \mathbf{x} . If $\mathbf{n}(\mathbf{s}_{\mathbf{x}}) \cdot (\mathbf{x} - \mathbf{s}_{\mathbf{x}}) < 0$, then $\mathbf{x} \in \Omega$. As we have seen in Section 3.2.2, the distance $\|\mathbf{x} - \mathbf{s}_{\mathbf{x}}\|$ determines whether $\mathbf{x} \in \Omega_N$, Ω_I or Ω_F . However, for large numbers of target points, a brute force calculation of closest points on $\hat{\Gamma}$ to all target points is prohibitively expensive. We present an accelerated algorithm combining Algorithm 1 and an FMM evaluation to require only constant work per target point.

3.6.1. Marking and culling far points

A severe shortcoming of Algorithm 1 is that its performance deteriorates as the distance from \mathbf{x} to $\hat{\Gamma}$ increases. Consider the case where $\hat{\Gamma}$ is a sphere with radius r with \mathbf{x} at its center. The first stage of Algorithm 1 returns a single quadrature patch that is distance r from \mathbf{x} ; the next stage will return all quadrature patches. This will take $O(N)$ time to check the distance to each patch. Even on more typical geometries, we observe poor performance of Algorithm 1 when \mathbf{x} is far from $\hat{\Gamma}$.

To address this, we use an additional FMM-based acceleration step to mark most points far from $\hat{\Gamma}$ before using applying Algorithm 1. Our approach is based on computing the generalized winding number [32] of $\hat{\Gamma}$ at the evaluation points. For closed curves in \mathbb{R}^2 , the *winding number* at a point counts the number of times the curve travels around that point. The *generalized winding number* of a surface $\hat{\Gamma}$ at a point $\mathbf{x} \in \mathbb{R}^3$ can be written as

$$\omega_{\hat{\Gamma}}(\mathbf{x}) = -\frac{1}{4\pi} \int_{\hat{\Gamma}} \frac{(\mathbf{x} - \mathbf{y}) \cdot \mathbf{n}}{\|\mathbf{x} - \mathbf{y}\|^3} d\mathbf{y}_{\hat{\Gamma}} \quad (12)$$

We recognize this integral as the double-layer potential in Eq. (4) for a Laplace problem with $\phi = 1$. Its values in \mathbb{R}^3 are [39]:

$$\omega_{\hat{\Gamma}}(\mathbf{x}) = \begin{cases} 1 & \mathbf{x} \in \Omega \setminus \hat{\Gamma} \\ 1/2 & \mathbf{x} \in \hat{\Gamma} \\ 0 & \mathbf{x} \in \mathbb{R}^3 \setminus \overline{\Omega} \end{cases} \quad (13)$$

Eq. (12) can be evaluated using the same surface quadrature in Eq. (8) using an FMM in $O(N)$ time. While the quadrature rule is inaccurate close to the surface, Ω_F is defined precisely as the zone where the quadrature rule is sufficiently accurate. For this reason, we use

$$|\omega_{\hat{\Gamma}}(\mathbf{x}) - 1| < \epsilon_{\text{target}} \quad (14)$$

to mark points $\mathbf{x} \in \Omega_F \subset \Omega$ and a similar relation

$$|\omega_{\hat{\Gamma}}(\mathbf{x})| < \epsilon_{\text{target}} \quad (15)$$

to mark points $\mathbf{x} \notin \Omega$. This approach is similar in spirit to the spectrally accurate collision detection scheme of [52, Section 3.5]. Unlike [52], however, we do *not* use singular integration to mark all points. This isn't possible since at this stage since we do not yet know which target points require singular integration. We use the FMM evaluation purely as a culling mechanism before applying the full marking algorithm.

Remark. Since the quadrature rule may be highly inaccurate for points close to the surface, due the near-singular nature of the integrand, $\omega_{\hat{\Gamma}}(\mathbf{x})$ may happen to be close to one or zero. We highlight that it is possible that points outside Ω_F may be mismarked, although we have not observed this in practice.

3.6.2. Full marking algorithm

We combine the algorithms of the previous two sections into a single marking pipeline for a general set of target points in \mathbb{R}^3 , by first applying the algorithm of Section 3.6.1 to mark all points satisfying Eq. (14) then passing the remaining points to Algorithm 1. The full marking algorithm is summarized as Algorithm 4.

4. Error analysis

As with other approximation-based quadrature methods, hedgehog has two primary sources of error: the quadrature error e_Q incurred as a result of evaluating potential at the check points and the extrapolation error e_E due to evaluating the polynomial approximation of the potential at the target point, assuming $\mathcal{P}_{\text{coarse}}$ is admissible. Let

$$e_Q(\mathbf{x}) = \left| \sum_{s=0}^p (u(c_s) - \hat{u}(c_s, \mathcal{P}_{\text{fine}})) \ell_s(t_{\mathbf{x}}) \right|, \quad (16)$$

$$e_E(\mathbf{x}) = \left| u(\mathbf{x}) - \sum_{s=0}^p u(c_s) \ell_s(t_{\mathbf{x}}) \right|, \quad (17)$$

$$e_{\text{hedgehog}}(\mathbf{x}) \leq e_Q(\mathbf{x}) + e_E(\mathbf{x}), \quad (18)$$

where $u(\mathbf{x})$ and $\hat{u}(\mathbf{x}, \mathcal{P}_{\text{fine}})$ are defined in Eqs. (4) and (8) and $\ell_s(t)$ is the s -th Lagrange polynomial defined on the points $\{0, 1, \dots, p\}$. We define $t_{\mathbf{x}}$ such that $\mathbf{x} = -\mathbf{n}(\mathbf{y})(R + t_{\mathbf{x}}r)$, so $t_{\mathbf{x}} = \frac{\|\mathbf{x} - \mathbf{y}\| - R}{r}$. In this section, we first prove that we achieve high-order accuracy with our singular/near-singular evaluation scheme in Section 3.1 with respect to extrapolation order p and quadrature order q . We then detail the impact of surface approximation on overall solution accuracy.

Algorithm 4: Mark points in regions Ω_F , Ω_I and Ω_N .

Data: An admissible set of quadrature patches \mathcal{P} , ϵ_{target} , target points \mathbf{X}
Result: A marked set of target points \mathbf{X}

```

1  $\phi_0 = 1$ 
2  $\omega_{\hat{\Gamma}} = \text{Laplace\_FMM}(\mathcal{P}, \mathbf{X}, \phi_0)$ 
3 for  $\mathbf{x} \in \mathbf{X}$  do
4   if  $|\omega_{\hat{\Gamma}}(\mathbf{x}) - 1| < \epsilon_{\text{target}}$  then
5     | Mark  $\mathbf{x}$  as inside  $\Omega$ .
6     | Mark  $\mathbf{x}$  as in  $\Omega_F$ .
7   else if  $|\omega_{\hat{\Gamma}}(\mathbf{x})| < \epsilon_{\text{target}}$  then
8     | Mark  $\mathbf{x}$  as outside  $\Omega$ .
9 for  $\mathbf{x} \in \mathbf{X}$  do
10  if  $\mathbf{x}$  is unmarked then
11    | Compute the closest point  $\mathbf{s}_x$  to  $\mathbf{x}$  with Algorithm 1
12    |  $d_{\min} = \|\mathbf{s}_x - \mathbf{x}\|_2$ 
13    | if  $d_{\min} \leq L_{\mathbf{s}_x}$  then
14      |   | Mark  $\mathbf{x}$  as in  $\Omega_N$ 
15    | else
16      |   | Mark  $\mathbf{x}$  as in  $\Omega_I$ 
17    | if  $\mathbf{n}(\mathbf{s}_x) \cdot (\mathbf{x} - \mathbf{s}_x) < 0$  then
18      |   | Mark  $\mathbf{x}$  as inside  $\Omega$ 
19    | Mark  $\mathbf{x}$  as outside  $\Omega$ 

```

4.1. Quadrature error

We briefly state a tensor-product variation of known Clenshaw-Curtis quadrature error results as applied to smooth functions in 3D. This estimate is derived based on assumptions detailed in Appendix A that, in general, is difficult to verify in practice and may not hold for all functions we consider. For this reason, we refer to it as a heuristic.

Heuristic 4.1. Let the boundary $\hat{\Gamma}$ be discretized by quadrature patches over the domains $[-h, h]$ and the boundary condition f in Eq. (2) be at least C^k . Apply the q -th order Clenshaw-Curtis quadrature rule to the double-layer potential $u(\mathbf{x})$ given in Eq. (7) and let \mathbf{x} be in the interior of Ω . Then for all sufficiently large q :

$$e_Q(\mathbf{x}) \lesssim \frac{128h^{k+1}}{15\pi k(2q+1-k)^k} \tilde{V}, \quad (19)$$

where

$$\tilde{V} = \max_{i=1,\dots,N} \max_{\alpha,\beta \leq k} \left\| \frac{\partial^{\alpha+\beta}}{\partial u^\alpha \partial v^\beta} \left(\frac{\partial G(\mathbf{x}, \mathbf{P}_i(s, t))}{\partial \mathbf{n}} \phi(\mathbf{P}_i(s, t)) g_{\mathbf{P}_i}(s, t) \right) \right\|_T, \quad (20)$$

$g_{\mathbf{P}}$ is the determinant of the metric tensor of a patch \mathbf{P} implicit in Eq. (7), \lesssim means “approximately less than or equal to,” and $\|\zeta\|_T = \|\zeta'/\sqrt{1-x^2}\|_1$.

This heuristic captures the qualitative behavior of the error. We present the derivation of Heuristic 4.1 in Appendix A. This heuristic is insufficient for direct application to Eq. (7). As $\mathbf{x} \rightarrow \hat{\Gamma}$, the value of k required in Heuristic 4.1 grows rapidly due to growing higher order derivatives of the integrand. Such large values of q and k imply that smooth quadrature rules are cost-prohibitive; this is the problem that singular/near-singular quadrature schemes like hedgehog aim to address. Moreover, this estimate is too loose to determine whether hedgehog or smooth quadrature is required to evaluate the potential. The assumption in Section 3.2.2 addresses this problem by providing a cheap, reasonably robust criterion for refinement that is motivated by existing analyses [3,6] instead of relying on Heuristic 4.1.

4.2. Extrapolation error

A reasonable critique of hedgehog is its reliance on an equispaced polynomial interpolant to extrapolate values of u to the target point. Despite using the first-kind barycentric interpolation formula [71], polynomial interpolation and extrapolation in equispaced points is well-known for an exponentially growing Lebesgue constant and poor stability properties as the number of points p increases [51,66]. Recently [20] demonstrated stable extrapolation in equispaced $p+1$ points using least-squares polynomials of degree \sqrt{p} . However, these results are asymptotic in nature and don’t tell the full story for small to moderate values of p , as in the hedgehog context.

We begin our discussion with a simple representative experiment in equispaced extrapolation. Fig. 6 depicts a minimal extrapolation setup in 3D of a simple singular function $\mu(t) = 1/\|t - q\|$ along a line, with $q = (\rho, 0, 0)$ and $\rho = -1$. We

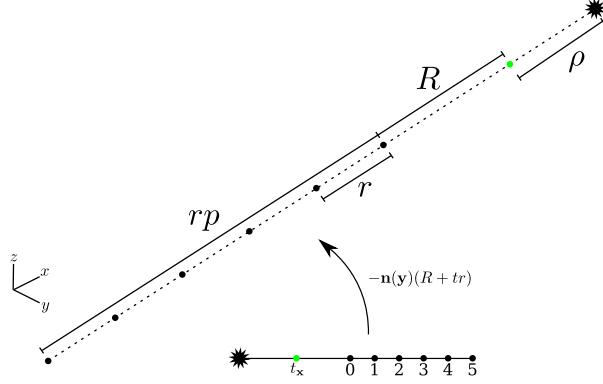


Fig. 6. DIAGRAM OF EXTRAPOLATION SETUP. The toy setup used to study the extrapolation error of a singular function. We choose a simple point singularity $\mu(t) = \frac{1}{\|t-q\|}$ where $q = (\rho, 0, 0)$ (black star) with $\rho = -1$. We choose samples at the points $t_i = (R + ir, 0, 0)$ for $i = 0, \dots, p$ (black dots) and extrapolate the values $\mu(t_0), \dots, \mu(t_p)$ to $t = 0$ (green dot).

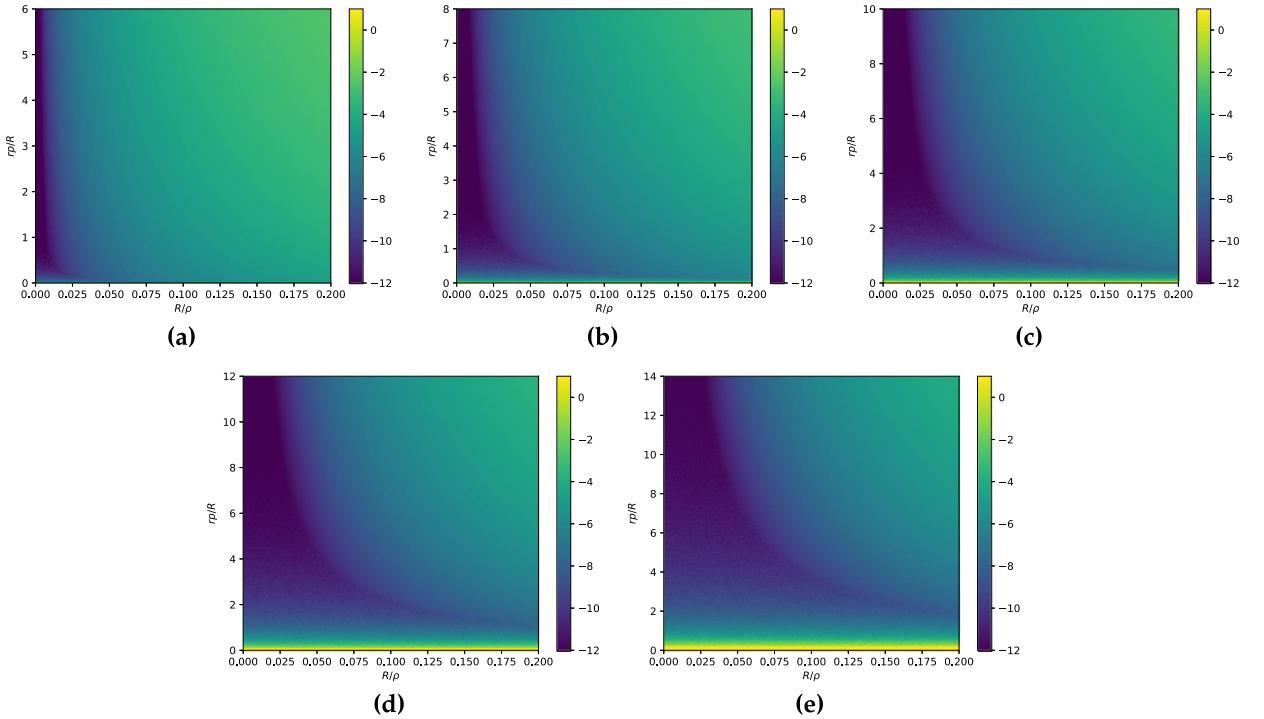


Fig. 7. EMPIRICAL EXTRAPOLATION ERROR BEHAVIOR. We sweep over a range of R and r values to vary Fig. 6 and plot the log of the relative error in Figs. 7a to 7e, for values $p = 6, 8, 10, 12, 14$, in increasing order, from (a) to (e). In these figures, the x-axis is the extrapolation distance R normalized by ρ and the y-axis is the ratio rp/R . The top of the y-axis corresponds to $r = R$; $rp/R = 1$ corresponds to our choice of the parameter a . Assuming that $\rho = O(L)$, $r/R = a/b$ and $R/\rho = b/\lambda$ for some constant λ .

extrapolate exact values of μ from p points, located at $t_i = (R + ir, 0, 0)$, to the origin. This closely mimics the worse-case extrapolation error in 1D of a function analytic in a Bernstein ellipse with a real axis intercept of $\rho + R + rp/2$. We repeat this for a large range of values of r and R for various values of p . The log of the relative error is plotted in Figs. 7a to 7e as a function of the relative extrapolation interval size rp/R and the scaled extrapolation distance R/ρ .

As mentioned in [53, Section 3.4], the adaptive refinement of $\mathcal{P}_{\text{coarse}}$ resolves the boundary data f , and therefore u and ϕ , on the length scale L of the patch. This means we can reasonably assume that the distance of the nearest singularity is $O(L)$ from $\hat{\Gamma}$, i.e., $\rho = \lambda L$ for some λ . In the context of hedgehog, we know that $R = bL(P)$ and $r = aL(P)$. Figs. 7a to 7e are a study of extrapolation error as a function of a/b , b/λ and p .

There are several important observations to make from these plots:

- Extrapolation error decreases as R/ρ decreases, as expected.

- For a fixed value of R/ρ , the extrapolation error *decreases* rapidly as rp decreases, up to a certain value r^*p . This is somewhat counterintuitive, since this means placing points closer together and extrapolating a further distance relative to rp . For a fixed p in exact arithmetic, letting the interpolation interval size tend to zero produces an order p Taylor expansion of the solution u centered at the interval's origin, which accounts for this phenomenon.
- Beyond r^*p , the extrapolation error *increases*. The effects of finite precision eventually pollute the convergence behavior described above. Moreover, the spacing r^* appears to be a function of p . For $p = 6$, r can be reduced to $1/p$ without any numerical issues, but by $p = 14$, only $r > \frac{1}{2}$ is a safe choice for extrapolation.

We do not aim to rigorously analyze these phenomena in this work. We highlight them to provide empirical evidence that equispaced extrapolation is a reasonable, but not optimal, choice for our problem of singular/near-singular integration and to provide some intuition for our parameter choices.

The following simple result describes the behavior of the extrapolation error in Eq. (17).

Theorem 4.2. *Let $u(\mathbf{c}(t))$ be the solution to Eq. (2) given by Eq. (4), restricted to the line $\mathbf{c}(t)$ in 3D intersecting \mathbf{x} , let $\mathbf{c}(t)$ be given by*

$$\mathbf{c}(t) = \mathbf{s}_x - (R + tr)\mathbf{n}(\mathbf{s}_x), \quad (21)$$

where \mathbf{s}_x is the closest point on $\hat{\Gamma}$ to \mathbf{x} , $R = bL_{s_x}$, $r = aL_{s_x}$, $\mathbf{n}(\mathbf{s}_x)$ is the outward surface normal at \mathbf{s}_x , and let $|u^{(p)}(\mathbf{c}(t))|$ be bounded above by C_p on the interval $[-R, R + pr]$. Let $\mathfrak{P}(t)$ be the p -th order polynomial interpolant of $u(\mathbf{c}(t))$ constructed from the check points $\mathbf{c}_0, \dots, \mathbf{c}_p$, where $\mathbf{c}_i = \mathbf{c}(i)$. Then the extrapolation error associated with hedgehog behaves according to:

$$|u(\mathbf{c}(t_x)) - \mathfrak{P}(t_x)| \leq \frac{C_p}{(p+1)!} |R + rp|^p = \frac{C_p}{(p+1)!} |b + ap|^p \cdot |L|^p, \quad (22)$$

where $t_x = \frac{\|\mathbf{x} - \mathbf{s}_x\| - R}{r}$.

Proof. We know that for a smooth function f and points x_0, \dots, x_p in a 1D interval I_0 , for some $\xi \in I_0$, the following relation holds for all $x \in I_0$:

$$f(x) - \mathfrak{P}(x) = \frac{f^{(p)}(\xi)}{(p+1)!} \prod_{i=0}^p (x - x_i). \quad (23)$$

Let \mathfrak{P} be the p th order polynomial interpolating the points x_0, \dots, x_p . In the hedgehog setup, since $R + rp$ is the distance of the furthest check point to \mathbf{y} , we know that $x - x_i < R + rp$ for each i . Since $f(t) = u(\mathbf{c}(t))$ is harmonic, and therefore C^∞ , in Ω , $|f^{(p)}(\xi)|$ can be uniformly bounded on I_0 by some constant C_p . Noting that $R = bL$ and $r = aL$ yields our result. \square

For fixed values of a and b , as we let $L \rightarrow 0$, the extrapolation error is bounded by $O(L^p)$. In practice, however, this means that we can choose a and b to minimize the constant factor $|b + ap|^p$ in Theorem 4.2. Since $p > 1$, a must be chosen to balance out the contribution of p , yet our extrapolation study shows that we can't simply set $a = 0$. We therefore choose $a \leq 1/p$ for $p = 6$ and 8, motivated by Figs. 7a and 7b. Moreover, since $b < 1$, we can choose $a \leq b/p$, which allows a and b to decay at the same rate. The advantage of choosing $a \leq b/p$ is that b is a single parameter that controls the accuracy of hedgehog. Since we have fixed the quadrature order $q = 20$ to satisfy the assumption in Section 3.2.2, a smaller value of b will trigger more upsampling in Algorithm 3, keeping quadrature error fixed while reducing extrapolation error.

It is important to keep in mind that Theorem 4.2 only provides insight for moderate values of p ; our conclusions are largely irrelevant for large p . We use $p = 6$ and $a = b/6$, leaving the construction of an optimal extrapolation scheme to future work.

4.3. Limitations

Our error discussion reveals several limitations of our method. The first and most apparent shortcoming is that extrapolation instability fundamentally limits convergence order. However, for reasonable orders of convergence, up to 14, we have discussed an empirical scheme to choose parameters to maximize the available convergence behavior. Moreover, low-order surface geometries used in engineering applications will likely limit the convergence rate before it is limited by the extrapolation order, making this a non-issue in practical scenarios.

Another downside of the chosen extrapolation approach is lack of direct extension of hedgehog to oscillatory problems like the Helmholtz equation. Due to the limitation on the values of p , we can't guarantee the ability to resolve high-frequency oscillations in the solution. A new extrapolation procedure is required to do so robustly without compromising efficiency.

In [68], the authors demonstrate a relationship between the truncation error of a QBX expansion and the local curvature of $\hat{\Gamma}$. Our scheme also is susceptible to this form of error and we do not address nor analyze this in this work. This is a subtle problem that requires a detailed analysis of the surface geometry with respect to the chosen extrapolation scheme.

Another limitation is the lack of an accurate error estimate to serve as an upsampling criteria in place of the criteria in Section 3.2.2, such as [35]. Extending [35] to 3D surfaces is non-trivial and whether the size of $\mathcal{P}_{\text{fine}}$ would be reduced enough to outweigh the added cost of the additional Newton iterations required by their scheme remains to be seen.

Finally, for certain accuracy targets and geometries, the algorithm above may lead to an impractically high number of patches in $\mathcal{P}_{\text{coarse}}$ and $\mathcal{P}_{\text{fine}}$. Geometries with nearly-touching non-local regions, as shown in Fig. 12, will see large amounts of refinement. If the nearly-touching embeddings γ_r are close enough, i.e., less than 10^{-10} apart, there is little hope of an accurate solution with a fixed computational budget. We allow the user to enforce a minimal patch size L_{\min} , limiting the time and memory consumption at the expense of not reaching the requested target accuracy.

5. Complexity

In this section, we summarize the complexity of the algorithms required by hedgehog. We present a detailed complexity analysis in [46, Section 3]. The input to our overall algorithm is a domain boundary Γ with N_{init} patches and boundary condition f . The parameters that directly impact complexity are:

- The number of patches N after admissibility refinement. This is a function of N_{init} , the geometry of Γ , the definition of f , and the choices of parameters a and b in check point construction.
- Quadrature order q and the degree of smoothness k of Γ and f . We assume that k is sufficiently high to obtain optimal error behavior for a given q by letting $k = 2q$ in Eq. (20).
- hedgehog interpolation order p .
- The numbers of evaluation points in different zones \mathcal{N}_{far} , $\mathcal{N}_{\text{inter}}$, and $\mathcal{N}_{\text{near}}$, with $\mathcal{N}_{\text{tot}} = \mathcal{N}_{\text{far}} + \mathcal{N}_{\text{inter}} + \mathcal{N}_{\text{near}}$.

The complexity is also affected by the geometric characteristics of Γ as described in [46, Section 3].

- *Admissibility.* The complexity of this step is $O(N_{\text{init}} \log N_{\text{init}})$, with constants dependent on α_0 , β_0 and C_J . The logarithmic factor is due to use of an AABB tree for closest surface point queries.
- *Upsampling.* The complexity of upsampling is $O(\hat{m}N \log(N))$, where \hat{m} is the largest upsampling ratio. The logarithmic factor appears for similar reason to admissibility, with constants that depend on geometric parameters and the boundary condition through the error estimate of Section 4. We show that the upsampling ratio is independent of N in [46, Section 3].
- *Point marking.* Identifying which zone an evaluation point belongs to (Ω_F , Ω_I or Ω_N) depends on N and the total number of points to be classified $\mathcal{N}_{\text{tot}} = \mathcal{N}_{\text{far}} + \mathcal{N}_{\text{inter}} + \mathcal{N}_{\text{near}}$. The complexity is $O(\mathcal{N}_{\text{tot}} \log N)$ with constants dependent on geometric parameters, due to the cost of closest surface point queries.
- *Far, intermediate and near zone integral evaluation.* The complexity of these components depends on N and \mathcal{N}_{far} , $\mathcal{N}_{\text{inter}}$ and $\mathcal{N}_{\text{near}}$ respectively, with the general form $O(s_1 N + s_2 \mathcal{N}')$, where \mathcal{N}' is the number of evaluation points in the corresponding class. For the far field, $s_1 = s_2 = 1$. For the intermediate evaluation, $s_1 = \hat{m}q^2$ and $s_2 = 1$; finally, for the near zone, $s_1 = \hat{m}q^2$ and $s_2 = p$. If b is chosen appropriately, the intermediate and near zone error is ϵ_{target} .
- *GMRES solve.* Due to the favorable conditioning of the double-layer formulation in Eq. (5), GMRES converges rapidly to a solution in a constant number of iterations for a given Γ that is independent of N . This means that the complexity to solve Eq. (5) is asymptotically equal (up to a constant dependent on Γ) to the complexity equal to a near-zone evaluation with $\mathcal{N}_{\text{near}} = N(q+1)^2$.
- *Evaluation on uniform point distribution* In many applications, one would like the value of the solution u due to a density ϕ at a collection of points uniformly distributed throughout the domain Ω . When the number of such targets is chosen to match the resolution of the surface discretization, the overall complexity of solution evaluation is $O((\hat{m} + \hat{m}q^2)q^2N + N^{3/2})$.

6. Results

We now demonstrate the accuracy and performance of hedgehog to evaluate singular/near-singular layer potentials on various complex geometries to solve the integral equation in Eq. (5) and evaluate the solution as defined in Eq. (4).

6.1. Classical convergence with patch refinement

We will first demonstrate the numerical convergence behavior of hedgehog. As discussed in [36, Section 3.1], approximation-based schemes such as hedgehog do not converge classically but do so up to a controlled precision if r and R scale with proportional to the patch size. In order to observe classical convergence as we refine $\mathcal{P}_{\text{coarse}}$, we must allow R and r to decrease slower than $O(L)$, such as with rate $O(\sqrt{L})$. In this section, we choose the hedgehog parameters a and b proportional to $1/\sqrt{L}$ to achieve this and demonstrate numerical convergence with refinement of L .

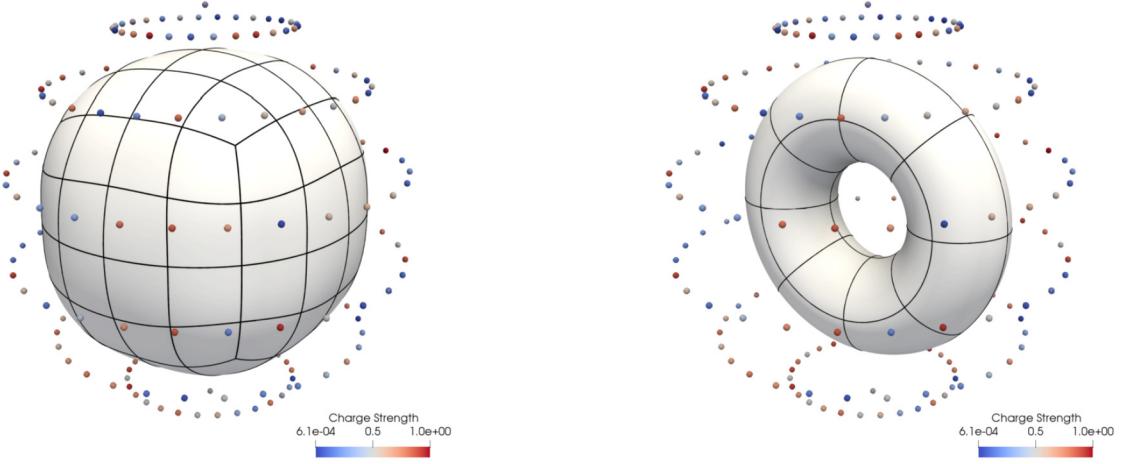


Fig. 8. Geometry and singularities used for Green's Identity convergence tests. Shown are polynomial patches defining boundary geometry (black lines) and point singularities placed on the surface on a sphere of radius one. Singularity strengths are randomly selected values in $[0, 1]$; shown is the strength intensity for Laplace problems, which varies from blue to red. We use 96 20th-order polynomial patches for the spheroid (left) and 32 cubic patches for the torus (right).

Table 1

ℓ^∞ RELATIVE ERROR IN GREEN'S IDENTITY VERSUS NUMBER OF PATCHES. The solution to Eq. (2) due to a known function u_c , shown in Fig. 8 is computed via Green's Identity. We evaluate the single- and double-layer potentials with hedgehog due to the Dirichlet and Neumann boundary data and compare against the known value of u_c on the boundary. Each column is the result of an additional level of uniform quadrissection of the patches in $\mathcal{P}_{\text{coarse}}$. The final column (EOC) is the estimated convergence order, computed via least-squares log-log fit of the error as a function of max patch size.

Geometry	PDE	Relative ℓ^∞ error (Number of patches)				EOC
Spheroid (Fig. 8-left)	Laplace	1.06×10^{-4} (96)	4.78×10^{-6} (384)	9.14×10^{-8} (1536)	4.35×10^{-9} (6144)	4.77
	Elasticity	1.68×10^{-3} (96)	6.94×10^{-5} (384)	1.53×10^{-6} (1536)	1.33×10^{-8} (6144)	5.74
	Stokes	1.92×10^{-3} (96)	7.95×10^{-5} (384)	1.74×10^{-6} (1536)	1.53×10^{-8} (6144)	5.72
Torus (Fig. 8-right)	Laplace	2.05×10^{-3} (32)	7.52×10^{-5} (128)	3.79×10^{-6} (512)	8.48×10^{-8} (2048)	5.45
	Elasticity	4.38×10^{-2} (32)	1.17×10^{-3} (128)	5.08×10^{-5} (512)	1.42×10^{-6} (2048)	5.09
	Stokes	5.03×10^{-2} (32)	1.33×10^{-3} (128)	5.81×10^{-5} (512)	1.65×10^{-6} (2048)	5.09

In our examples, we use analytic solutions to Eq. (2) obtained as sums of point charge functions of the form

$$u_c(\mathbf{x}) = \sum_{i=1}^m G(\mathbf{x}, \mathbf{y}_i) \psi_i \quad (24)$$

where the charge locations \mathbf{y}_i with strengths ψ_i are outside of Ω . To construct specific solutions, we sample a sphere of radius one with point charges, as shown in Figs. 8 and 9. We choose charge strengths ψ_i randomly from $[0, 1]^d$, where $d = 1$ for Laplace problems and $d = 3$ for Stokes and elasticity problems.

We use the multipole order $m = 20$ with 5000 points per leaf box for the kernel-independent FMM. This ensures that the FMM error does not dominate; sufficiently large number of points per leaf box is needed to minimize the additional error due to tree depth. We choose a high quadrature order $q = 20$, or 400 quadrature points per patch in $\mathcal{P}_{\text{coarse}}$, relative to overall convergence order to satisfy the assumption in Section 3.2.2. We also use two levels of uniform upsampling to demonstrate convergence.

6.1.1. Green's identity

We report the accuracy of the hedgehog evaluation scheme in Table 1, where we verify Green's Identity for a random known function u_c in Eq. (24). We evaluate the Dirichlet and Neumann boundary data due to u_c at the discretization points of $\hat{\Gamma}$ and use one-sided hedgehog to evaluate the corresponding single- and double-layer potentials at the same discretization points. With each column of Table 1, we subdivide $\mathcal{P}_{\text{coarse}}$ to more accurately resolve the boundary condition. The error shown in Table 1 is the ℓ^∞ -relative error in the solution value

$$\frac{\left\| \hat{S} \left[\frac{\partial u_c}{\partial \mathbf{n}} \right] (\mathbf{x}) - \hat{D} [u_c] (\mathbf{x}) - u_c (\mathbf{x}) \right\|_\infty}{\| u_c \|_\infty}, \quad (25)$$

where \hat{S} and \hat{D} are the single- and double-layer singular integral operators discretized and evaluated with hedgehog. In these tests, we choose $p = 6$, $r = .004\sqrt{L}$ ($a = .004/\sqrt{L}$) and $R = .03\sqrt{L}$ ($b = .03/\sqrt{L}$). We observe roughly 5th order

Table 2

PERFORMANCE OF SINGULAR EVALUATION IN GREEN'S IDENTITY. For each test in Table 1, we report the number of target points evaluated with one-sided hedgehog per second per core.

Geometry	PDE	Target points/second/core			
Spheroid (Fig. 8-left)	Laplace	3684	5438	5077	5629
	Elasticity	1325	1731	1687	1790
	Stokes	1635	2075	2016	2120
Torus (Fig. 8-right)	Laplace	2729	3373	4564	5477
	Elasticity	984	1171	1347	1502
	Stokes	1134	1331	1609	1727

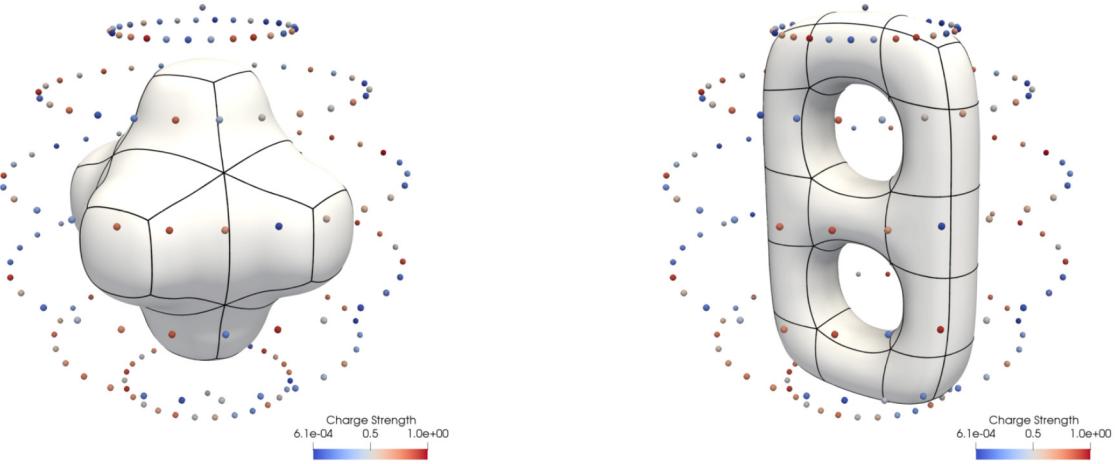


Fig. 9. Geometry and singularities used for solver convergence tests. Figures are similar to Fig. 8, but displaying geometries for testing the convergence of hedgehog within a GMRES solver. We use 30 16th-order polynomial patches for the pipe (left) and 50 20th-order patches for the genus two surface (right). Note the proximity of the singularities to the domain of the genus two surface; the nearest singularity is less than $.05L$ from $\hat{\Gamma}$.

Table 3

ℓ^∞ RELATIVE ERROR IN GMRES SOLVE AND SOLUTION EVALUATION VERSUS NUMBER OF PATCHES. We solve Eq. (2) by discretizing and evaluating the layer potential in the integral equation in Eq. (5) as described in Section 3.1. We use two-sided hedgehog inside of GMRES to solve for ϕ , then evaluate Eq. (8) with one-sided hedgehog at a new set of points on $\hat{\Gamma}$. Each column is the result of an additional level of uniform quadrisecion of the patches in $\mathcal{P}_{\text{coarse}}$. The final column (EOC) is the estimated convergence order, computed via least-squares log-log fit of the error as a function of max patch size.

Geometry	PDE	Relative ℓ^∞ error (Number of patches)				EOC
Spheroid (Fig. 8-left)	Laplace	2.70×10^{-6} (96)	1.92×10^{-7} (384)	4.47×10^{-9} (1536)	5.13×10^{-11} (6144)	5.35
Pipe (Fig. 9-left)	Laplace	5.99×10^{-4} (30)	3.03×10^{-5} (120)	6.68×10^{-7} (480)	2.27×10^{-8} (1920)	5.92
	Elasticity	7.17×10^{-2} (30)	3.57×10^{-3} (120)	8.90×10^{-5} (480)	4.14×10^{-6} (1920)	5.45
	Stokes	8.53×10^{-2} (30)	4.12×10^{-3} (120)	1.03×10^{-4} (480)	4.73×10^{-6} (1920)	5.43
Genus 2 (Fig. 9-right)	Laplace	4.00×10^{-2} (50)	1.25×10^{-4} (200)	1.54×10^{-6} (800)	5.73×10^{-10} (3200)	8.76
	Elasticity	9.20×10^{-2} (50)	1.05×10^{-3} (200)	1.00×10^{-5} (800)	9.44×10^{-8} (3200)	6.89
	Stokes	1.03×10^{-1} (50)	1.18×10^{-3} (200)	1.15×10^{-5} (800)	1.03×10^{-7} (3200)	6.88

convergence on both the spheroid and torus test geometries in Fig. 8 for each of the tested PDE's. In Table 2, we present the number of target points evaluated per second per core with one-sided hedgehog. We see that performance is best for Laplace and worst for elasticity problems, as expected.

6.1.2. Solution via GMRES

We report the accuracy of the hedgehog scheme when used to solve Eq. (2) via the integral equation in Eq. (5). Two-sided hedgehog is used in the matrix-vector multiply inside GMRES to solve Eq. (5) for the values of the density ϕ at the discretization points. Then one-sided hedgehog is used to evaluate Eq. (8) at a slightly coarser discretization. Since GMRES minimizes the residual at the original discretization of Eq. (5), this final step prevents an artificially accurate solution by changing discretizations. Table 3 lists the ℓ^∞ relative error values for the total solve and evaluation steps using Section 3.1 as we refine $\mathcal{P}_{\text{coarse}}$ by subdivision as in the previous section. In these tests, we choose $p = 6$, $r = .005\sqrt{L}$ ($a = .005/\sqrt{L}$), and $R = .03\sqrt{L}$ ($b = .03/\sqrt{L}$). As for previous examples, we observe at least 5th order convergence on all tested geometries in Fig. 9 and Fig. 8-left and all PDE's. We include the spheroid example as an additional demonstration of a high accuracy solution via GMRES with our approach. We report the number of target points evaluated per second per core with two-

Table 4

PERFORMANCE OF SINGULAR EVALUATION IN GMRES MATRIX-VECTOR MULTIPLY. For each test in Table 3, we report the number of target points per second per core evaluated with two-sided hedgehog in a single GMRES matrix-vector multiplication.

Geometry	PDE	Target points/second/core			
Spheroid	Laplace	2737	3149	2846	2950
Pipe (Fig. 8-left)	Laplace	3046	2178	2832	2982
	Elasticity	991	993	1189	1261
	Stokes	1048	1140	1335	1422
Genus 2 (Fig. 8-right)	Laplace	1862	2886	3122	2879
	Elasticity	729	1125	1255	1295
	Stokes	929	1304	1450	1504

sided hedgehog in Table 4. The results are similar to Table 2; the slower performance is because evaluation via two-sided hedgehog is more expensive than one-sided hedgehog.

6.2. Comparison with [75]

In this section, we compare our method to [75], a previously proposed high-order, kernel-independent singular quadrature method in 3D for complex geometries. These characteristics are similar to hedgehog shares these characteristics. [46, Section 4] presents additional comparisons.

The metric we are interested is *cost for a given relative error*. Assuming the surface discretization is $O(N)$, we measure the cost of a method as its total wall time during execution T divided by the total wall time of an FMM evaluation on the same $O(N)$ discretization, T_{FMM} . By normalizing by the FMM evaluation cost, we minimize the dependence of the cost on machine- and implementation-dependent machine-dependent parameters.

We run the tests in this section on the spheroid geometry shown in Fig. 8-left. We focus on the singular quadrature scheme of [75]. The near-singular quadrature of [75] is algorithmically similar to hedgehog, but since an expensive singular quadrature rule is used as a part of near-singular evaluation, it has a higher total cost. As a result, the accuracy and cost of near-singular evaluation of [75] is bounded by the accuracy and cost of the singular integration scheme.

To compare the full hedgehog method with [75], we fit polynomial patches to the C^∞ surface of [76], denoted Γ_b , to produce $\hat{\Gamma}$ during the first step of Section 3.4. We apply the remaining geometry preprocessing algorithms of Section 3.4 to $\hat{\Gamma}$ to produce $\mathcal{P}_{\text{coarse}}$. After producing $\mathcal{P}_{\text{fine}}$ with two levels of uniform upsampling, we solve Eq. (5) with two-sided hedgehog on $\hat{\Gamma}$ and evaluate the solution on the boundary with one-sided hedgehog. We then solve for the solution to Eq. (5) on Γ_b using [75].

For each of the tests in this section, we choose some initial spacing parameter h_0 to discretize the surface of [76], as in [75], and use the $16 \times$ upsampled grid and floating partition of unity radius proportional to $O(\sqrt{h})$, as in the original work. We apply hedgehog to $\hat{\Gamma}$ and the scheme of [75] to Γ_b with spacing $h_0/2^i$, for $i = 1, \dots, 4$.

As in the previous section, we choose the parameters r and R of hedgehog to be $O(\sqrt{L})$. For both quadrature methods, we use a multipole order of 16 for PVFMM with at most 250 points in each leaf box. The results are shown in Fig. 10. From left to right, each plot details the total cost of each scheme, the cost of each subroutine for hedgehog (denoted HH) and the singular quadrature scheme of [75] (denoted POU), and the relative error as a function of h and L , respectively, for all refinement levels. We plot the cost of both schemes the cost of each algorithmic step as a function of their computed relative error. In each figure, we present results for a Laplace problem (top) and an elasticity problem (bottom).

In Fig. 10, as expected, we observe a higher convergence rate for hedgehog compared to [75]. [75] outperforms hedgehog in terms of cost for all tested discretizations. We observe that the FMM evaluation in Fig. 10 accounts for at least 95% of the hedgehog cost. This means that a local singular quadrature method (based on corrections to an FMM evaluation, Section 1.2) of worse complexity can beat a global method, simply by virtue of reducing the FMM size. By noting the large difference between the hedgehog FMM cost and the hedgehog density interpolation, we can reasonably infer that a local hedgehog scheme should narrow this performance gap and outperform [75] for larger problems, assuming that switching to a local scheme does not dramatically affect error convergence.

6.3. Requested target precision vs. computed accuracy

In this section, we study the performance of the full algorithm outlined in Section 3. We test hedgehog on the torus domain shown in Fig. 8-right. We choose a reference solution of the form of Eq. (24) with a single point charge located at the origin, in the middle of the hole of the torus. We solve the integral equation with two-sided hedgehog and evaluate the singular integral on a distinct discretization with one-sided hedgehog. We choose $q = 20$, $p = 6$ and $a = b/6$. We select various values for ϵ_{target} using the plot in Fig. 7a to choose b to ensure sufficiently accurate extrapolation. We plot the results of our tests in Fig. 11.

We see in Fig. 11-left that we are consistently close to the requested target precision. We see a decline in target points per second per core as accuracy increases in Fig. 11-middle. This is explained by Fig. 11-right, which shows an increase in

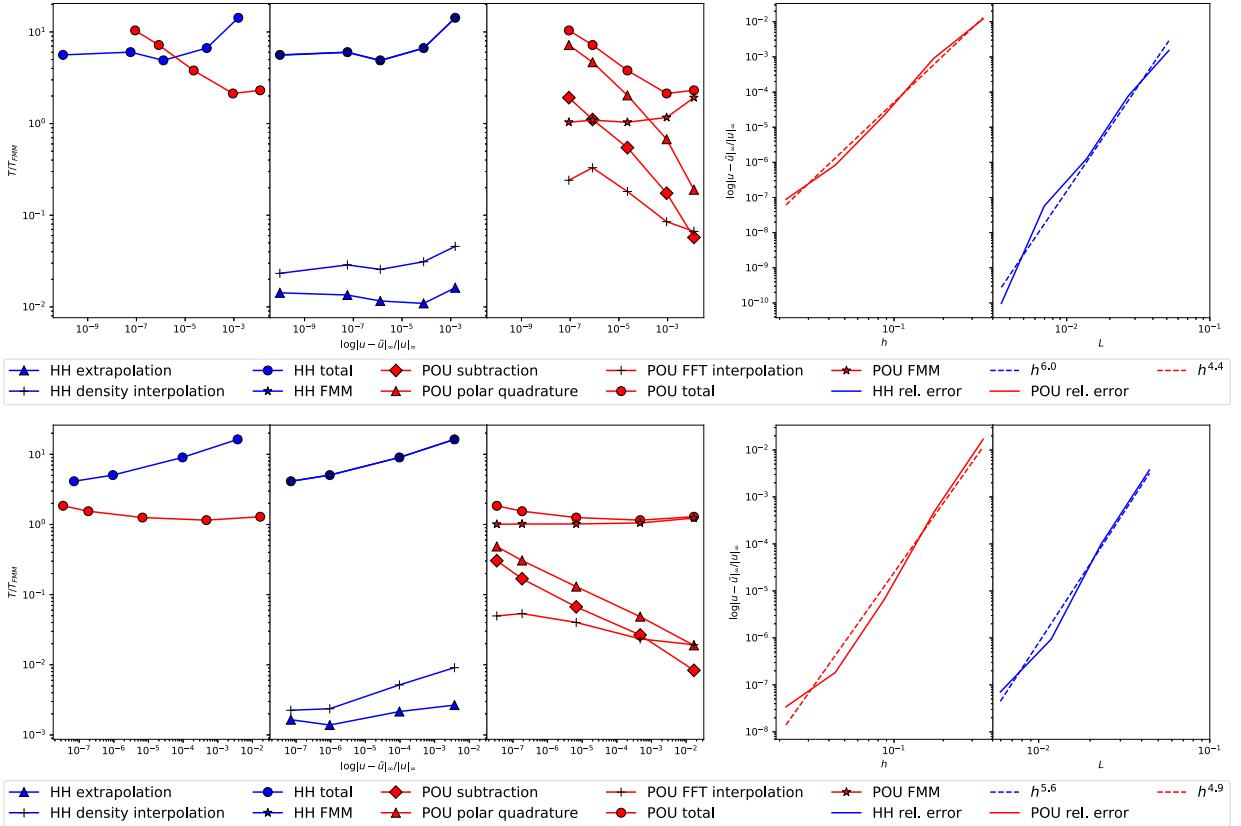


Fig. 10. COMPARISON OF HEDGEHOG ON POLYNOMIAL PATCHES (HH) VERSUS [75] ON THE SURFACE REPRESENTATION OF [75] (POU) SOLVING VIA GMRES FOR u_c . Laplace (top) and elasticity (bottom) problems solved on the spheroid shown in Fig. 8. From left to right, we plot the total cost of each scheme, the cost of each subroutine for hedgehog (blue) and the singular quadrature scheme of [75] (red), and the relative error as a function of h . We plot error convergence of [75] as a function of h and hedgehog as a function of L , due to the distinct discretizations. For hedgehog parameters, we choose $r = .013\sqrt{L}$, $R = .075\sqrt{L}$ for the Laplace problem; for the elasticity problem, we choose $r = .013\sqrt{L}$, $R = .08\sqrt{L}$. We choose $p = 6$ and $q = 15$ for both problems. For [75] the spacing is $h_0 = .35$. Note that in the hedgehog timing breakdown, since the FMM time is dominant, the FMM cost lies directly on top of the total cost.

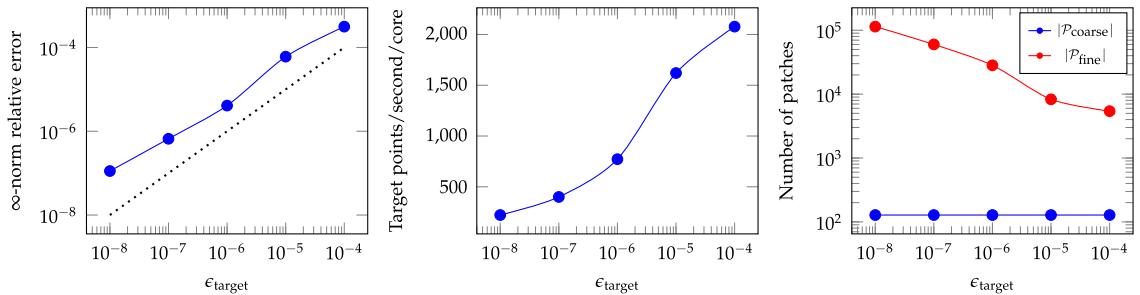


Fig. 11. PERFORMANCE OF FULL ALGORITHM. Left: ∞ -norm relative error in singular integral vs requested target accuracy (blue). The dotted line is the ideal behavior $y = x$. Middle: Performance in terms of target points evaluated per second per core with hedgehog. Right: Number of patches in $\mathcal{P}_{\text{coarse}}$ and $\mathcal{P}_{\text{fine}}$ computed by the preprocessing algorithms.

the size $\mathcal{P}_{\text{fine}}$ as $\mathcal{P}_{\text{coarse}}$ remains a fixed size. The initial 128 patches in $\mathcal{P}_{\text{coarse}}$ are enough to resolve the boundary condition and Γ , but we need greater quadrature accuracy for lower values of ϵ_{target} . Decreasing the number of points in passed to the FMM, i.e., decreasing the size of $\mathcal{P}_{\text{fine}}$, is the main way to improve performance of our method. This is further indication that a local version of hedgehog will outperform a global approach.

6.4. Full algorithm on interlocking torii

We now demonstrate the full algorithm pipeline on an exterior Laplace problem, whose boundary is defined by four interlocking torii shown in Fig. 12. The domain boundary is contained in the box $[-3.8, 2.4] \times [-1.1, 1.1] \times [-1, 1]$. The

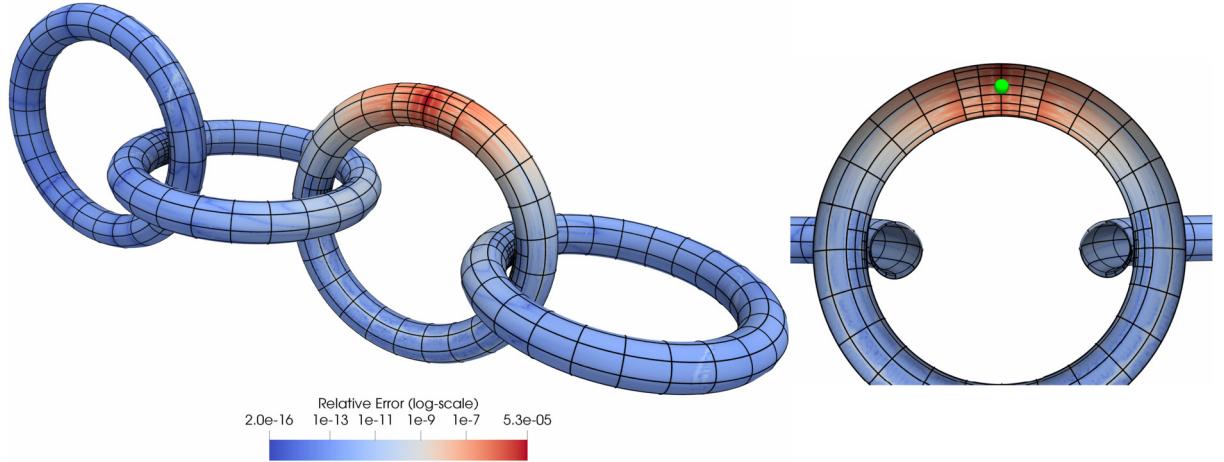


Fig. 12. ABSOLUTE ERROR OF GMRES SOLVE VIA HEDGEHOG ON INTERLOCKING TORII. Left: The admissible set of 1128 patches in $\mathcal{P}_{\text{coarse}}$ used to solve Eq. (5) is shown (black lines denote patch boundaries). The point charge generated the boundary condition is located within the second torus from the right. Right: a cross-section of the torii geometry through the xz -plane, showing the second torus from the right and the location of the singularity (green point).

shortest distance between two adjacent torii is less than 10% of a polynomial patch length defining the boundary. We again use a boundary condition of the form Eq. (24) with a single point charge located at $(0, .03, .875)$, inside the upper half of the second torus from the right in Fig. 12. This problem is challenging due to the nearly touching geometry of the torii, along with the singularity placed close to the boundary. We run the admissibility and adaptive upsampling algorithms outlined in Section 3, solve Eq. (5) using two-sided hedgehog, and evaluate the solution on the boundary using one-sided hedgehog. The absolute error in the ∞ -norm of the singular evaluation is plotted on the boundary surface.

Using $a = .1$, $b = .025$, $p = 6$ and $q = 20$, we achieve a maximum pointwise error of 1.29×10^{-5} . GMRES was able to reduce the residual by a factor of 10^{-13} over 109 iterations. There are 288768 quadrature points in the coarse discretization, 18235392 quadrature points in the fine discretization, and 3465216 check points used in the two-sided hedgehog evaluation inside GMRES. We evaluate the solved density at 451200 points on the boundary with one-sided hedgehog to produce the render in Fig. 12. On a machine with two Intel Xeon E-2690v2 3.0 GHz CPU's, each with 10 cores, and 100 GB of RAM, the GMRES solve and interior evaluation required 5.7 hours and can evaluate the singular integral at a rate of 1709 target points per second per core.

6.5. Solution on complex geometry

We have demonstrated in [42] a parallel implementation of Section 3.1, applied to simulating red blood cell flows. The surface geometry of the blood vessel shown in Fig. 13 is complex, with rapidly varying curvatures and geometric distortions due to singular vertices in the surface mesh. Since the surface is admissible, we are able to apply parallel hedgehog directly without geometric preprocessing to solve an interior Dirichlet Stokes problem. We use $a = .125$, $b = .125$, $p = 6$ and $q = 16$ as simulation parameters.

Using 32 machines each with twenty 2.6 GHz cores with 100 GB of RAM, we achieve a maximum pointwise error of 3×10^{-6} when solving a Stokes problem with constant density. We then place a random vector point charge two patch lengths away (relative to the patches in $\mathcal{P}_{\text{coarse}}$) from the domain boundary (on the left side of Fig. 13, solve Eq. (5) using two-sided hedgehog, and evaluate the solution on the boundary using one-sided hedgehog. The absolute error in the ∞ -norm of the singular evaluation is plotted on the boundary surface. There are 10,485,760 quadrature points in the coarse discretization, 167,772,160 quadrature points in the fine discretization, and 125,829,120 check points used in the two-sided hedgehog evaluation inside GMRES. We evaluate the solved density at 209,715,200 points on the boundary with one-sided hedgehog to produce the render in Fig. 12. We achieve a maximum pointwise error of 1.8×10^{-2} and can evaluate the singular integral at rate of 3529 target points per second per core.

7. Conclusion

We have presented hedgehog, a fast, high-order, kernel-independent, singular/near-singular quadrature scheme for elliptic boundary value problems in 3D on complex geometries defined by piecewise tensor-product polynomial surfaces. The primary advantage of our approach is *algorithmic simplicity*: the algorithm can be implemented easily with an existing smooth quadrature rule, a point FMM and 1D and 2D interpolation schemes. We presented fast geometry processing algorithms to guarantee accurate singular/near-singular integration, adaptively upsample the discretization and query local surface patches. We then evaluated hedgehog in various test cases, for Laplace, Stokes, and elasticity problems on various patch-based geometries and compared our approach with [75].

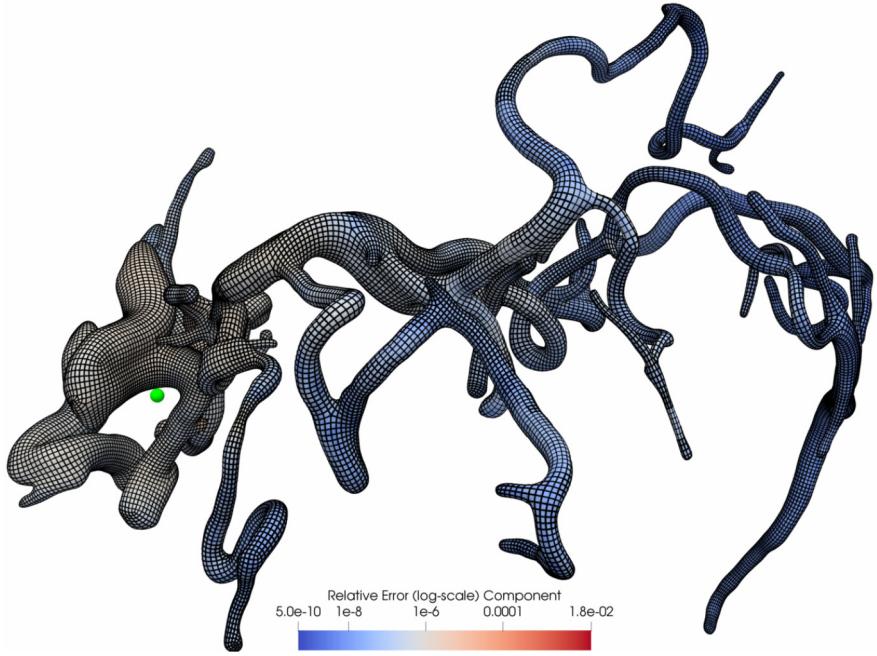


Fig. 13. ABSOLUTE ERROR OF GMRES SOLVE VIA HEDGEHOG ON COMPLEX BLOOD VESSEL GEOMETRY USED IN [42]. The blood vessel uses 40,960 8th order polynomial patches (black edges denote patch boundaries). The geometry is admissible by construction. The point charge is located on left side of the figure (green).

[42] demonstrates a parallel implementation of hedgehog, but the geometric preprocessing and adaptive upsampling algorithms presented in Section 3 are not parallelized. This is a requirement to solve truly large-scale problems that exist in engineering applications. Our method can also be easily restructured as a local method. The comparison in Section 6.2 highlights an important point: a local singular quadrature method can outperform a global method for moderate accuracies, *even when the local scheme is asymptotically slower*. This simple change can also dramatically improve both the serial performance and the parallel scalability of hedgehog shown in [42], due to the decreased communication of a smaller parallel FMM evaluation. The most important improvement to be made, however, is the equispaced extrapolation. Constructing a superior extrapolation procedure, optimized for the boundary integral context, is the main focus of our current investigations.

CRediT authorship contribution statement

Matthew J. Morse: Conceptualization, Data curation, Formal analysis, Investigation, Methodology, Project administration, Resources, Validation, Writing – original draft, Writing – review & editing. **Abtin Rahimian:** Conceptualization, Funding acquisition, Methodology, Project administration, Supervision. **Denis Zorin:** Conceptualization, Funding acquisition, Methodology, Project administration, Resources, Supervision, Writing – original draft, Writing – review & editing.

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 Michael O'Neil, Dhairyा Malhotra, Libin Lu, Alex Barnett, Leslie Greengard, Michael Shelley for insightful conversations, feedback and suggestions regarding this work. We would also like to thank the NYU HPC team, and Shenglong Wang in particular, for great support throughout the course of this work, and the helpful feedback of the anonymous reviewers. This work was supported by NSF grant DMS-1821334.

Appendix A. Derivation of Heuristic 4.1

We are interested in computing the error incurred when approximating a 2D surface integral with an interpolatory quadrature rule. In 1D on the interval $[-1, 1]$, we're interested in the quantity

$$R_q[f] = I[f] - Q_q[f] \quad (\text{A.1})$$

where

$$I[f] = \int_{-1}^1 f(x) dx \quad (A.2)$$

$$Q_q[f] = \sum_{i=0}^q f(x_i) w_i, \quad (A.3)$$

for quadrature weights w_i for a q -point quadrature rule. For a 2D double integral, we define a similar relationship between the remainder, the exact integral and the q th order quadrature rule:

$$R_q^{(2)}[f] = I^{(2)}[f] - Q_q^{(2)}[f] \quad (A.4)$$

where

$$I^{(2)}[f] = \int_{-1}^1 \int_{-1}^1 f(s, t) ds dt \quad (A.5)$$

$$Q_q^{(2)}[f] = \sum_{j=0}^q \sum_{i=0}^q f(s_i, t_j) w_i w_j. \quad (A.6)$$

For a function of two variables $f(s, t)$, we will denote $I_s[f] = \int_{-1}^1 f(s, \cdot) ds$ as integration with respect to the s variable only, which produces a function of t . The same subscript notation applies to $R_{q,s}[f]$ and $Q_{q,s}[f]$. We use similar notation for t : we apply the 1D functional to the variable in the subscript, producing a 1D function in the remaining variable. We observe that

$$I^{(2)}[f] = \int_{-1}^1 \left(\int_{-1}^1 f(s, t) ds \right) dt = \int_{-1}^1 I_s[f] dt = I_t[I_s[f]] \quad (A.7)$$

Following the discussion in [3], we substitute into Eq. (A.7) and have

$$I^{(2)}[f] = I_t[R_{q,s}[f] + Q_{q,s}[f]] \quad (A.8)$$

$$= R_{q,t}[R_{q,s}[f] + Q_{q,s}[f]] + Q_{q,t}[R_{q,s}[f] + Q_{q,s}[f]] \quad (A.9)$$

$$= R_{q,t}[R_{q,s}[f]] + Q_{q,s}[R_{q,t}[f]] + Q_{q,t}[R_{q,s}[f]] + Q_{q,t}[Q_{q,s}[f]] \quad (A.10)$$

We assume that the higher-order “remainder of remainder” term contributes negligibly to the error. Although it has been shown that this term has a non-trivial contribution to a tight error estimate [23], we are able to provide a sufficiently tight upper bound. For large q , the quadrature rule approaches the value of the integral, i.e., $Q_{q,\beta} \approx I_\beta$ for $\beta = s, t$, we’re left with:

$$I^{(2)}[f] \approx I_s[R_{q,t}[f]] + I_t[R_{q,s}[f]] + Q_q^{(2)}[f], \quad (A.11)$$

and hence:

$$R_q^{(2)}[f] \lesssim I_s[R_{q,t}[f]] + I_t[R_{q,s}[f]], \quad (A.12)$$

where \lesssim means “approximately less than or equal to.” From [64, Theorem 5.1], we recall that for a 1D function θ defined on $[-1, 1]$, if $Q_q[\theta]$ is computed with Clenshaw-Curtis quadrature, θ is C^k and $\|\theta^{(k)}\|_T < V$ on $[-1, 1]$ for real finite V , then for sufficiently large q , the following inequality holds

$$R_q[\theta] \leq \frac{32V}{15\pi k(2q+1-k)^k}, \quad (A.13)$$

where $\|\alpha(x)\|_T = \|\alpha'/\sqrt{1-x^2}\|_1$. We’re interested in integrating a function $\tilde{\theta}$ over an interval $[-h, h]$ for various h . If $\tilde{\theta}$ is C^k and $\|\tilde{\theta}\|_T < V'$ on $[-h, h]$ for a real constant V' independent of h , then we can define $\theta(x) = \tilde{\theta}(hx)$ on $[-1, 1]$ and apply Eq. (A.13):

$$R_q[\tilde{\theta}] \leq \frac{32h^{k+1}V'}{15\pi k(2q+1-k)^k}. \quad (A.14)$$

This follows directly from the proof of [64, Theorem 4.2] applied to θ by replacing θ with $\tilde{\theta}(hx)$ and noting that $\theta^{(k)}(x) = h^k \tilde{\theta}^{(k)}(hx)$. The change of variables produces the first power of h , while each of the k integration by parts produces an additional power of h . In the context of hedgehog, the size of h is proportional to the edge length of the subdomain D_i outlined in Section 2.2.

Applying Eq. (A.14) to Eq. (A.12), and again letting $f(s, t) = \Theta(hs, ht)$, gives us

$$R_q^{(2)}[f] \lesssim \frac{32h^{k+1}}{15\pi k(2q+1-k)^k} [I_s[V'_t(s)] + I_t[V'_s(t)]] \quad (\text{A.15})$$

where $V'_t(s) = \max_t \|\Theta^{(k)}(hs, ht)\|_T$ and $V'_s(t) = \max_s \|\Theta^{(k)}(hs, ht)\|_T$ for fixed values of s, t . If we can choose a \tilde{V} that is strictly greater than $V'_s(t)$ and $V'_t(s)$ for any s, t in $\mathcal{I}^{(2)}$, we are left with

$$R_q^{(2)}[f] \lesssim \frac{128h^{k+1}\tilde{V}}{15\pi k(2q+1-k)^k}. \quad (\text{A.16})$$

Applying this to the integration of double layer potentials, we can simply let \tilde{V} be the largest variation of the k th partial derivatives of the integrand of any single patch in Eq. (7). In fact, we know that this value is achieved at the projection of \mathbf{x} on the patch P_i closest to \mathbf{x} , i.e., $(s^*, t^*) = \operatorname{argmin}_{\mathcal{I}^{(2)}} \|\mathbf{x} - P_i(s, t)\|_2$. We can also choose $h = \max_i h_i$ to observe standard high-order convergence as a function of patch domain size, which we summarize in the following theorem. The smoothness and bounded variation assumptions required to apply Eq. (A.13) to our layer potential follow directly from the smoothness of $u(\mathbf{x})$ in Ω . Our heuristic directly follows.

References

- [1] Mustafa Abduljabbar, Mohammed Al Farhan, Noha Al-Harthi, Rui Chen, Rio Yokota, Hakan Bagci, David Keyes, Extreme scale FMM-accelerated boundary integral equation solver for wave scattering, *SIAM J. Sci. Comput.* 41 (3) (2019) C245–C268.
- [2] Kendall Atkinson, Weimin Han, Numerical solution of Fredholm integral equations of the second kind, in: *Theoretical Numerical Analysis*, Springer, 2009, pp. 473–549.
- [3] Ludvig af Klinteberg, Anna-Karin Tornberg, Error estimation for Quadrature by Expansion in layer potential evaluation, *Adv. Comput. Math.* 43 (1) (2017) 195–234.
- [4] Ludvig af Klinteberg, Anna-Karin Tornberg, Adaptive Quadrature by Expansion for layer potential evaluation in two dimensions, *SIAM J. Sci. Comput.* 40 (3) (2018) A1225–A1249.
- [5] Bradley K. Alpert, Hybrid Gauss-trapezoidal quadrature rules, *SIAM J. Sci. Comput.* 20 (5) (1999) 1551–1584.
- [6] Alex H. Barnett, Evaluation of layer potentials close to the boundary for Laplace and Helmholtz problems on analytic planar domains, *SIAM J. Sci. Comput.* 36 (2) (2014) A427–A451.
- [7] Alex H. Barnett, Timo Betcke, Stability and convergence of the method of fundamental solutions for Helmholtz problems on analytic domains, *J. Comput. Phys.* 227 (14) (2008) 7003–7026.
- [8] Alex Bespalov, Timo Betcke, Alexander Haberl, Dirk Praetorius, Adaptive BEM with optimal convergence rates for the Helmholtz equation, *Comput. Methods Appl. Mech. Eng.* 346 (2019) 260–287.
- [9] J. Thomas Beale, A grid-based boundary integral method for elliptic problems in three dimensions, *SIAM J. Numer. Anal.* 42 (2) (2004) 599–620.
- [10] James Bremer, Zydrunas Gimbutas, A Nyström method for weakly singular integral operators on surfaces, *J. Comput. Phys.* 231 (14) (2012) 4885–4903.
- [11] James Bremer, Zydrunas Gimbutas, On the numerical evaluation of the singular integrals of scattering theory, *J. Comput. Phys.* 251 (2013) 327–343.
- [12] Timo Betcke, Alexander Haberl, Dirk Praetorius, Adaptive boundary element methods for the computation of the electrostatic capacity on complex polyhedra, preprint, arXiv:1901.08393, 2019.
- [13] Oscar P. Bruno, Leonid 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.
- [14] Oscar P. Bruno, Stéphane K. Lintner, A high-order integral solver for scalar problems of diffraction by screens and apertures in three-dimensional space, *J. Comput. Phys.* 252 (2013) 250–274.
- [15] J. Thomas Beale, Wenjun Ying, Jason R. Wilson, A simple method for computing singular or nearly singular integrals on closed surfaces, *Commun. Comput. Phys.* 20 (3) (2016) 733–753.
- [16] Stéphanie Chaillat, Luca Desiderio, Patrick Ciarlet, Theory and implementation of H-matrix based iterative and direct solvers for Helmholtz and elastodynamic oscillatory kernels, *J. Comput. Phys.* 351 (2017) 165–186.
- [17] Stéphanie Chaillat, Marion Darbas, Frédérique Le Louër, Fast iterative boundary element methods for high-frequency scattering problems in 3D elastodynamics, *J. Comput. Phys.* 341 (2017) 429–446.
- [18] Camille Carvalho, Shilpa Khatri, Arnold D. Kim, Asymptotic analysis for close evaluation of layer potentials, *J. Comput. Phys.* 355 (2018) 327–341.
- [19] Camille Carvalho, Shilpa Khatri, Arnold D. Kim, Asymptotic approximations for the close evaluation of double-layer potentials, preprint, arXiv:1810.02483, 2018.
- [20] Laurent Demanet, Alex Townsend, Stable extrapolation of analytic functions, preprint, arXiv:1605.09601, 2016.
- [21] Charles L. Epstein, Leslie Greengard, Andreas Klockner, On the convergence of local expansions of layer potentials, *SIAM J. Numer. Anal.* 51 (5) (2013) 2660–2679.
- [22] David Elliott, Barbara M. Johnston, Peter R. Johnston, Clenshaw–Curtis and Gauss–Legendre quadrature for certain boundary element integrals, *SIAM J. Sci. Comput.* 31 (1) (2008) 510–530.
- [23] David Elliott, Barbara M. Johnston, Peter R. Johnston, A complete error analysis for the evaluation of a two-dimensional nearly singular boundary element integral, *J. Comput. Appl. Math.* 279 (2015) 261–276.
- [24] Gerald Farin, *Curves and Surfaces for Computer Aided Geometric Design: A Practical Guide*, Academic Press Professional, Inc., San Diego, CA, USA, 1988.
- [25] Mahadevan Ganesh, Ivan G. Graham, A high-order algorithm for obstacle scattering in three dimensions, *J. Comput. Phys.* 198 (1) (2004) 211–242.
- [26] Leslie Greengard, Vladimir Rokhlin, A fast algorithm for particle simulations, *J. Comput. Phys.* 73 (2) (1987) 325–348.
- [27] Abinand Gopal, Lloyd N. Trefethen, Solving Laplace problems with corner singularities via rational functions, preprint, arXiv:1905.02960, 2019.
- [28] Thomas J.R. Hughes, John A. Cottrell, Yuri Bazilevs, Isogeometric analysis: CAD, finite elements, NURBS, exact geometry and mesh refinement, *Comput. Methods Appl. Mech. Eng.* 194 (39–41) (2005) 4135–4195.

[29] Johan Helsing, Rikard Ojala, On the evaluation of layer potentials close to their sources, *J. Comput. Phys.* 227 (5) (2008) 2899–2921.

[30] Jeremy G. Hoskins, Vladimir Rokhlin, Kirill Serkh, On the numerical solution of elliptic partial differential equations on polygonal domains, *SIAM J. Sci. Comput.* 41 (4) (2019) A2552–A2578.

[31] George C. Hsiao, Wolfgang L. Wendland, *Boundary Integral Equations*, Springer, 2008.

[32] Alec Jacobson, Ladislav Kavan, Olga Sorkine-Hornung, Robust inside-outside segmentation using generalized winding numbers, *ACM Trans. Graph.* 32 (4) (2013) 33.

[33] Seppo Järvenpää, Matti Taskinen, Pasi Ylä-Oijala, Singularity extraction technique for integral equation methods with higher order basis functions on plane triangles and tetrahedra, *Int. J. Numer. Methods Eng.* 58 (8) (2003) 1149–1165.

[34] Seppo Järvenpää, Matti Taskinen, Pasi Ylä-Oijala, Singularity subtraction technique for high-order polynomial vector basis functions on planar triangles, *IEEE Trans. Antennas Propag.* 54 (1) (2006) 42–49.

[35] Ludvig af Klinteberg, Alex H. Barnett, Accurate quadrature of nearly singular line integrals in two and three dimensions by singularity swapping, preprint, arXiv:1910.09899, 2019.

[36] Andreas Klöckner, Alexander Barnett, Leslie Greengard, Michael O’Neil, Quadrature by expansion: a new method for the evaluation of layer potentials, *J. Comput. Phys.* 252 (2013) 332–349.

[37] Shilpa Khatri, Arnold D. Kim, Ricardo Cortez, Camille Carvalho, Close evaluation of layer potentials in three dimensions, *J. Comput. Phys.* 423 (2020) 109798.

[38] Sharad Kapur, Vladimir Rokhlin, High-order corrected trapezoidal quadrature rules for singular functions, *SIAM J. Numer. Anal.* 34 (4) (1997) 1331–1356.

[39] Rainer Kress, *Linear Integral Equations*, Applied Mathematical Sciences, vol. 82, 1999.

[40] Yuxiang Liu, Alex H. Barnett, Efficient numerical solution of acoustic scattering from doubly-periodic arrays of axisymmetric objects, *J. Comput. Phys.* 324 (2016) 226–245.

[41] Bruno Lévy, *Geogram*, 2015.

[42] Libin Lu, Matthew J. Morse, Abtin Rahimian, Georg Stadler, Denis Zorin, Scalable simulation of realistic volume fraction red blood cell flows through vascular networks, preprint, arXiv:1909.11085, 2019.

[43] Dhairyra Malhotra, George Biros, PVFMM: a parallel kernel independent FMM for particle and volume potentials, *Commun. Comput. Phys.* 18 (3) (2015) 808–830.

[44] Dhairyra Malhotra, Antoine Cerfon, Lise-Marie Imbert-Gérard, Michael O’Neil, Taylor states in stellarators: a fast high-order boundary integral solver, preprint, arXiv:1902.01205, 2019.

[45] Solomon Grigorevich Mikhlin, *Integral Equations: And Their Applications to Certain Problems in Mechanics, Mathematical Physics and Technology*, vol. 4, Elsevier, 2014.

[46] Matthew J. Morse, Abtin Rahimian, Denis Zorin, Supplementary material for: a robust solver for elliptic PDEs in 3D complex geometries, <https://cims.nyu.edu/gcl/papers/2020-qbkix3d-supplementary.pdf>, 2020.

[47] Jean-Claude Nédélec, *Acoustic and Electromagnetic Equations*, Applied Mathematical Sciences, vol. 144, 2001.

[48] Constantine Pozrikidis, *Boundary Integral and Singularity Methods for Linearized Viscous Flow*, Cambridge University Press, 1992.

[49] Costas Pozrikidis, *Boundary Integral and Singularity Methods for Linearized Viscous Flow*, Cambridge University Press, 1992.

[50] Vladimir Zalmanovich Parton, Petr Il’ič Perlin, *Integral Equations in Elasticity*, Imported Pubn, 1982.

[51] Rodrigo B. Platte, Lloyd N. Trefethen, Arno B.J. Kuijlaars, Impossibility of fast stable approximation of analytic functions from equispaced samples, *SIAM Rev.* 53 (2) (2011) 308–318.

[52] Bryan Quaife, George Biros, High-volume fraction simulations of two-dimensional vesicle suspensions, *J. Comput. Phys.* 274 (2014) 245–267.

[53] Abtin Rahimian, Alex Barnett, Denis Zorin, Ubiquitous evaluation of layer potentials using Quadrature by Kernel-independent Expansion, *BIT Numer. Math.* 58 (2) (2018) 423–456.

[54] Manas Rachh, Andreas Klöckner, Michael O’Neil, Fast algorithms for Quadrature by Expansion I: globally valid expansions, *J. Comput. Phys.* 345 (2017) 706–731.

[55] Manas Rachh, Kirill Serkh, On the solution of Stokes equation on regions with corners, preprint, arXiv:1711.04072, 2017.

[56] Hanan Samet, *Foundations of Multidimensional and Metric Data Structures*, Morgan Kaufmann, 2006.

[57] Wojciech Śmigaj, Timo Betcke, Simon Arridge, Joel Phillips, Martin Schweiger, Solving boundary integral problems with BEM++, *ACM Trans. Math. Softw.* 41 (2) (2015) 6.

[58] Kirill Serkh, On the solution of elliptic partial differential equations on regions with corners II: detailed analysis, *Appl. Comput. Harmon. Anal.* (2017).

[59] Kirill Serkh, On the solution of elliptic partial differential equations on regions with corners III: curved boundaries, in preparation, 2018.

[60] Kirill Serkh, Vladimir Rokhlin, On the solution of elliptic partial differential equations on regions with corners, *J. Comput. Phys.* 305 (2016) 150–171.

[61] Kirill Serkh, Vladimir Rokhlin, On the solution of the Helmholtz equation on regions with corners, *Proc. Natl. Acad. Sci. USA* 113 (33) (2016) 9171–9176.

[62] Michael Siegel, Anna-Karin Tornberg, A local target specific quadrature by expansion method for evaluation of layer potentials in 3D, *J. Comput. Phys.* 364 (2018) 365–392.

[63] Svetlana Tlupova, J. Thomas Beale, Regularized single and double layer integrals in 3D Stokes flow, *J. Comput. Phys.* (2019).

[64] Lloyd N. Trefethen, Is Gauss quadrature better than Clenshaw-Curtis?, *SIAM Rev.* 50 (1) (2008) 67–87.

[65] Matthias Taus, Gregory J. Rodin, Thomas J.R. Hughes, Isogeometric analysis of boundary integral equations: high-order collocation methods for the singular and hyper-singular equations, *Math. Models Methods Appl. Sci.* 26 (08) (2016) 1447–1480.

[66] Lloyd N. Trefethen, J.A.C. Weideman, Two results on polynomial interpolation in equally spaced points, *J. Approx. Theory* 65 (3) (1991) 247–260.

[67] Matt Wala, Andreas Klöckner, A fast algorithm with error bounds for Quadrature by Expansion, *J. Comput. Phys.* 374 (2018) 135–162.

[68] Matt Wala, Andreas Klöckner, A fast algorithm for Quadrature by Expansion in three dimensions, *J. Comput. Phys.* 388 (2019) 655–689.

[69] Matt Wala, Andreas Klöckner, Optimization of fast algorithms for global Quadrature by Expansion using target-specific expansions, *J. Comput. Phys.* (2019) 108976.

[70] Matt Wala, Andreas Klöckner, On the approximation of local expansions of Laplace potentials by the fast multipole method, preprint, arXiv:2008.00653, 2020.

[71] Marcus Webb, Lloyd N. Trefethen, Pedro Gonnet, Stability of barycentric interpolation formulas for extrapolation, *SIAM J. Sci. Comput.* 34 (6) (2012) A3009–A3015.

[72] Bowei Wu, Hai Zhu, Alex Barnett, Shravan Veerapaneni, Solution of stokes flow in complex nonsmooth 2d geometries via a linear-scaling high-order adaptive integral equation scheme, *J. Comput. Phys.* (2020) 109361.

[73] Hong Xiao, Zydrunas Gimbutas, A numerical algorithm for the construction of efficient quadrature rules in two and higher dimensions, *Comput. Math. Appl.* 59 (2) (2010) 663–676.

- [74] Lexing Ying, George Biros, Denis Zorin, A kernel-independent adaptive fast multipole algorithm in two and three dimensions, *J. Comput. Phys.* 196 (2) (2004) 591–626.
- [75] Lexing Ying, George Biros, Denis Zorin, A high-order 3D boundary integral equation solver for elliptic PDEs in smooth domains, *J. Comput. Phys.* 219 (1) (2006) 247–275.
- [76] Lexing Ying, Denis Zorin, A simple manifold-based construction of surfaces of arbitrary smoothness, in: *ACM Transactions on Graphics (TOG)*, vol. 23, ACM, 2004, pp. 271–275.
- [77] Jürgen Zechner, Benjamin Marussig, Gernot Beer, Thomas-Peter Fries, The isogeometric Nyström method, *Comput. Methods Appl. Mech. Eng.* 308 (2016) 212–237.