



A fast, high-order scheme for evaluating volume potentials on complex 2D geometries via area-to-line integral conversion and domain mappings

Thomas G. Anderson*, Hai Zhu, Shravan Veerapaneni

Department of Mathematics, University of Michigan, United States of America

ARTICLE INFO

Article history:

Received 22 March 2022

Accepted 9 October 2022

Available online 14 October 2022

Keywords:

Newton potential

Inhomogeneous PDEs

Fast algorithms

Potential theory

Singular and near-singular
multi-dimensional quadrature

ABSTRACT

While potential theoretic techniques have received significant interest and found broad success in the solution of linear partial differential equations (PDEs) in mathematical physics, limited adoption is reported in the case of nonlinear and/or inhomogeneous problems (i.e. with distributed volumetric sources) owing to outstanding challenges in producing a particular solution on complex domains while simultaneously respecting the competing ideals of allowing complete geometric flexibility, enabling source adaptivity, and achieving optimal computational complexity. This article presents a new high-order accurate algorithm for finding a particular solution to the PDE by means of a convolution of the volumetric source function with the Green's function in complex geometries. Utilizing volumetric domain decomposition, the integral is computed over a union of regular boxes (lending the scheme compatibility with adaptive box codes) and triangular regions (which may be potentially curved near boundaries). Singular and near-singular quadrature is handled by converting integrals on volumetric regions to line integrals bounding a reference volume cell using cell mappings and elements of the Poincaré lemma, followed by leveraging existing one-dimensional near-singular and singular quadratures appropriate to the singular nature of the kernel. The scheme achieves compatibility with fast multipole methods (FMMs) and thereby optimal asymptotic complexity by coupling global rules for target-independent quadrature of smooth functions to local target-dependent singular quadrature corrections, and it relies on orthogonal polynomial systems on each cell for well-conditioned, high-order and efficient (with respect to number of required volume function evaluations) approximation of arbitrary volumetric sources. Our domain discretization scheme is naturally compatible with standard meshing software such as Gmsh, which are employed to discretize a narrow region surrounding the domain boundaries. We present 8th-order accurate results, demonstrate the success of the method with examples showing up to 12-digit accuracy on complex geometries, and, for static geometries, our numerical examples show well over 99% of evaluation time of the particular solution is spent in the FMM step.

© 2022 Elsevier Inc. All rights reserved.

* Corresponding author.

E-mail address: tganders@umich.edu (T.G. Anderson).

1. Introduction

This article describes a fast, high-order accurate numerical scheme for evaluating the *volume potential*, also known as the *Newton potential*, given by

$$\mathcal{V}[f](\mathbf{r}_0) = \int_{\Omega} G(\mathbf{r}, \mathbf{r}_0) f(\mathbf{r}) dA(\mathbf{r}), \quad \mathbf{r}_0 \in \Omega, \quad (1.1)$$

where Ω is an irregular two-dimensional domain, f is a given source density function and G is typically the free-space Green's function for an underlying linear, constant-coefficient elliptic partial differential equation (PDE) operator, such as for the Laplace, Stokes or (modified) Helmholtz equations. Domain convolutions of this form are a key ingredient when solving inhomogeneous PDEs *via* potential theory [1]. Historically, the community has tackled a variety of fundamental challenges in the use of potential theoretic methods for addressing such problems.

While integral equation methods offer a reduction in dimensionality in the associated homogeneous problem, they lead to dense operators upon discretization, whose efficient application has formed a significant body of work over the recent decades (such as the FMM [2]). With regards to discretization, as the kernels which arise in boundary integral formulations for homogeneous PDE boundary value problems are singular, the integral equations that result involve singular and nearly-singular integrals. Singular quadrature has formed the basis of significant inquiry with many successful schemes proposed; see [3] for a recent review of the situation in two dimensions (note that in three dimensions the landscape is more challenging). Furthermore, *nearly singular* integrals arise from the nature of these kernels in the presence of arbitrary target points close to $\partial\Omega$, as generally occur in complex geometries e.g. when boundary components are close to touching. These challenges each have a counterpart in the volume potential problem. Recently, a concerted effort in the field produced several robust strategies for high-accuracy evaluation of nearly singular integrals in two dimensions [4–11]. For example, the method of [6], which we employ for the homogeneous solver component of this work, achieves spectral accuracy in evaluating layer potentials close to smooth boundaries, requiring only a modest number of discretization nodes per domain inclusion for full accuracy in double precision. These advances motivate us to revisit the volume potential evaluation problem as a prominent remaining task.

Most fundamentally for inhomogeneous problems, the volume discretization needs to be adapted to the irregular domain. We first mention that for the case of Ω a unit box or a hierarchical finite union of scaled and translated boxes, there exist fast, adaptive, highly-accurate and computationally scalable algorithms [12–15]. In a similar vein, reference [16] solves local elliptic problems on each box using the action of the operator on orthogonal polynomials and obtains high-order accuracy with very limited computational expense per degree of freedom. Taken together, these “box code” methods have proven highly effective because they inherently exploit the translational invariance of the Green function of the PDE, allowing repeated use of what are essentially lookup tables for singular and near-singular evaluation points within a vicinity of a given box. However, achieving the same level of accuracy and efficiency when Ω is an arbitrary complex geometry has been a sustained challenge. Naïvely, for a box code in the presence of complex geometry some boxes inevitably are ‘cut’ by the boundary and evaluating the contribution to the volume potential (1.1) from such irregular ‘cut cells’ is the primary challenge to be met (such issues arise in other contexts utilizing regular grids in the presence of embedded boundaries see e.g. [17]). Attempts have been made to address irregular geometries by using (generally low-order accurate) extrapolation or local extension for the function f over the resulting cut cells in combination with extensive adaptivity near to domain boundaries to achieve desired tolerances [18]. As a result, while the resulting methods effectively make use of fast algorithms to reduce the computational burden, the required number of degrees of freedom appears to be significant.

In seeking to overcome this issue, most existing works avoid direct evaluation of (1.1) in an irregular domain; instead, one class of algorithms employs (local) volumetric PDE solvers—finite difference [19–22] and finite element [23] methods—in an embedded regular domain and obtains corrected stencils near to the boundary using Taylor expansions and local geometric information as well as, in some cases, jump relations of layer potentials. In particular, the embedded boundary integral approach of [23] solves the inhomogeneous PDE problem on a rectangular domain that embeds Ω ; local low-order extrapolation for bulk forces and use of jump relations result in a second-order accurate method, which appears quite efficient. The method of reference [22], meanwhile, couples to geometric multigrid solvers for the bulk and is also restricted to the Poisson equation; this work in fact includes a timing comparison to the box code method of reference [24] (discussed in the next paragraph), which shows superior speed per degree of freedom. We note, however, that only first-order convergence is demonstrated in [22], with no claim made that the accuracy in the comparison example matches that of the box code with continuous extension, and we further observe that while the method implicitly requires extension of the source density outside the domain, this point is unaddressed (i.e. known source functions are chosen that are continuous across the boundary and their values outside used).

Smooth extension or continuation methods for the volumetric source form another broad class of methods, in which the irregular domain Ω , again, is embedded in a box B , the function f is extended onto $B \setminus \Omega$, and standard box codes are applied on the whole of B (inevitably incurring some performance penalty due to the increased degrees of freedom over the enlarged domain). It is essential in these methods that the extended function be smooth for high-order accuracy. In [24], a harmonic extension of f is performed using a boundary integral approach and is coupled to an adaptive box code. While the approach generalizes beyond the Poisson equation solver described in this work, its use for the construction of

higher-order smooth approximations does depend upon higher-order derivatives of the source function f along the domain boundary. More significantly, it requires solution of a separate harmonic equation in order to generate the continuously extended source function, with higher-order continuity for higher-order convergence requiring solution of progressively more complex and costly high-order harmonic equations. Indeed the use of adaptivity in [24] functions at least in part as a means to achieve high accuracy in the presence of limited smoothness in the extension, as experienced previously in [18].

Other approaches to extension have traded off adaptivity for high-order accuracy and compatibility with FFTs on uniform grids; one recently introduced technique is the two-dimensional Fourier continuation method [25] which as well as being a general-purpose numerical tool has recently been demonstrated to give accurate smooth periodic extensions suitable for use in elliptic solvers, at least in some simple geometrical contexts. In [26], a function extension scheme based on the use of radial basis functions (RBFs) is proposed, termed as the *partition of unity extension* (PUX) method. PUX lays down a set of disks along the domain boundary, each covering uniform points both inside and outside the domain, and solves overdetermined linear systems on each disk to generate smooth extensions to Cartesian grid points within these disks laying outside the domain, which are then smoothly mollified to zero away from the boundary. Although PUX has shown success on a variety of examples [27,28], several hurdles remain for a scalable implementation including the number of parameters and the heuristic nature by which they are determined (e.g., the partition radius and shape parameters of the RBFs), poor conditioning of RBF methods and a reliance on uniform grids.

In this article, we take a more direct approach to solving inhomogeneous linear PDEs, by computing explicitly the volume integral (1.1) by means of numerical quadrature over complex geometry (an example is shown in Fig. 1). This mathematically obvious avenue has in the past been considered computationally impractical, with recent contributions [24,27,28] noting that such an approach is to be avoided because quadratures for irregular cells are seemingly challenging and less amenable to fast algorithms (reference [29] does evaluate (1.1), in a method coupled to the FMM). Indeed, multidimensional quadratures over general regions are much less developed (though see [30–32]) though there exist ad-hoc quadratures over a variety of specific regions [33–37]; see [38] for an extensive review. This relative lack of development holds especially so for singular integrands; in this more specialized context, in addition to the older work [36] on triangles we find the work [39,40] for locally-corrected quadratures on rectangles and work [41,42] for discretizations of boundary integral operators on mapped triangles that form surfaces in three dimensions.

To address the challenges of multi-dimensional quadrature over potentially arbitrary regions, a variety of works have converted a volumetric integral of interest *with a smooth integrand* to integration along domain boundaries (potentially after some degree of domain-decomposition), after which numerical quadrature rules are used for various resulting one-dimensional integrals. Much of this work is focused on cubature for polyhedra, where we note contributions for this purpose [43,44] based on theorems of vector calculus. One approach [45] treats volumes bounded by rational parametric curves using Green's theorem, while others [46,47] produce volume quadratures over implicitly-defined surfaces and volumes by a recursive dimensional-reduction algorithm and ultimately result in one-dimensional integrals that can be treated with Gaussian quadrature.

We also note a variety of volumetric-via-surface-integral quadratures that, like ours, use ideas from classical proofs of the Poincaré lemma (which, in brief, is applicable to domains which are star-shaped with respect to *some* interior point—but not necessarily to all points—and leads to representations of volume integrals as surface integrals with an iterated integral whose associated physical integration points lay on rays from the boundary to the star-point). Recent work of [48] results in integral representations similar in some respects to our own but only smooth integrands are considered there. Other works such as [49] have focused on addressing the Newton potential problem specifically, but a target-centered coordinate system with the star-point as the origin results in quadrature points that lay outside the domain for regions that are non-convex (i.e. containing regions, and therefore potential evaluation points, with respect to which the domain is not star-shaped). This fact implicitly imposes upon such methods a reliance on function extension in order to obtain values of the source density f at non-physical quadrature nodes (or, alternatively, moment matching)—such tasks (which, in view of previous discussion

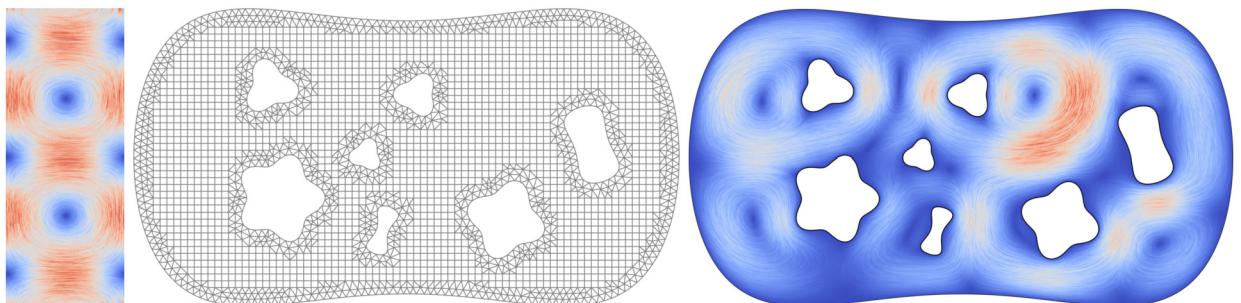


Fig. 1. Solution of the inhomogeneous Stokes equations [1] on an irregular domain with no-slip boundaries. Left: contour plot of the imposed forcing function, $f(x, y) = (-\sin x \cos y, \cos x \sin y)^T$. Center: volume mesh employed in our Newton potential (1.1) evaluation scheme comprised of a regular grid in the bulk and a narrow triangulated region around the domain boundaries. Right: streamlines of the solution with color indicating the magnitude of the velocity field.

on function extension can be seen to be challenging) were not addressed in those works, instead relying on assumed-known analytic expressions outside of the domain. Related efforts to avoid limitations of this kind with schemes designed to rely on points only inside the computational domain are described in [50,51] (and include acceleration via FMM), but numerical results demonstrate first-order convergence (in the simplest cases) or no convergence at all. The dual reciprocity method (see [52] for references to this body of literature and see [53,54] for some related works) is another method popular in the engineering literature that employs surface integrals and delivers low to modest accuracy, with similar challenges as mentioned above. Our proposed scheme, while still based in spirit on the Poincaré lemma and while it also utilizes surface integrals, does not suffer from any of these difficulties (cf. Remark 3): all quadrature nodes lay inside the domain and the scheme leads to high-order accuracy.

Our work generalizes the volume potential scheme developed in the second author's thesis [55] (also see [56]), wherein, Poincaré's lemma and recursive product integration rules were employed for a few pairs of PDE kernels and approximating bases. In a similar vein, prior work in [57] utilizes congruent integral representations that are limited to the Laplace kernel and to convex spatial domains—work which, additionally, lists several other interesting methods for the production of particular solutions. On the other hand, the methods proposed here generate numerical quadratures for regions arising from domain decomposition of the integral (1.1), which are tailored to kernel functions with a variety of singular behaviors (i.e. not restricted to PDE Green function kernels) and lead to efficient algorithms that are compatible with fast algorithms (such as FMMs). An important goal is to retain compatibility with adaptive box codes [12,15] due to their unparalleled speed and maturity—which we achieve by triangulating only a small boundary-fitted region—but we also seek high-order convergence and speed for repeated application of the volume potential with different source densities f . A secondary objective is the use of components (e.g. meshing, choice of basis, quadrature generation schemes) that will generalize naturally to three dimensional volume quadratures.

Synopsis. The proposed method proceeds by meshing a thin region near to the boundary using a combination of curvilinear (with one side conforming to the boundary represented as parametric curves) and straight triangles—the number of triangular regions growing *linearly* as the mesh is refined with the number of regular boxes growing quadratically (for the uniform gridding of the bulk that we consider here). Each cell (or element) is represented as a map from a reference element (either of a simplex or a box, with straight triangles corresponding to a simple affine map) and all physical target points are characterized by their position in reference space relative to cells for the purposes of singular and near-singular quadrature. Singular and near-singular quadrature on these irregular domains (and indeed even over boxes), in turn, is performed with high-order accuracy by use of a Poincaré lemma-type idea that re-writes the integrals (in reference space) on standard triangles (or boxes) over the boundary of that domain, followed by the use of one-dimensional quadratures that are adapted to the singular behavior of the kernel function (typically, the PDE Green function). The method possesses optimal asymptotic complexity (costing $\mathcal{O}(N)$ to obtain the Newton potential at, say, all N quadrature points in a domain) as a result of use of point-FMMs for summation of target-independent quadratures for smooth integrands and via coupling to inexpensive local singular corrections that can be represented as a linear map of small size.

Advantages. Our proposed methodology builds naturally on existing and growing bodies of work concerning each of *i*) one-dimensional singular and near-singular quadrature as well as *ii*) interpolation and quadrature of smooth functions over convex polyhedra by means of orthogonal polynomials. It is fully compatible with adaptive schemes for spatially-concentrated sources. Problem geometry is exactly captured with our approach which incorporates the local boundary information and integrates it with existing software for high-quality mesh-generation. High-order accuracy is easy to achieve and requires only a one-dimensional quadrature rule adequate for the singularity of the given kernel function and a known orthogonal polynomial system with associated smooth quadrature rule on the reference cell. The previous two points (meshing and orthogonal polynomial systems) provide direct means to generalize this work to three-dimensions. Finally, the methods are straightforward to integrate with existing acceleration techniques.

Limitations. At the present moment our solver discretizes regular portions of the domain with uniformly sized boxes, which can be inefficient for functions with significant variation. This issue is easily addressed since our scheme is designed to be drop-in compatible with adaptive box codes for the bulk region that are highly effective in this context. A more serious limitation is the cost of generating local singular corrections; while the regular bulk region does not contribute to these costs (so that asymptotically the set-up costs pertain only to thin boundary regions), and we present techniques in Section 4 to manage the remaining computational burden, computation of these quantities still forms the majority of the up-front cost of the method. Lastly, we restrict our attention in this work to static geometries only.

Outline. This article proceeds in Section 2 with necessary preliminaries that present boundary integral equation formulations for some common linear PDEs, demonstrates the role that volume potentials play in their solution, and then presents in Sections 3 and Section 4 the proposed methodology: Section 3 covers volumetric meshing, then the use of smooth quadratures and interpolation on the coordinate-mapped triangles that naturally arise, followed by a description of novel singular and near-singular quadrature techniques, while Section 4, in turn, focuses on implementation and efficiency considerations of the method, both for repeated evaluation of (1.1) and for efficient up-front generation of necessary corrections to the smooth quadratures (as well as some limited discussion of our approach to the routine task of quadratures on the bulk region). Various numerical examples are presented in Section 5 demonstrating the properties of the method and a brief summary with concluding remarks is given in Section 6.

2. Potential theory for inhomogeneous PDEs

This article is primarily concerned with methods to obtain solutions $u : \Omega \rightarrow \mathbb{R}$ to the elliptic PDE boundary value problem

$$Lu(\mathbf{r}) = f(\mathbf{r}), \quad \mathbf{r} \in \Omega, \quad (2.1a)$$

$$u(\mathbf{r}) = g(\mathbf{r}), \quad \mathbf{r} \in \Gamma, \quad (2.1b)$$

where $\Gamma = \partial\Omega$ denotes the (assumed piecewise-smooth) boundary of a complex geometry $\Omega \subset \mathbb{R}^2$. Here, it is assumed that L is a linear operator with a known translation-invariant Green function $G(\mathbf{r}, \mathbf{r}_0) = G(\mathbf{r} - \mathbf{r}_0)$, as occurs e.g. for the Poisson, Helmholtz, modified Helmholtz, and Stokes equations. A standard solution technique for the boundary value problem (2.1) is to exploit linearity and seek a particular solution u_P to the problem

$$Lu_P(\mathbf{r}) = f(\mathbf{r}), \quad \mathbf{r} \in \Omega, \quad (2.2)$$

with no care given to boundary conditions for u_P on Γ , and then subsequently solve the augmented boundary value problem

$$Lu_H(\mathbf{r}) = 0, \quad \mathbf{r} \in \Omega, \quad (2.3a)$$

$$u_H(\mathbf{r}) = g(\mathbf{r}) - u_P(\mathbf{r}), \quad \mathbf{r} \in \Gamma, \quad (2.3b)$$

whereby the solution to (2.1) is given by

$$u(\mathbf{r}) = u_H(\mathbf{r}) + u_P(\mathbf{r}), \quad \mathbf{r} \in \Omega. \quad (2.4)$$

Depending on the operator L and the domain Ω a variety of boundary integral equation formulations may be appropriate, each requiring their own numerical analysis—such questions are not of concern here. It suffices to note that all require boundary values of a particular solution u_P which can be obtained by evaluation of the volume potential (1.1). This article is thus concerned with the efficient and accurate evaluation of this integral for targets $\mathbf{r}_0 \in \Gamma$ (to provide boundary values of u_P as data for Equation (2.3b)) and for $\mathbf{r}_0 \in \Omega$ (for evaluation of the full solution u at desired evaluation points $\mathbf{r} \in \Omega$). Nevertheless, we briefly outline for definiteness the integral equation formulations used in this article.

2.1. Boundary integral formulation

Consider the Poisson equation, for which $L = -\Delta$. We utilize standard representation formulas for the solution u_H at a point \mathbf{r}_0 in the domain Ω , expressed in terms of the double-layer potential

$$u_H(\mathbf{r}_0) = \mathcal{D}[\varphi](\mathbf{r}_0) := \int_{\Gamma} \frac{\partial G(\mathbf{r}, \mathbf{r}_0)}{\partial \mathbf{n}(\mathbf{r})} \varphi(\mathbf{r}) d\sigma(\mathbf{r}) \quad (2.5)$$

induced by the boundary integral density φ , where $G(\mathbf{r}, \mathbf{r}_0)$ denotes the Green function of Equation (2.1), $G(\mathbf{r}, \mathbf{r}_0) = -\frac{1}{2\pi} \log(|\mathbf{r} - \mathbf{r}_0|)$. As is well-known, using the representation formula (2.5) and enforcing the boundary conditions (2.3b) leads via the jump relations of the double layer potential [58] to the integral equation

$$\left(\pm \frac{1}{2} I + D \right) [\varphi](\mathbf{r}) = g - u_P, \quad \mathbf{r} \in \Gamma_{\pm}, \quad (2.6)$$

for a function φ which must be satisfied in order for Equation (2.5) to yield a solution to the boundary value problem (2.3) for $L = -\Delta$. Here Γ_+ (resp. Γ_-) denotes that section of the boundary $\partial\Omega$ with respect to which the domain lays exterior (interior), and we denote by D the double-layer boundary integral operator

$$D[\psi](\mathbf{r}_0) := \int_{\Gamma} \frac{\partial G(\mathbf{r}, \mathbf{r}_0)}{\partial \mathbf{n}(\mathbf{r})} \psi(\mathbf{r}) d\sigma(\mathbf{r}), \quad \mathbf{r}_0 \in \partial\Omega. \quad (2.7)$$

The above formulation works as written for the modified Helmholtz equation ($L = -\Delta + \lambda^2$) as well by replacing the Green's function with $G(\mathbf{r}, \mathbf{r}_0) := \frac{\lambda^2}{2\pi} K_0(\lambda|\mathbf{r} - \mathbf{r}_0|)$. We do not discuss here subtler points regarding treatment of nullspaces for these operators, see [59] for a treatise on this topic.

3. A volume potential scheme for thin boundary-fitted regions

This section describes the basic elements of the proposed methodology for volume potential evaluation. It describes first an automatic method for the construction of a boundary-fitted region consisting of both curvilinear triangles (that are fitted to the boundary of inclusions defined by parametric curves) and straight triangles. Then, domain mappings are presented that map a reference cell to a triangular region and describe smooth quadrature and interpolation schemes for these regions (our approach to the routine task [12] of quadrature for regular source boxes is deferred and only briefly mentioned in Section 4.2). Finally, in Section 3.3 we describe our approach to singular and near-singular quadrature based on expressing domain integrals over a reference cell in terms of line integrals on its boundary.

3.1. Meshing

This section describes a technique to generate a boundary-fitted mesh for a (possibly multiply-) connected volumetric region Ω with boundary formed by a collection of curves $\Gamma = \bigcup_{i=0}^{N_\Gamma} \Gamma_i$, with the method proceeding in an identical manner whether the domain is an unbounded (exterior) or bounded (interior) one. The meshing algorithm we describe depends on parametrizations of the boundary curves, and to this end we introduce some useful notation. We assume for each curve Γ_i ($i = 1, \dots, N_\Gamma$) that we have access to a global parametrization $\gamma_i : [0, 2\pi] \rightarrow \Gamma_i$, $\gamma_i = \gamma_i(t)$, and denote by \mathbf{n}_i the associated normal vector $\mathbf{n}_i = \mathbf{n}_i(t)$ directed into the domain Ω . We also denote by $\ell_i(t)$ the arclength of the portion of the curve Γ_i traced out by $\gamma_i(\tau)$ for $0 \leq \tau \leq t$, and, abusing notation slightly, we call $\ell_i = \ell_i(2\pi)$ the total arclength of curve Γ_i . The algorithm fills the bulk of the domain away from the boundary with boxes in a uniform background mesh, assumed to be of size h .

The boundary-fitted mesh will be in the form of a tessellation involving N_t triangular regions \mathcal{T}_k and N_b regular boxes \mathcal{B}_k ,

$$\Omega = \bigcup_{k=1}^K \mathcal{C}_k = \left(\bigcup_{k=1}^{N_t} \mathcal{T}_k \right) \bigcup \left(\bigcup_{k=1}^{N_b} \mathcal{B}_k \right), \quad (3.1)$$

which together comprise $K = N_t + N_b$ cells, ordered first by triangles so that

$$\mathcal{C}_k = \begin{cases} \mathcal{T}_k, & 1 \leq k \leq N_t, \\ \mathcal{B}_{k-N_t}, & N_t + 1 \leq k \leq K, \end{cases} \quad (3.2)$$

where the box \mathcal{B}_k has center \mathbf{o}_k .

Corresponding to the curve Γ_i with arclength ℓ_i there will be a contribution of approximately ℓ_i/h boundary-fitted (curved) triangles, reflecting general uniformity in the size of mesh cells that abut the boundary; other strategies incorporating considerations of local curvature or source adaptivity are also possible but beyond the scope of the present discussion (see also Remark 2). For each of the N_Γ curves, the algorithm proceeds in three steps to generate a boundary-fitted mesh: (1) Firstly, by identification of boundary ‘knot’ points which segment each curve parametrization into N_{γ_i} sections of approximately equal-arclength ($\Delta\ell \approx h$), then (2) Secondly, by generation of boundary-fitted mesh cells conforming to the knot points, and (3) Finally, by generation of “buffer” zones that connect the boundary-fitted mesh to the background “bulk” mesh.

First, we describe a method for generating a sequence $\{\kappa_j\}_{j=1}^{N_{\gamma_i}}$ of parametric knot points $\kappa_j \in [0, 2\pi]$. Starting from $t = 0$, knot points are progressively laid down so that each segment of the curve formed by consecutive knot points contains an arclength of approximately $\Delta\ell \approx h$; care is taken that all knot points are reasonably spaced (in particular, that knots κ_j do not lay close to $t = 2\pi$ —which may lead to close knot points in view of the periodicity of $\gamma_i(t)$). Then, from each pair of knot points κ_j and κ_{j+1} , $j = 1, \dots, N_{\gamma_i} - 1$, a curvilinear triangle is generated by connecting the two knot points by the boundary curve, and then further connecting to a third point in the volume generated by projecting a point a distance h from the boundary in the normal direction $\mathbf{n}_i((\kappa_j + \kappa_{j+1})/2)$. The three points forming the new triangle thus are $\gamma_i(\kappa_j)$, $\gamma_i(\kappa_{j+1})$, and $\gamma_i(m_j) + h\mathbf{n}_i(m_j)$, where $m_j = (\kappa_j + \kappa_{j+1})/2$ is the curve-parametric midpoint of the new triangle being formed. The curved cells that result from this procedure are shown in Fig. 2, shaded grey.

Since the mesh cells formed in such a manner do not conform in any way to any uniform (or quadtree) background mesh, a watertight mesh is generated by means of a global ‘buffer’ zone \mathcal{F} that lays flush with these cells and serves to separate them from the background mesh. This region is defined by excluding any box from the background mesh laying within distance $h\delta$ of any boundary-fitting triangles (the selection $\delta = 0.8$ was made for the experiments in this article), the process repeating for each of the curves Γ_i , and yielding N_Γ local buffer zones \mathcal{B}_i ($i = 1, \dots, N_\Gamma$) which satisfy $\bigcup_i \mathcal{B}_i = \mathcal{F}$. In the simple case that all such regions \mathcal{B}_i are pairwise disjoint, it is a routine meshing task to generate N_Γ high-quality triangulations for each as each is bounded by two polygons; we use the Gmsh [60] software suite for this task which results in a total of N_t (curved and straight) triangular mesh cells.

One possibility that must be addressed, then, is the situation that arises when ‘buffer’ zones from two or more curves Γ_i and Γ_j , $i \neq j$, are overlapping ($\mathcal{B}_i \cap \mathcal{B}_j \neq \emptyset$) as occurs naturally when Γ_i and Γ_j are not sufficiently well-separated. The algorithm handles this situation by utilizing a queue-like system that progresses iteratively through all remaining unmeshed curves to identify overlapping \mathcal{B}_i , merging these, and triangulating the result. Selecting the first remaining curve Γ_i and

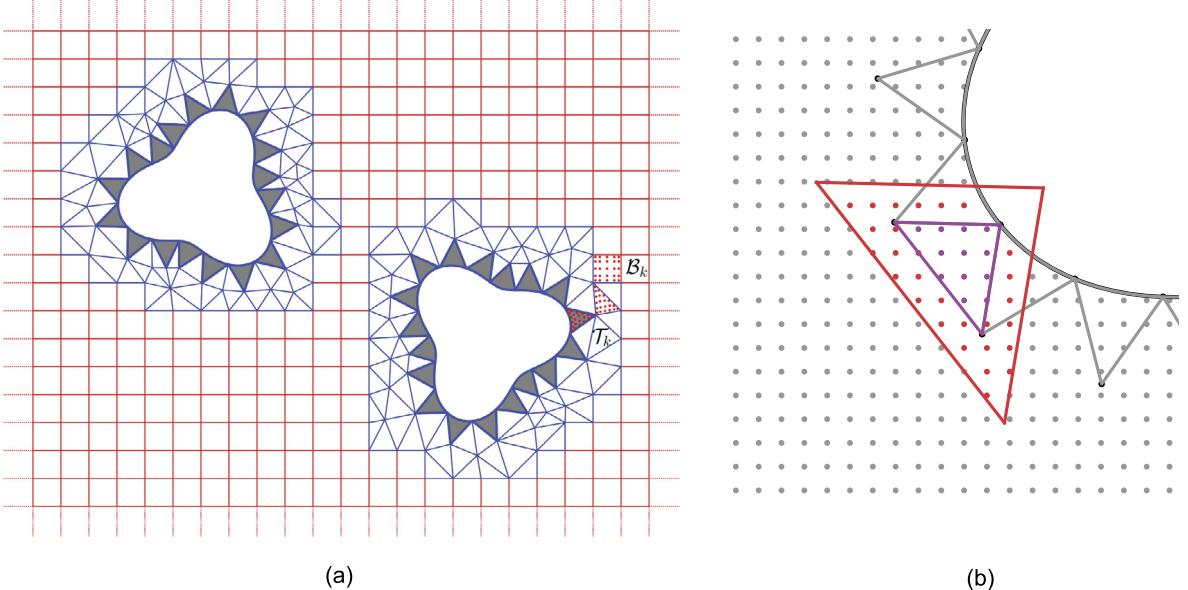


Fig. 2. (a) The volumetric discretization that arises from the present methodology applied to the exterior of two smooth bounded curves. Curved triangles that abut the curve are shaded while the buffer zone triangle cells are identified by blue edges. In three of the cells the interpolation nodes are denoted with red points. (b) The self-near-far separation rule for a straight triangular cell (with cell boundary plotted in purple and vertices in black) applied to uniformly-distributed volumetric targets. The region bounded by the red triangle contains all near-singular targets (marked red), the triangle itself bounds all singular targets (marked purple), while the remainder of targets (marked gray) are considered to be far quadrature to be treated with a smooth quadrature rule. (For interpretation of the colors in the figure(s), the reader is referred to the web version of this article.)

letting $\mathcal{F}_i = B_i$, for any curve Γ_j ($j > i$) whose buffer zone B_j satisfies $\mathcal{F}_i \cap B_j \neq \emptyset$ we let $\mathcal{F}_i = \mathcal{F}_i \cup B_j$ and remove Γ_j from the queue of remaining curves. This process terminates when $j = N_\Gamma$ after which the index i is increased to the next remaining curve and the process continues until no curves are remaining ($i = j = N_\Gamma$). The resulting sets \mathcal{F}_i are pairwise disjoint and satisfy $\cup_i \mathcal{F}_i = \mathcal{F}$ (the union taken over all i for which \mathcal{F}_i is defined); each are triangulated individually, again with Gmsh. This procedure thus robustly handles the possibility of arbitrarily-many curves that lay close to each other.

A final depiction of the region surrounding one or more curves Γ_i is given in Fig. 2 for well-separated curves (the buffer zone cells \mathcal{F} are shaded blue), while examples of the buffer-zone merging process described in the previous paragraph can be seen in the mesh in Fig. 8. The set of regular cells is defined, finally, by the remaining part of the volume, per Equation (3.1); in Fig. 2 these regular cells are colored red (note that regular cells may be enclosed by any given union-of-buffer-zones region \mathcal{F}_i and are still handled by the box code).

A final element of the mesh construction process is a self-near-far separation rule, which classifies a given target \mathbf{r}_0 as a singular, near-singular, or smooth quadrature point with respect to each mesh cell. A target $\mathbf{r}_0 \notin \mathcal{C}_k$ is considered a near-singular target of \mathcal{C}_k if it lays within a polygon, with sides parallel to the physical cell (of course, this prescription is extended in an obvious manner to treat curvilinear triangles) and with sides laying a distance Dh (we use $D = 4/10$ for experiments in this article) from the boundary of \mathcal{C}_k . We denote by $\mathcal{C}_k^{\text{near}}(\mathbf{r}_0)$ the index set of cells whose associated separation rules classify \mathbf{r}_0 as a near-singular target, with $\mathcal{C}^{\text{near}}(\mathbf{r}_0)$ the union of all cells \mathcal{C}_k with index in $\mathcal{C}_k^{\text{near}}(\mathbf{r}_0)$. We denote by $\mathcal{C}^{\text{self}}(\mathbf{r}_0)$ the cell to which \mathbf{r}_0 lays in, and define further $\mathcal{C}^{\text{far}}(\mathbf{r}_0) = \Omega \setminus (\mathcal{C}^{\text{near}}(\mathbf{r}_0) \cup \mathcal{C}^{\text{self}}(\mathbf{r}_0))$ as the remainder of Ω which is well-separated from \mathbf{r}_0 (see the depiction in Fig. 2). Therefore, the volume potential can be decomposed as

$$\mathcal{V}[f](\mathbf{r}_0) = \left(\int_{\mathcal{C}^{\text{near}}(\mathbf{r}_0)} + \int_{\mathcal{C}^{\text{self}}(\mathbf{r}_0)} + \int_{\mathcal{C}^{\text{far}}(\mathbf{r}_0)} \right) G(\mathbf{r}, \mathbf{r}_0) f(\mathbf{r}) dA(\mathbf{r})$$

over near-singular, singular, and smooth regions of Ω , respectively, for each target \mathbf{r}_0 .

Remark 1. It is trivial to define for each triangular cell a boundary parametrization $Z_k : [0, 2\pi] \rightarrow \partial \mathcal{T}_k$, which could in principle be utilized by the same methodology outlined in Section 3.3 to generate physical-space volumetric quadratures for targets located in or near $\partial \mathcal{T}_k$. In practice, however, all quadrature will occur on the standard simplex $\widehat{\mathcal{T}}_0$ or the unit box $\widehat{\mathcal{B}}_0$ (see also Remark 3). It should be noted as well that the boundary of every box \mathcal{B}_k in fact has the same parametrization, up to a translation, with similar implications for generating physical-space quadratures.

Remark 2. The prescription we give here for generating a high-quality volumetric mesh is certainly not the only method to obtain a compatible mesh for volume potential evaluation, but was effective for forming meshes for the examples in this article. Local curvature could be usefully incorporated into the algorithm described above in a blend of strategies based on arclength and changes in curvature, to achieve higher quality meshes. We mention as well the “TriWild” method of reference [61] for producing a valid volumetric mesh consisting of straight and curved triangles from a collection of curves whose union forms Γ , which has demonstrated success on a variety of real-world benchmark problems and is fully compatible with our curved triangle approach to the boundary-fitted region. (At the present moment implementations of the TriWild technique will result in only third-order approximations to the true boundary curve, motivating our exact approach to geometry representation, but it appears that this is not an essential limitation of the general TriWild technique; moreover, it appears straightforward to utilize slight perturbations of TriWild meshes and retain full geometric accuracy.)

3.2. Quadratures and interpolation for smooth integrands

In this section we describe first the overall approach to computing the volume potential $\mathcal{V}[f](\mathbf{r}_0)$ at arbitrary target points \mathbf{r}_0 via domain decomposition and then discuss quadrature of a potential that contains *smooth* integrands (as well as the related source density interpolation problem) over a domain-decomposed cell \mathcal{C}_k . Efficient and accurate evaluation follows from a combination of (i) Standard fast summation technologies which, on the one hand allow for efficient computation of sums arising from target-independent quadratures and which are, on the other hand, necessarily inaccurate in the vicinity of a given target point \mathbf{r}_0 for a singular kernel, as well as (ii) Accurate singular and near-singular *target-dependent* quadrature corrections applied locally in the vicinity of the given target point \mathbf{r}_0 . While the focus in this section is on point (i), the methodology does rely on a fixed orthogonal basis used to approximate the source function f to high-order accuracy and for this reason we discuss also interpolation of smooth functions. Our approach is similar in spirit to that of reference [62].

We first sketch the unifying singular correction and functional approximation strategy. The Newton potential of a function ρ at a given target point $\mathbf{r}_0 \in \Omega$ can be written as

$$\mathcal{V}[\rho](\mathbf{r}_0) = \sum_{k=1}^K \mathcal{V}_k[\rho](\mathbf{r}_0), \quad \text{where} \quad \mathcal{V}_k[\rho](\mathbf{r}_0) = \int_{\mathcal{C}_k} G(\mathbf{r}, \mathbf{r}_0) \rho(\mathbf{r}) dA(\mathbf{r}).$$

Quadrature will be performed using coordinate mappings \mathbf{R}^k from a reference cell $\widehat{\mathcal{C}}_k$ —either a unit square $\widehat{\mathcal{B}}_0$ or the unit simplex $\widehat{\mathcal{T}}_0$ (box regions are identical by translational invariance, so the use of mappings in this case is done merely for notational consistency). We will denote by

$$\widehat{\mathcal{V}}_k[\rho](\xi_0) = \int_{\widehat{\mathcal{C}}_k} G\left(\mathbf{R}^k(\xi), \mathbf{R}^k(\xi_0)\right) J^k(\xi) \rho(\xi) dA(\xi), \quad \text{where} \quad J^k(\xi) = \det \frac{\partial \mathbf{R}^k}{\partial \xi}, \quad (3.3)$$

a volume potential in ξ -reference space, $\xi = (\xi, \eta)^T$. In a slight abuse of notation we will when convenient use the vector notation $\mathbf{R}^k(\xi)$ instead of the notation $\mathbf{R}^k(\xi, \eta)$ for the same function \mathbf{R}^k (and similarly for other functions). It is known that if \mathbf{R}^k is a C^1 -invertible mapping onto \mathcal{C}_k , then letting $\xi_0^k = (\mathbf{R}^k)^{-1}(\mathbf{r}_0)$ denote the location of a physical target $\mathbf{r}_0 \in \Omega$ in the reference space of the k^{th} mesh cell the identity

$$\mathcal{V}_k[\rho](\mathbf{r}_0) = \widehat{\mathcal{V}}_k[\rho \circ \mathbf{R}^k](\xi_0^k) \quad (3.4)$$

holds [63, Thm. 5.5 & Add. 5.6] for an integrable function ρ on \mathcal{C}_k .

The two sections below outline (i) interpolation by a standard basis on both the straight and curved cells present in the volumetric mesh as well as (ii) the specifics of smooth quadratures and their coupling to fast algorithms. (Similar matters for the simple case of the reference box $\widehat{\mathcal{B}}_0$ are addressed in Section 4.2.)

3.2.1. Smooth quadratures and interpolation on triangles

This section is concerned firstly with mappings of the standard simplex into curvilinear mapped triangles and then with interpolation and quadrature of arbitrary *smooth* functions over these regions; methods for singular corrections for target points $\mathbf{r}_0 \in \Omega$ near to or contained in \mathcal{T}_k are deferred to Sections 3.3 and 4.

Mapped triangles

An arbitrary mapped triangle \mathcal{T}_k can be represented (Fig. 3) using a coordinate transformation from the standard 2-simplex

$$\widehat{\mathcal{T}}_0 = \{(\xi, \eta) : \xi, \eta, 1 - \xi - \eta \geq 0\}. \quad (3.5)$$

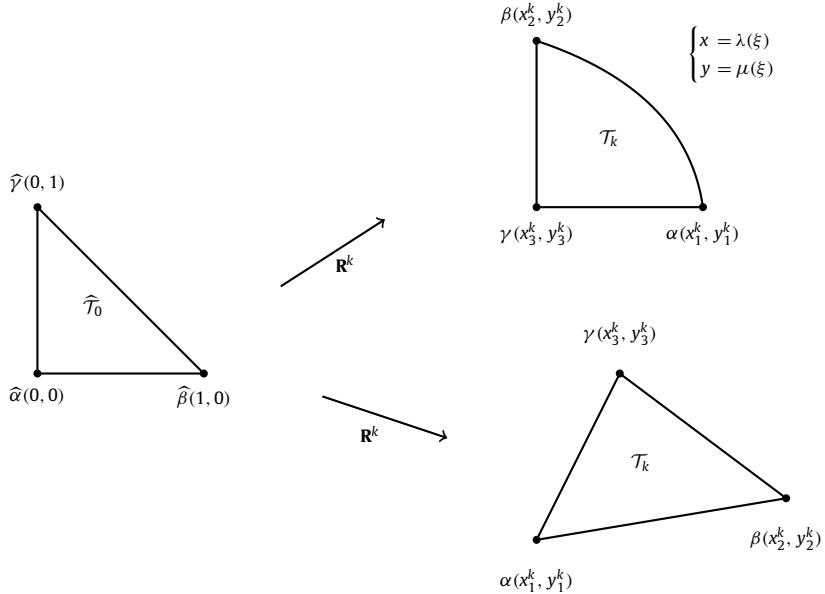


Fig. 3. Depiction of reference space and mapped triangles. Left: $\widehat{\mathcal{T}}_0$ in (ξ, η) -parameter space. Right: Curved (top) and straight (bottom) mapped triangles \mathcal{T}_k in physical space; note that the map $\mathbf{R}^k : \widehat{\mathcal{T}}_0 \rightarrow \mathcal{T}_k$ for a straight-edged triangle \mathcal{T}_k takes the particularly simple form of an affine map.

Noting that from (3.2) the cell \mathcal{C}_k and transformation \mathbf{R}^k correspond to the cell \mathcal{T}_k , for $k \leq N_t$, the transformation for each cell \mathcal{T}_k can be written for such k as

$$\mathbf{R}^k(\xi) = \mathbf{R}^k(\xi, \eta) = \begin{pmatrix} x^k(\xi, \eta) \\ y^k(\xi, \eta) \end{pmatrix}, \quad \xi = \begin{pmatrix} \xi \\ \eta \end{pmatrix}, \quad (3.6)$$

where $x^k(\xi, \eta)$ and $y^k(\xi, \eta)$ denote the x - and y -coordinate mappings from $\widehat{\mathcal{T}}_0$ to \mathcal{T}_k . The mappings \mathbf{R}^k differ in the case that \mathbf{R}^k alternatively maps into a straight-edged or a curved triangle, with the mappings x^k and y^k in the straight-edged triangle case given as affine maps defined such that they correctly map the coordinates of the corners of $\widehat{\mathcal{T}}_0$ to those of \mathcal{T}_k . The meshing strategy described in Section 2 also introduces curved triangles, which in our context consist specifically of curvilinear regions with two straight edges and a single curved edge. For such curved triangles \mathcal{T}_k with corners (x_1^k, y_1^k) , (x_2^k, y_2^k) , and (x_3^k, y_3^k) we let, without loss of generality, the curved edge connect (x_1^k, y_1^k) and (x_2^k, y_2^k) , and introduce the transformation

$$\begin{cases} x^k(\xi, \eta) = (1 - \xi - \eta)x_1^k + \xi x_2^k + \eta x_3^k + \frac{1 - \xi - \eta}{1 - \xi}(\lambda(\xi) - (1 - \xi)x_1^k - \xi x_2^k), \\ y^k(\xi, \eta) = (1 - \xi - \eta)y_1^k + \xi y_2^k + \eta y_3^k + \frac{1 - \xi - \eta}{1 - \xi}(\mu(\xi) - (1 - \xi)y_1^k - \xi y_2^k), \end{cases} \quad (3.7)$$

which can easily be seen is a C^1 -invertible map of the standard simplex $\widehat{\mathcal{T}}_0$ onto \mathcal{T}_k . Here, $\lambda : [0, 1] \rightarrow \mathbb{R}$ and $\mu : [0, 1] \rightarrow \mathbb{R}$ are parametrizations of the individual coordinates of the curved edge connecting (x_1^k, y_1^k) and (x_2^k, y_2^k) that satisfy $\lambda(0) = x_1^k$, $\lambda(1) = x_2^k$, $\mu(0) = y_1^k$, $\mu(1) = y_2^k$. This procedure for mapping $\widehat{\mathcal{T}}_0$ to an arbitrary deformed triangle \mathcal{T}_k is known as the blending function method [64,65] originally introduced in the finite element literature; see also [66].

The mappings \mathbf{R}^k are used in the method to relate integrals over triangular regions \mathcal{T}_k to integrals over the standard simplex $\widehat{\mathcal{T}}_0$: since the mapping \mathbf{R}^k given by (3.6) is a C^1 -invertible mapping over $\widehat{\mathcal{T}}_0$ it is known [63, Thm. 5.5 & Add. 5.6] that the integral over \mathcal{T}_k can be expressed as

$$\int_{\mathcal{T}_k} \rho(\mathbf{r}) dA(\mathbf{r}) = \int_{\widehat{\mathcal{T}}_0} \rho(\mathbf{R}^k(\xi, \eta)) |J^k(\xi, \eta)| d\xi d\eta, \quad (3.8)$$

for an integrable function $\rho : \mathcal{T}_k \rightarrow \mathbb{R}$, where $|J^k|$ denotes the absolute value of the Jacobian determinant of the mapping \mathbf{R}^k ; clearly, Equation (3.8) is equivalent to (3.4) for this class of cell. This elementary change of variables not only enables the use of an orthogonal basis for every cell \mathcal{T}_k , but it also enables efficient means to generate singular corrections (see Section 4.1.1).

Smooth interpolation and quadrature

At the core of the present numerical method for volume potential evaluation are robust and efficient, high-order interpolation and quadrature schemes on each of the cell types. Here we present such a scheme for the simplex using the

well-known Koornwinder polynomial system. Denoting by $P_n^{(\alpha, \beta)}(\xi)$ (see [67, §22] for details) the Jacobi polynomial of degree n that satisfies the ODE

$$(1-x)^2 y'' + (\beta - \alpha - (\alpha + \beta + 2)x)y' + n(n + \alpha + \beta + 1)y = 0, \quad -1 < x < 1,$$

we will make extensive use of the Koornwinder polynomials which are defined, up to maximal total degree p , by

$$K_{nm} = \gamma_{nm} P_{n-m}^{(0, 2m+1)}(1-2\eta) P_m^{(0, 0)}\left(\frac{2\xi}{1-\eta} - 1\right) (1-\eta)^m; \quad m = 0, \dots, n, \quad \text{and} \quad n = 0, \dots, p,$$

where the weights γ_{nm} are chosen so that

$$\int_{\widehat{\mathcal{T}}_0} K_{nm}^2(\xi, \eta) d\xi d\eta = 1.$$

As is well-known [68], the $p(p+1)/2$ polynomials $\{K_{nm} : 0 \leq m \leq n, 0 \leq n < p\}$ form an orthogonal basis for the space \mathcal{P}_{p-1} of polynomials of total degree less than p on the simplex $\widehat{\mathcal{T}}_0$. We will call an element ρ of \mathcal{P}_{p-1} ,

$$\rho(\xi, \eta) = \sum_{n=0}^{p-1} \sum_{m=0}^n a_{nm} K_{nm}(\xi, \eta),$$

a p -th order Koornwinder expansion in keeping with classical results on the error in interpolation of smooth functions by polynomials of degree less than p .

The coefficients of a polynomial $\rho \in \mathcal{P}_{p-1}$ can be related to its values on a discrete set of interpolation nodes

$$I_{\mathcal{T}, p} = \{(\xi_{p,i}, \eta_{p,i}) : 1 \leq i \leq N_p\}, \quad N_p = p(p+1)/2. \quad (3.9)$$

In detail, denoting by \mathbf{a} and $\boldsymbol{\rho}$ the vectors with elements $(\mathbf{a})_{nm} = a_{nm}$ and $(\boldsymbol{\rho})_i = \rho(\xi_{p,i}, \eta_{p,i})$ the coefficients will satisfy

$$\mathbf{V}_p \mathbf{a} = \boldsymbol{\rho}, \quad (3.10)$$

where $(\mathbf{V}_p)_{nm,i} = K_{nm}(\xi_{p,i}, \eta_{p,i})$ is the so-called coefficients-to-values map associated with the Koornwinder polynomials on the nodal set $I_{\mathcal{T}, p}$. Provided the nodal set $I_{\mathcal{T}, p}$ is such that \mathbf{V}_p is nonsingular we denote the values-to-coefficients map $\mathbf{C}_p = \mathbf{V}_p^{-1}$ —while in our case \mathbf{V}_p will always be invertible due to a specific choice of nodal set $I_{\mathcal{T}, p}$, see also [69] for a sufficient condition for a generic nodal set $I_{\mathcal{T}, p}$ to yield an invertible matrix \mathbf{V}_p .

Recent contributions [30] have developed nodal sets $I_{\mathcal{T}, p}$ leading to favorable conditioning of the maps \mathbf{V}_p and \mathbf{C}_p that we utilize, and, making this selection we henceforth denote by $I_{\mathcal{T}, p}$ the Vioreanu-Rokhlin nodes for interpolation by polynomials with total degree less than p . It will also be convenient to introduce the oversampling matrix \mathbf{O}_{p_1, p_2} which maps values at the Vioreanu-Rokhlin nodes $I_{\mathcal{T}, p_1}$ in a Koornwinder expansion of order p_1 to that same expansion's values at the Vioreanu-Rokhlin nodes $I_{\mathcal{T}, p_2}$ for an expansion of order $p_2 \geq p_1$. For a generic smooth function $\rho : \widehat{\mathcal{T}}_0 \rightarrow \mathbb{R}$ we have the approximation

$$\rho(\xi, \eta) \approx \rho_p(\xi, \eta) = \sum_{n=0}^{p-1} \sum_{m=0}^n a_{nm} K_{nm}(\xi, \eta), \quad (3.11)$$

accurate to p -th order. The expansion coefficients a_{nm} in a p -th order Koornwinder expansion are determined for generic smooth functions and polynomials alike via the solution to the system (3.10).

Having considered interpolation, we turn to quadratures of smooth functions: a target \mathbf{r}_0 -independent high-order quadrature rule is required on the simplex for integrals such as those in Equation (3.8) with ρ smooth. (The \mathbf{r}_0 -independence of the rule—that the source nodes are independent of the target—is required for compatibility with fast summation techniques.) For this task, we turn to generalized Gaussian quadrature rules, in which context we recall that one-dimensional Gaussian quadrature yields a specific set $\{(\xi_i, w_i) : i = 1, 2, \dots, N\}$ of N quadrature nodes and weights that can integrate exactly all polynomials of degree at most $2N - 1$ (the weights w_i following from a specific choice of nodal set). In higher dimensions, it appears that perfect Gaussian quadratures are unfortunately not available, but *generalized* Gaussian quadratures have been introduced [31,30] which, for N quadrature nodes in $d > 1$ dimensions, rather than exactly integrating dN functions (as would be the case for a perfect Gaussian quadrature rule), instead only integrate some number of functions greater than N . Fortunately, the interpolation nodes $I_{\mathcal{T}, p}$ introduced above also have associated with them corresponding quadrature weights, thereby furnishing us with a highly-efficient set of quadrature nodes and weights (the *efficiency* of a modified Gaussian quadrature is the ratio of the number of functions integrated exactly to the ideal “Gaussian” number dN)—for details see [30], but we caution that this does not yield a quadrature rule of order p [30, §5]. Thus, for a smooth function ρ on the domain \mathcal{T}_k and in view of (3.8) we use the interpolatory quadrature rule

$$\int_{\mathcal{T}_k} \rho(\mathbf{r}) dA(\mathbf{r}) \approx \sum_{j=1}^{N_p} \rho\left(\mathbf{R}^k(\xi_j, \eta_j)\right) J^k(\xi_j, \eta_j) w_j, \quad \text{where } (\xi_j, \eta_j) \in I_{\mathcal{T},p}, \quad (3.12)$$

termed so because the quadrature nodes coincide with the set $I_{\mathcal{T},p}$ of interpolation nodes.

Oversampled smooth quadratures

An ideal smooth quadrature rule for $\mathcal{V}_k[f](\mathbf{r}_0)$ delivers accurate approximations with error on the level of interpolation of the function f over \mathcal{T}_k , but, unfortunately, large gradients in the integrand due to G over cells \mathcal{T}_k near to \mathbf{r}_0 can lead to quadrature error that dominates that of interpolation of f at any fixed order p . Oversampled quadratures can address this problem, limiting the number of nodal points at which f is required while delivering high accuracy, and we describe next their use. Defining the vector of function samples \mathbf{f}_p^k by

$$(\mathbf{f}_p^k)_i = f\left(\mathbf{R}^k(\xi_{p,i}, \eta_{p,i})\right), \quad \text{where } (\xi_{p,i}, \eta_{p,i}) \in I_{\mathcal{T},p} \text{ for } i = 1, \dots, N_p,$$

so that

$$f(\mathbf{R}^k(\xi, \eta)) \approx \sum_{n=0}^{p-1} \sum_{m=0}^n a_{nm}^k K_{nm}(\xi, \eta) \quad \text{with } \mathbf{V}_p \mathbf{a}^k = \mathbf{f}_p^k,$$

interpolated approximate values of the source f are obtained at nodes $I_{\mathcal{T},q}$ via the oversampling map \mathbf{O}_{p_1, p_2} with $p_1 = p$ and $p_2 = q \geq p$ (see the previous section on smooth interpolation),

$$\mathbf{f}_q^k = \mathbf{O}_{p,q} \mathbf{f}_p^k,$$

where the integer q represents the degree of oversampled quadrature. The quadrature rule (3.12) applied to $\rho(\mathbf{r}) = G(\mathbf{r}, \mathbf{r}_0)f(\mathbf{r})$ for each \mathcal{T}_k satisfying $\mathcal{T}_k \subset \mathcal{C}^{\text{far}}(\mathbf{r}_0)$ is then

$$\mathcal{V}_k[f](\mathbf{r}_0) \approx \sum_{j=1}^{N_q} G(\mathbf{R}^k(\xi_j, \eta_j), \mathbf{r}_0) (\mathbf{f}_q^k)_j J^k(\xi_j, \eta_j) w_j. \quad (3.13)$$

One detail to note is that while quadrature occurs in reference (ξ, η) space in which context the function is not immediately obviously amenable to standard fast algorithms, the sum in (3.13) is nevertheless compatible with FMMs in physical space by viewing the quantity $J^k(\xi_j, \eta_j) w_j$ as modified weights for the discrete inner product.

3.3. A singular quadrature scheme for mapped simplices

This section is devoted to a description of a singular quadrature scheme that can evaluate volumetric integrals over arbitrary star-shaped regions of the plane. While the singular quadrature methods outlined below are quite generic geometrically (though see Remark 3), for our purposes here singular integral evaluation will be required over one of precisely two volumetric regions $\widehat{\mathcal{C}}$: a box $\widehat{\mathcal{C}} = \widehat{\mathcal{C}}_k = \widehat{\mathcal{B}}_0$ and a standard simplex $\widehat{\mathcal{C}} = \widehat{\mathcal{C}}_k = \widehat{\mathcal{T}}_0$ (as is suggested by the notation, we suppress the dependence on k in this section whenever possible). This section begins by presenting the basic theory of converting convolutions of the Green's function with a source function ρ over certain star-shaped domains to integrals over each of those domain's boundaries (using ideas related to Poincaré's lemma), then discusses treatment of the singular kernel in Section 3.3.2 and finally addresses some subtler geometric details in Section 3.3.3. In principle the function ρ could incorporate directly the actual source function f (with mapped argument) arising from the PDE boundary value problem Equation (2.1a) but in practice will be chosen to be related to a member of a given family of orthogonal polynomials, as detailed in Section 3.2.

3.3.1. Poincaré's lemma and volume-to-boundary integral conversion for $\widehat{\mathcal{V}}_k$

It is first useful to recall some terminology. A region \mathcal{S} is called star-shaped if there exists *some* point $\mathbf{r}_* \in \mathcal{S}$ such that every line segment that connects \mathbf{r}_* to any other point $\mathbf{r} \in \mathcal{S}$ *lays entirely in* \mathcal{S} (the region is called star-shaped with respect to \mathbf{r}_*). On the other hand, a star-shaped region is convex if and only if it is star-shaped with respect to every $\mathbf{r}_* \in \mathcal{S}$.

Inspired by elements of the proof of Poincaré's lemma, which utilizes certain maps involving iterated boundary integrals, we make use of an extension of these ideas for integrals involving singular functions, but first state its main holding in a simpler lemma limited to smooth functions. A proof is given in the appendix, using the perhaps more familiar tools of vector calculus (cf. elements of this lemma using the language of differential forms in [70]).

Lemma 1. *Let the closed region $\mathcal{S} \subset \mathbb{R}^2$ be star-shaped with respect to the origin with a piecewise smooth boundary $\partial\mathcal{S}$ and assume $\rho : \mathcal{S} \rightarrow \mathbb{R}$ is continuously differentiable on \mathcal{S} . Then the relation*

$$\int_{\mathcal{S}} \rho(\mathbf{r}) dA = \oint_{\partial\mathcal{S}} \left(\int_0^1 t \rho(t\mathbf{r}) dt \right) \mathbf{r} \times \boldsymbol{\tau} ds \quad (3.14)$$

holds, where s denotes the arclength, $\boldsymbol{\tau}$ denotes the unit tangent vector of $\partial\mathcal{S}$, and $\mathbf{r} \times \boldsymbol{\tau}$ is understood as a scalar.

The integrals resulting from Lemma 1 have at times been referred to as dilation integrals (with dilation with respect to the star point), with e.g. reference [48] relying precisely (up to a fixed geometry-dependent translation) on the relation (3.14) to generate accurate quadratures of smooth functions. However, in the context of volume potentials with weakly singular kernels, such approaches [57,49,53,54,50,51] have previously faced a number of challenges related to star-shapedness, location of quadrature points, and presence of singularities (see the discussion in the Introduction); in some cases [57,49] the applicability is restricted to Poisson problems.

Remark 3. For a convex reference cell $\mathcal{S} = \widehat{\mathcal{C}}$, which in the proposed method are the only domains over which quadrature is performed, Corollary 1 below applies to every point $\xi_* \in \widehat{\mathcal{C}}$, while this may unfortunately not hold for a generic star-shaped region. It is a significant strength of the present mapped-domain approach, in contrast to the possibility of applying this lemma in physical space (for which case difficulties may arise when targets lay in certain subsets of nonconvex cells), that the Poincaré lemma-related ideas can be successfully applied for every point ξ_* . More generally, methods based on transformation of volume integrals to surface integrals potentially suffer when integration domains are non-convex as they generally result in quadrature points laying outside the integration domain result [57,45,49,46,53,54], a long-recognized issue in multidimensional quadrature [38,71]. As a more minor matter, the presence of such quadrature nodes implies the existence of negative quadrature weights, which is generally considered unfavorable in view of stability concerns.

The following corollary of Lemma 1 is used in what follows to express the volume potentials in terms of boundary integrals; its proof is given in the appendix.

Corollary 1. Let $K = K(\mathbf{r})$ denote a weakly-singular kernel function which is continuously differentiable for $\mathbf{r} \in \mathbb{R}^2 \setminus \{\mathbf{0}\}$, assume the closed region $\mathcal{S} \subset \mathbb{R}^2$ has a piecewise smooth boundary $\partial\mathcal{S}$, and assume the function $\rho : \mathcal{S} \rightarrow \mathbb{R}$ is continuously differentiable on \mathcal{S} . Then for each $\mathbf{r}_* \in \mathcal{S}$ such that \mathcal{S} is star-shaped with respect to \mathbf{r}_* , we have

$$\int_{\mathcal{S}} K(\mathbf{r} - \mathbf{r}_*) \rho(\mathbf{r}) dA = \oint_{\partial\mathcal{S}} \left(\int_0^1 t K(t(\mathbf{r} - \mathbf{r}_*)) \rho(t(\mathbf{r} - \mathbf{r}_*) + \mathbf{r}_*) dt \right) ((\mathbf{r} - \mathbf{r}_*) \times \boldsymbol{\tau}) ds. \quad (3.15)$$

The proposed methodology identifies for every singular and near-singular target $\mathbf{r}_0 \in \Omega$ a point $\mathbf{r}_* = \mathbf{r}_*(\mathbf{r}_0) \in \mathcal{C}_k$ using the rule

$$\xi_* = \operatorname{argmin}_{\xi \in \widehat{\mathcal{C}}_k} |\mathbf{R}^k(\xi) - \mathbf{r}_0| \quad \text{and} \quad \mathbf{r}_* = \mathbf{R}^k(\xi_*). \quad (3.16)$$

Having identified $\xi_* \in \widehat{\mathcal{C}}_k$ we now apply Corollary 1 and then describe quadrature rules for evaluating certain resulting integrals in the reference domain $\widehat{\mathcal{C}} = \widehat{\mathcal{C}}_k$.

Remark 4. A subtle but important point for near-singular targets \mathbf{r}_0 is that the star point \mathbf{r}_* must be selected as the closest point in $\partial\mathcal{C}_k$ to \mathbf{r}_0 despite quadrature occurring over $\widehat{\mathcal{C}}$, as this point corresponds to the minimum distance d that arises as an argument to the singular kernel (cf. Equation (3.24) and Fig. 5). By adapting to this point the quadrature rule is able to deliver optimal accuracy.

Writing $\widehat{\mathcal{V}}_k[\rho]$ in the form

$$\widehat{\mathcal{V}}_k[\rho] = \int_{\widehat{\mathcal{C}}} G \left(\mathbf{R}^k(\xi - \xi_* + \xi_*) - \mathbf{R}^k(\xi_0) \right) J^k(\xi) \rho(\xi) dA$$

and making the selections of region $\mathcal{S} = \widehat{\mathcal{C}}$, kernel $K(\boldsymbol{\tau}) = G(\mathbf{R}^k(\boldsymbol{\tau} + \xi_*) - \mathbf{R}^k(\xi_0))$ and smooth source $J^k(\xi) \rho(\xi)$ in Corollary 1, we obtain

$$\begin{aligned} \widehat{\mathcal{V}}_k[\rho](\xi_0) = & \oint_{\partial\widehat{\mathcal{C}}} \left(\int_0^1 t G(\mathbf{R}^k(\xi_* + t(\xi - \xi_*)), \mathbf{R}^k(\xi_0)) \times \right. \\ & \left. \times J^k(\xi_* + t(\xi - \xi_*)) \rho(\xi_* + t(\xi - \xi_*)) dt \right) ((\xi - \xi_*) \times \boldsymbol{\tau}) ds. \end{aligned} \quad (3.17)$$

The expression (3.17) for the volume potential $\widehat{\mathcal{V}}_k[\rho]$ can be written in the simplified form

$$\widehat{\mathcal{V}}_k[\rho](\xi_0) = \oint_{\partial\widehat{\mathcal{C}}} I^k[\rho](\xi, \xi_0) ((\xi - \xi_*) \times \boldsymbol{\tau}) ds, \quad (3.18)$$

where $I^k[\rho](\xi, \xi_0)$ is defined as the inner integral of the iterated integral:

$$I^k[\rho](\xi, \xi_0) = \int_0^1 t G(\mathbf{R}^k(\xi_* + t(\xi - \xi_*)), \mathbf{R}^k(\xi_0)) J^k(\xi_* + t(\xi - \xi_*)) \rho(\xi_* + t(\xi - \xi_*)) dt. \quad (3.19)$$

The presence of the Green function (near-)singularity at $t = 0$, for all ξ, ξ_0 , in the one-dimensional integral in (3.19) suggests that one-dimensional quadrature schemes for the integral $I^k[\rho](\xi, \xi_0)$ may be effective, which we outline in the next section. Before discussing these quadrature schemes we first introduce the quadrature scheme for the outer integral in (3.18).

Since $\partial\widehat{\mathcal{C}}$ is piecewise smooth, it follows that $I^k[\rho](\xi, \xi_0)$ is piecewise smooth as a function of $\xi \in \partial\widehat{\mathcal{C}}$, implying that the integrand of the circulation integral in (3.18) is likewise piecewise smooth. The quadrature rule we propose for $\widehat{\mathcal{V}}_k[\rho](\xi_0)$ is developed by segmenting the boundary $\partial\widehat{\mathcal{C}}$ into $M_{\partial\widehat{\mathcal{C}}}$ intervals, and obtaining a composite Gauss-Legendre rule by application of a P -point Gauss-Legendre quadrature rule on each interval, with P a fixed integer (we use the selection $P = 16$ for the numerical examples in this article). We thus have the quadrature rule

$$\widehat{\mathcal{V}}_k[\rho](\xi_0) \approx \sum_{j=1}^{PM_{\partial\widehat{\mathcal{C}}}} I^k[\rho](\xi_j, \xi_0) ((\xi_j - \xi_*) \times \boldsymbol{\tau}_j) w_j, \quad (3.20)$$

where $\{\xi_j\}_{j=1}^{PM_{\partial\widehat{\mathcal{C}}}} \subset \partial\widehat{\mathcal{C}}$ and $\{w_j\}_{j=1}^{PM_{\partial\widehat{\mathcal{C}}}}$ are the quadrature nodes and weights. The error arising from use of this quadrature can depend in an important manner on the proximity of ξ_0 relative to the boundary $\partial\widehat{\mathcal{C}}$, and, in particular, we find that a uniform distribution of composite Gauss-Legendre intervals is not sufficient to achieve desired accuracies. Further details on boundary discretization are given in Section 3.3.3 where an optimal interval distribution is described which may depend on the target ξ_0 .

3.3.2. Singular and near-singular quadrature rules for $I^k[\rho]$

While use of Corollary 1 transforms the volume integral over the cell $\widehat{\mathcal{C}}$ to an iterated integral over the boundary of the cell $\partial\widehat{\mathcal{C}}$, the integrand of the inner integral remains potentially non-smooth (indeed, potentially singular depending on the strength of the kernel singularity). The one-dimensional quadrature scheme outlined in what follows is tailored to the known singular behavior of the kernel: crucially, only the endpoint *asymptotic behavior* of the integrand is relevant to quadrature rule selection, making the proposed scheme applicable to many kernel functions, including, but not limited to, those arising in a variety of Green functions of mathematical physics (in particular, the kernel function could be more singular than the Green functions typically encountered in elliptic PDEs).

Remark 5. For convenience, we refer to the one-dimensional integrals $I^k[\rho]$ as being singular (for $\xi_0 \in \widehat{\mathcal{C}}_k$) or near-singular (for $\xi_0 \notin \widehat{\mathcal{C}}_k$), even though, in view of the t -factor present in the integrand of $I^k[\rho]$, it is possible that the integrand is not truly singular; for example, for the elliptic PDEs we consider in this article, the small- t asymptotic behavior of G in I^k is merely logarithmic ($s(t) = \log t$ in (3.21) and $S(t) = \log t$ in (3.23) below) and the integrand of $I^k[\rho]$ is thus merely non-smooth. Indeed the t -factor, which could be loosely viewed as an analogue of weights arising in polar changes of variables, allows the treatment of highly singular kernels by the proposed methodology. In any case, for the PDE kernels considered in the examples of this article, singular and near-singular quadrature rules are still required for high-order accuracy (i.e. the t -factor is implicitly treated as part of the smooth component of the integrand ($\phi(t)$ in (3.21) and $k_2(t)$ in (3.23) below)).

In detail, the primary difficulty in accurate evaluation of the $I^k[\rho](\xi, \xi_0)$ integral is that, in view of the nature of the integrand in I^k , the target point ξ_0 will lie precisely at (in the case $\xi_0 \in \widehat{\mathcal{C}}$) or instead near (in the case $\xi_0 \notin \widehat{\mathcal{C}}$) one end of the line segment connecting ξ_* ($t = 0$) and $\xi \in \partial\widehat{\mathcal{C}}$ ($t = 1$). In the singular case, there is possibly an integrable singularity in the integrand at the left ($t = 0$) endpoint of the integration interval due to the singular nature of the Green function G . Problems remain even in the case when there is a positive distance of ξ_0 to $\widehat{\mathcal{C}}$, since steep gradients, again at $t = 0$,

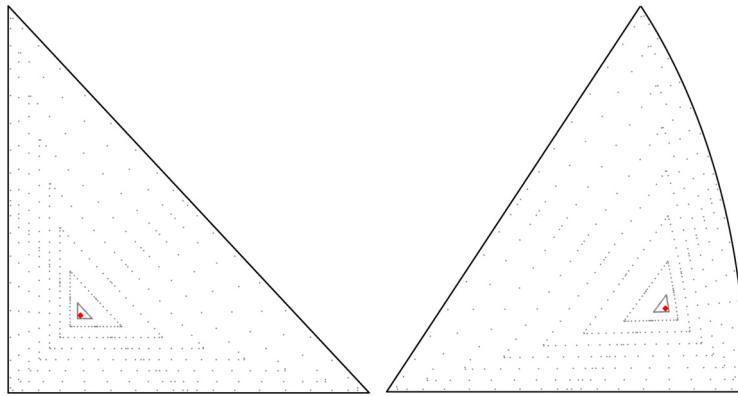


Fig. 4. In reference (left) and physical (right) space for a curvilinear cell, the high-order quadrature nodes (arising from a 5th-order corrected trapezoidal rule for the I^k integral and a 10th-order composite Gauss-Legendre rule for the integral over $\partial\widehat{C}$) are displayed that result from the proposed quadrature scheme applied to a singular target point $\mathbf{r}_0 \in \mathcal{C}$ (marked in red).

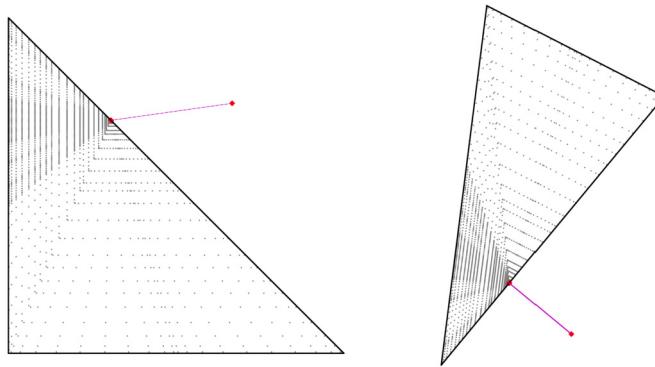


Fig. 5. In reference (left) and physical (right) space for a straight triangle cell, the high-order quadrature nodes (arising from a 10th-order modified Gaussian rule for the I^k integral and a 10th-order composite Gauss-Legendre rule for the integral over $\partial\widehat{C}$) are displayed that result from the proposed quadrature scheme applied to a near-singular target point $\mathbf{r}_0 \notin \mathcal{C}$ (marked in red). The dashed line segment connects the target \mathbf{r}_0 to the nearest physical point $\mathbf{r}_* \in \mathcal{C}$ which is selected as the star-point (also marked in red); in reference domain $\xi_* = (\mathbf{R})^{-1}(\mathbf{r}_*)$ is not the closest point to ξ_0 since the map \mathbf{R} is affine (see also Remark 4).

can be challenging to accurately resolve with any fixed target-independent quadrature scheme. The high-order accurate quadrature scheme described in this article depends on a corresponding high-order accurate quadrature rule for the integral I^k , of which a wide variety of suitable schemes for singularity behavior of various types have been developed over many years, e.g. [72–74] (see [75] and references therein for discussion of early work in this direction). The integrals I^k , when the integrand is non-smooth, are amenable to the use of existing corrected trapezoidal quadrature rules for functions with known (singular) endpoint behavior and we utilize in this article the rules of reference [72] for singular quadrature—rules for integrating smooth functions multiplied by singular functions of logarithmic and (integrable) inverse-power type; for near-singular quadrature we use the rules introduced in reference [74].

We first detail the application of the Alpert [72] quadrature rule to I^k for self-interaction (singular) terms. The Alpert rules provide endpoint-corrected trapezoidal quadrature nodes and weights for integrals of functions $h(t) : (0, 1] \rightarrow \mathbb{R}$ of the form

$$h(t) = \phi(t)s(t) + \psi(t), \quad (3.21)$$

where $\phi(t), \psi(t) \in C^k[0, 1]$ and $s(t) \in C(0, 1]$ is an integrable function that is singular at $t = 0$. Letting the set of pairs $\{(t_i, v_i), 1 \leq i \leq m_P\}$ denote, for a given integer P (the selection $P = 16$ is made in all numerical results in this article), the nodes and weights of a P -th order convergent Alpert quadrature rule for a given $s(t)$, we obtain from (3.19)

$$I^k[\rho](\xi, \xi_0) \approx \sum_{i=1}^{m_P} t_i G \left(\mathbf{R}^k(\xi_* + t_i(\xi - \xi_*)), \mathbf{R}^k(\xi_0) \right) J^k(\xi_* + t_i(\xi - \xi_*)) \rho(\xi_* + t_i(\xi - \xi_*)) v_i, \quad (3.22)$$

which is a P -th order approximation to $I^k[g]$ and where we used the selection $\xi_* = \xi_0$ that has been made for the case of singular target points ($\xi_0 \in \widehat{\mathcal{C}}$).

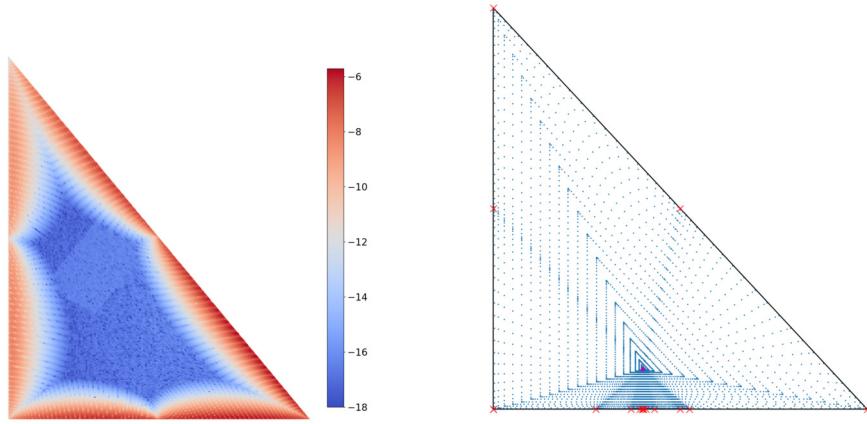


Fig. 6. Left: Additional error (in evaluation of a logarithmic-kernel volume potential) that arises from a (naïve) uniform distribution of intervals for purposes of a boundary discretization in Equation (3.18), for each target point in the interior of $\widehat{\mathcal{C}}$ a unit simplex (demonstrating the need for some special care in handling target points laying close to a boundary); the color indicates the base-10 logarithm of the ‘error’, compared to the results of the proposed non-uniform boundary interval distribution that is *target-adapted*. Right: Singular volumetric quadrature nodes (blue) and boundary interval endpoints (red) resulting from the proposed target-adapted methodology when the target (magenta) lays close to the boundary of the reference cell $\widehat{\mathcal{C}}$; the proposed boundary interval distribution produces uniformly-accurate volume potential values over the unit simplex to an error level of $\approx 10^{-14}$ (see Table 1).

In the near-singular case ($\zeta_0 \notin \widehat{\mathcal{C}}$) we turn to modified Gaussian quadrature rules for integrands with known singular behavior $S(t)$ as $t \rightarrow 0^+$. Such methods provide, for a given d which lays in intervals of the form $[10^{-q-1}, 10^{-q}]$ for positive integer q , quadrature nodes and weights $\{(\tilde{t}_i, \tilde{v}_i), 1 \leq i \leq m'_p\}$ so that the near-singular rule

$$\int_0^1 (k_1(t) + k_2(t)S(t+d)) dt \approx \sum_{i=1}^{m'_p} (k_1(\tilde{t}_i) + k_2(\tilde{t}_i)S(\tilde{t}_i+d)) \tilde{v}_i, \quad (3.23)$$

holds to high accuracy (to within an accuracy ε of $\varepsilon \approx 10^{-15}$) for k_1 and k_2 polynomials of degree at most P (were the quadrature rule exact for such polynomials it would be a perfect Gaussian, hence the ‘modified’ moniker). Methods and theory are described in [74] for generation of such quadratures for a variety of singular behaviors; for the two-dimensional elliptic PDE demonstrations of the present work, we rely on pre-computed rules [3] for the case $S(t) = \log(t)$, where for the selection $P = 10$ made everywhere in this article a total of $m'_p = 24$ nodes are required. These quadrature rules are parametrized by the “near-singularity distance” d defined by the expression

$$d = \left| \mathbf{R}^k(\zeta_*) - \mathbf{R}^k(\zeta_0) \right|, \quad (3.24)$$

which is the minimum distance that can occur for arguments to the kernel for this integration domain and target \mathbf{r}_0 (cf. (3.16)). We thus have the quadrature rule for $I^k[\rho](\zeta, \zeta_0)$,

$$I^k[\rho](\zeta, \zeta_0) \approx \sum_{i=1}^{m'_p} \tilde{t}_i G \left(\mathbf{R}^k(\zeta_* + \tilde{t}_i(\zeta - \zeta_*)), \mathbf{R}^k(\zeta_0) \right) J^k(\zeta_* + t(\zeta - \zeta_*)) \rho(\zeta_* + \tilde{t}_i(\zeta - \zeta_*)) \tilde{v}_i. \quad (3.25)$$

3.3.3. Close evaluation

The quadrature rule (3.20) can lead to a highly-accurate quadrature rule for the representation (3.18) of $\widehat{V}_k[\rho](\zeta_0)$, provided an appropriate distribution of quadrature nodes ζ_j on the boundary $\partial\widehat{\mathcal{C}}$ are selected in the quadrature rule (3.20) for (3.18). However, simply using a uniform distribution of boundary intervals (e.g. for $\widehat{\mathcal{C}} = \widehat{\mathcal{T}}_0$ a triangle parametrized by $Z = Z(t)$, $Z : [0, 2\pi] \rightarrow \partial\widehat{\mathcal{C}}$, using interval endpoints equi-spaced in t in each of $[0, 2\pi/3]$, $[2\pi/3, 4\pi/3]$, and $[4\pi/3, 2\pi]$) will not always lead to a high-order approximation of the desired volume potential. Possible loss of accuracy can be understood by observing, in the test depicted in Fig. 6, that we are building a quadrature rule on $\widehat{\mathcal{C}}$ centered around the target ζ_* , with the value at each quadrature node ζ_j in (3.18) being given by the integral $I^k[\rho](\zeta, \zeta_0)$ along a ray from ζ_* to ζ_j . However, an equi-arc length interval distribution does not uniformly cover the angular variable in the coordinate system with origin ζ_* , in which context steep gradients arise in $I^k[\rho](\zeta, \zeta_0)$ with respect to ζ , and poor accuracy in the quadrature rule (3.20) can result—a deficiency that we now remedy.

Since potential loss of accuracy in the quadrature rule (3.20) arises due to the described inadequate coverage of the angular variable in the coordinate system with origin ζ_* , with this effect being most pronounced at the closest point $\zeta_1 \in \partial\widehat{\mathcal{C}}$ to ζ_* , we refine the boundary discretization in a vicinity of ζ_1 and thus recover volume quadrature nodes that are suitably equi-distributed. After first identifying the parametric location t_1 of ζ_1 in the parametrization of $\widehat{\mathcal{C}}$, the method proceeds

Table 1

Single-cell test of the proposed singular quadrature scheme: evaluation of a volume potential over a straight (mapped) triangle Ω with vertices located at $(-0.618, -0.312)$, $(-0.825, -0.311)$, and $(-0.802, -0.516)$. For each integer p , $\|e\|_{\ell^\infty}$ denotes maximum error in the volume potential $\mathcal{V}[K_{nm} \circ \mathbf{R}^{-1}](\mathbf{r}_0)$ of the $p(p+1)/2$ -numbered polynomials $K_{nm} \in \{K_{nm} : 0 \leq m \leq n, 0 \leq n < p\}$, each over all of the $p(p+1)/2$ -numbered points $\mathbf{r}_0 = \mathbf{R}(\xi_0)$, $\xi_0 \in I_{\mathcal{T},p}$ (see (3.9)). Ground ‘truth’ is a highly-adaptive (and highly-expensive) multidimensional quadrature rule unrelated to the methods of this article, whose associated error is no greater than $\approx 10^{-14}$. The minimum distance from any point ξ_0 to $\partial\widehat{\mathcal{T}}_0$ is denoted by \widehat{d}_{\min} .

p	$\ e\ _{\ell^\infty}$	\widehat{d}_{\min}
4	4.0e-15	7.1e-02
5	5.0e-15	4.8e-02
6	7.0e-15	3.1e-02
7	1.0e-14	2.3e-02
8	1.1e-14	1.9e-02
9	1.1e-14	1.5e-02
10	1.2e-14	1.3e-02

by placing a graded sequence of intervals near to t_1 . Letting t_1^- and t_1^+ denote the parametric location of the nearest (still uniformly-distributed) composite interval endpoints that surround t_1 , the method introduces additional composite regions with endpoints in the set

$$\Pi(t_1) := \left\{ t^j : t^j := t_1 \pm |t_1^+ - t_1^-| R^j, \quad j = 1, \dots, N \right\}; \quad \text{with} \quad R = \frac{1}{4}, \quad N = 5.$$

A depiction of the boundary discretization and associated volumetric quadrature nodes can be seen in Fig. 6.

Table 1 displays results of a test which demonstrates that this refinement strategy for the volume potential yields accuracy of approximately 13 digits for targets close to cell boundaries. The test consists of evaluation of a Laplace volume potential, $G(\mathbf{r}, \mathbf{r}_0) = -\frac{1}{2\pi} \log |\mathbf{r} - \mathbf{r}_0|$, over a single (mapped) triangle Ω at specific points corresponding to Koornwinder interpolation nodes, with associated map denoted by $\mathbf{R} : \widehat{\mathcal{T}}_0 \rightarrow \Omega$. This experiment has direct relevance to evaluation of the volume potential problem (1.1) since, as detailed in Section 4.1.1 below, values of $\mathcal{V}[f]$ at these nodes can be used to produce high-order accurate interpolation of the volume potential throughout a given cell \mathcal{T}_k . The results of the experiment summarized in Fig. 8 provide a somewhat more challenging test as it includes some target points that lay up to two orders of magnitude closer to the cell boundary.

3.3.4. Final rule for local corrections

Having developed singular and near-singular quadrature rules for the one-dimensional integrals I^k (3.19), we provide the full singular and near-singular rules for the volume potential $\mathcal{V}_k[\rho]$. In the singular case ($\xi_0 \in \widehat{\mathcal{C}}$), inserting the quadrature rule (3.22) into (3.20) we have

$$\begin{aligned} \widehat{\mathcal{V}}_k[\rho](\xi_0) \approx \sum_{j=1}^{PM_{\partial\mathcal{C}}} \left(\sum_{i=1}^{m_p} t_i G\left(\mathbf{R}^k(\xi_* + t_i(\xi_j - \xi_*)), \mathbf{R}^k(\xi_0)\right) \times \right. \\ \left. \times J^k(\xi_* + t_i(\xi_j - \xi_*)) \rho(\xi_* + t_i(\xi_j - \xi_*)) v_i \right) ((\xi_j - \xi_*) \times \boldsymbol{\tau}_j) w_j. \end{aligned} \quad (3.26)$$

Similarly, for the near-singular case ($\xi_0 \notin \widehat{\mathcal{C}}$), inserting the quadrature rule (3.25) into (3.20) we have

$$\begin{aligned} \widehat{\mathcal{V}}_k[\rho](\xi_0) \approx \sum_{j=1}^{PM_{\partial\mathcal{C}}} \left(\sum_{i=1}^{m'_p} \tilde{t}_i G\left(\mathbf{R}^k(\xi_* + \tilde{t}_i(\xi_j - \xi_*)), \mathbf{R}^k(\xi_0)\right) \times \right. \\ \left. \times J^k(\xi_* + \tilde{t}_i(\xi_j - \xi_*)) \rho(\xi_* + \tilde{t}_i(\xi_j - \xi_*)) \tilde{v}_i \right) ((\xi_j - \xi_*) \times \boldsymbol{\tau}_j) w_j. \end{aligned} \quad (3.27)$$

Examples of the resulting quadrature nodes (for the case of the reference domain a unit simplex $\widehat{\mathcal{T}}_0$) for each of the quadrature rules (3.26) and (3.27) are provided in Fig. 4 and Fig. 5, respectively.

4. Efficient generation and application of singular corrections

This section completes the description of a singular quadrature-corrected scheme for the volume potential with an emphasis on *efficiency* that has not as-yet been considered per se: while on the one hand for cells \mathcal{T}_k that are well-separated from the target point \mathbf{r}_0 the smooth quadrature scheme described in Section 3.2 that leads to the approximation (3.13) is accurate (and amenable to FMM acceleration), and the resulting rules for singular and near-singular quadrature corrections (respectively (3.26) and (3.27)) from Section 3.3 accurately evaluate contributions to $\mathcal{V}[f]$ from nearby cells, on the other hand we have still yet to describe an efficient scheme. By ‘efficient’ we really mean simultaneously that

- (i) The method is ‘node-efficient’, that is, efficient with respect to the required number of degrees of freedom per cell,
- (ii) The required local corrections are cheap to *apply* (i.e. cheap relative to the FMM call), and
- (iii) Any pre-computed local corrections are themselves cheap to *generate*.

To motivate our approach we note that simply applying the methods of Section 3.2 to the source density f is *inefficient* since, while accurate, the associated points where values of the source density f are required, the singular quadrature nodes, are both numerous and dependent on the target itself. Our approach, rather, is to apply the singular quadrature methods in Section 3.3 to the Koornwinder basis elements (which, recalling Section 3.2.1 provide a high-order basis for approximation of arbitrary smooth functions on \mathcal{T}_k). The singular corrections are highly efficient for repeated application of the Newton potential since they are (a) *Local* in the sense that the only cells which contribute to a singular or near-singular correction at a target \mathbf{r}_0 are those cells which are a subset of $\mathcal{C}^{\text{self}}(\mathbf{r}_0) \cup \mathcal{C}^{\text{near}}(\mathbf{r}_0)$ and are also (b) *Data-sparse*, as a linear map from source function f values at Koornwinder interpolation nodes directly to local correction values.

In more detail, as a consequence of the local correction methods described in this section and the smooth oversampled quadratures described in Section 3.2 we conclude that the total degrees of freedom (source function evaluation points) of the scheme for (1.1) number

$$\text{ndofs} = \frac{p(p+1)}{2} N_t + p^2 N_b, \quad (4.1)$$

while the source points for the FMM number

$$\text{nsrcs} = \frac{q(q+1)}{2} N_t + q^2 N_b, \quad q \geq p, \quad (4.2)$$

both for the order p scheme over a mesh with N_t triangular regions and N_b boxes. Denoting by $\Pi_k(\mathbf{r}_0) : \mathbb{R}^{N_s} \rightarrow \mathbb{R}$ the linear map from N_s source function values on a cell \mathcal{C}_k to scalar correction (i.e. the high-order accurate value of \mathcal{V}_k obtained via singular quadrature less the inaccurate contribution from the oversampled smooth quadrature described in Section 3.2.1 for triangles and Section 4.2 below for boxes) for the volume potential on \mathcal{C}_k evaluated at \mathbf{r}_0 , we have either $N_s = p(p+1)/2$ (in the case $\widehat{\mathcal{C}}_k = \widehat{\mathcal{T}}_0$) or $N_s = p^2$ (in the case $\widehat{\mathcal{C}}_k = \widehat{\mathcal{B}}_0$) with, of course, an $\mathcal{O}(1)$ number of nontrivial correction maps $\Pi_k(\mathbf{r}_0)$ per target point due to the $\mathcal{O}(1)$ cardinality of $\mathcal{C}_k^{\text{near}}(\mathbf{r}_0)$ (see Section 3.1). Assembling all nontrivial corrections Π_k into a sparse matrix results in a final singular correction scheme where the overwhelming majority of computational effort is spent in the highly-efficient FMM stage, and which requires a conservative number of source evaluation nodes; the overall performance of the solver is demonstrated in Table 3.

Remark 6. Numerical experiments show that a fixed, twofold ($q = 2p$) oversampling allows for smooth quadrature error for $\mathcal{V}_k[f]$ that is, as desired, dominated by interpolation error of f for the smooth (Stokes and Laplace) kernels and for the error levels presented in this article (thus, for such kernels the selection $q = 2p$ was made in producing the numerical results), though see also [62] for other oversampling schemes. In particular, to achieve higher accuracies than those presented in our numerical experiments, an increase in the fixed oversampling rate is sometimes useful—increasing slightly the number of FMM source points, but not the overall number of degrees of freedom, per (4.1) and (4.2). Furthermore, oversampling and adequate identification of an appropriately-sized near-field region is more critical to obtain high accuracies for more sharply-peaked kernels, such as Helmholtz kernels. We leave detailed consideration of these matters to future study.

Point (iii) above relates to a question of perennial concern in potential theoretic methods: the efficient generation of corrections for singular and near-singular targets (see reference [62] where adaptive quadrature is performed for each near-singular target, at significant cost, and see also references [10,11,41,42,76,77]); the cost of generating such corrections can be burdensome, even while the resulting per-use costs, say, in an iterative solver seem highly favorable. (In our context the cost is due both to the generation of target-adapted quadrature nodes and weights and also to the evaluation of the orthogonal polynomial family at these nodes.) We address this issue by showing that as a result of our use of mappings from a common reference-space domain, the reference-space quadratures and orthogonal polynomial values can be re-used across elements even as the mappings vary, with the favorable implication that the only required per-element computation for generating the (near-)singular corrections consist of evaluation of the Green function.

4.1. Triangles and Koornwinder systems

For a given target \mathbf{r}_0 and for each nearby cell \mathcal{T}_k , i.e. \mathcal{T}_k satisfying $\mathcal{T}_k \subset \mathcal{C}^{\text{self}}(\mathbf{r}_0) \cup \mathcal{C}^{\text{near}}(\mathbf{r}_0)$, the method begins by approximating the source function f on \mathcal{T}_k using the p -th order Koornwinder expansion

$$f(\mathbf{R}^k(\xi, \eta)) \approx \sum_{n=0}^{p-1} \sum_{m=0}^n a_{nm}^k K_{nm}(\xi, \eta), \quad (4.3)$$

where the coefficients \mathbf{a}^k solve the system

$$\mathbf{V}_p \mathbf{a}^k = \mathbf{f}_p^k, \quad (4.4)$$

with $(\mathbf{f}_p^k)_i = f(\mathbf{R}^k(\xi_{p,i}, \eta_{p,i}))$, $i = 1, 2, \dots, N_p$, the vector of function values at the Koornwinder interpolation nodes of (3.9). Substituting this expansion into the volume potential (3.4) we obtain the high-order approximation to $\mathcal{V}_k[f](\mathbf{r}_0)$,

$$\begin{aligned} \mathcal{V}_k[f](\mathbf{r}_0) &\approx \int_{\widehat{\mathcal{T}}_0} G(\mathbf{R}^k(\xi, \eta), \mathbf{R}^k(\xi_0^k, \eta_0^k)) \sum_{n=0}^{p-1} \sum_{m=0}^n a_{nm}^k K_{nm}(\xi, \eta) J^k(\xi, \eta) d\xi d\eta \\ &= \sum_{n=0}^{p-1} \sum_{m=0}^n a_{nm}^k \int_{\widehat{\mathcal{T}}_0} G(\mathbf{R}^k(\xi, \eta), \mathbf{R}^k(\xi_0^k, \eta_0^k)) K_{nm}(\xi, \eta) J^k(\xi, \eta) d\xi d\eta \\ &= \sum_{n=0}^{p-1} \sum_{m=0}^n a_{nm}^k \widehat{\mathcal{V}}_k[K_{nm}](\xi_0^k), \end{aligned} \quad (4.5)$$

where

$$\xi_0^k = (\xi_0^k, \eta_0^k)^T = (\mathbf{R}^k)^{-1}(\mathbf{r}_0)$$

denotes the \mathcal{T}_k -reference-space location of the target \mathbf{r}_0 . Writing the volume potential $\mathcal{V}_k[f](\mathbf{r}_0)$ in terms of reference-space potentials $\widehat{\mathcal{V}}_k$ allow for singular and near-singular correction of the smooth quadrature rule (3.13) for cells \mathcal{T}_k that lay close to \mathbf{r}_0 . The corrections can be pre-computed for each element in the p -th order Koornwinder using the methods of Section 3.3—specifically, if $\mathbf{r}_0 \in \mathcal{T}_k$ then rule (3.26) is used to compute $\widehat{\mathcal{V}}_k[K_{nm}](\xi_0^k)$ while otherwise rule (3.27) is used.

This completes the description of how singular and near-singular corrections can be generically pre-computed; in what follows we consider optimizations that can be obtained for the fortunately-typical case of fixed reference-space target locations. We describe here the case for prescribed reference-space target points $\xi_0 = (\xi_0, \eta_0) \in I_{\mathcal{T},p}$, where $I_{\mathcal{T},p}$ are the Vioreanu-Rokhlin interpolation nodes (see also Section 3.2.1); this ensures the possibility of high-quality interpolation of the resulting function $\mathcal{V}[f]$ throughout the domain (which is useful, for example, in the context of a non-linear or time-dependent PDE)—see also the experiment in Section 5 corresponding to Table 3. The ideas are not restricted to a specific set of interpolation nodes, and the method proposed in Section 4.1.1 could be used e.g. for fixed reference-space locations on the curved boundary component of a curvilinear cell—i.e. for evaluation of $\mathcal{V}[f](\mathbf{r}_0)$ with $\mathbf{r}_0 \in \partial\Omega$. Optimized pre-computation of singular corrections are generalized later to near-singular targets.

4.1.1. Efficient pre-computation of singular and near-singular corrections

An observation that leads to substantial efficiency gains for the proposed method is that the quadratures for every cell \mathcal{T}_k occur on the same reference cell $\widehat{\mathcal{T}}_0$. Since every triangular mesh cell is mapped to the standard simplex $\widehat{\mathcal{T}}_0$ and the target nodes are fixed in reference space, the same (expensive to generate) singular quadratures and Koornwinder polynomial values at these nodes can be used for each target point for every cell. To see the implications for computational cost savings, observe that from (4.5),

$$\mathcal{V}_k[f](\mathbf{r}_0) \approx \sum_{n=0}^{p-1} \sum_{m=0}^n a_{nm}^k \widehat{\mathcal{V}}_k[K_{nm}](\xi_0),$$

with

$$\widehat{\mathcal{V}}_k[K_{nm}](\xi_0) = \int_{\widehat{\mathcal{T}}_0} G(\mathbf{R}^k(\xi), \mathbf{R}^k(\xi_0)) K_{nm}(\xi) J^k(\xi) dA(\xi). \quad (4.6)$$

Clearly, the singular point is $\xi = \xi_0$ for every cell \mathcal{T}_k , meaning that the singular quadrature weights and nodes are also identical even as the mapping \mathbf{R}^k and associated Jacobian J^k change. Applying the singular quadrature rule (3.26) with

$\rho = K_{nm} J^k$ and $\widehat{\mathcal{C}} = \widehat{\mathcal{T}}_0$ results in a set $\{(\chi_i, \omega_i) : i = 1, \dots, N^Q\}$, $N^Q = PM_{\partial S}m_P$, of pairs of quadrature nodes $\chi_i = (\xi_i, \eta_i)$ and weights ω_i pairs adequate for discretization of $\widehat{\mathcal{V}}_k$ that are independent of k (here, after appropriately-relabeling the index i in (3.26) to ℓ , the weights ω_i are given by $\omega_i = t_\ell v_\ell((\xi_j - \xi_*) \times \tau_j) w_j$ for appropriate j and ℓ). Re-writing (3.26) with this notation we have the singular corrections given by the ω_j -weighted inner product of Green function and Jacobian values with Koornwinder polynomial values,

$$\widehat{\mathcal{V}}_k[K_{nm}](\xi_0) \approx \sum_{i=1}^{N^Q} G(\mathbf{R}^k(\chi_i), \mathbf{R}^k(\xi_0)) J^k(\chi_i) K_{nm}(\chi_i) \omega_i, \quad (4.7)$$

from which it is clear that the only quantities that need to be recomputed for each cell are the mapped Green function values $G(\mathbf{R}^k(\chi_i), \mathbf{R}^k(\xi_0)) J^k(\chi_i)$; note, further, that J^k is constant for straight triangles. It is essential in the independence of the node-weight pairs (χ_i, ω_i) with respect to the mapping k (and thereby to the cell) that for a given target ξ_0 the singular quadrature rule developed in Section 3.3 is determined solely by the small-argument asymptotic behavior of the kernel function.

The additional challenge that unstructured meshes pose for near-singular corrections as opposed to the singular corrections discussed previously is that near-singular target locations $\mathbf{r}_0 = \mathbf{R}^j(\xi_0^j)$ (where $\xi_0^j = \xi_0$ is, say, one of the Koornwinder interpolation nodes in cell $\widehat{\mathcal{T}}_j$ —which are always in fixed locations in the $\widehat{\mathcal{T}}_j$ -reference space) unfortunately do not lay at arbitrary reference space locations ξ_0^k relative to the source cell \mathcal{T}_k ($k \neq j$) under consideration. For this reason it is no longer possible to repeatedly use the same fixed set of reference-space target points to directly generate the corrections as for the $j=k$ case described previously; fortunately, however, similar ideas are still applicable with essentially the same effect. Despite the arbitrary (reference-space) location of near-singular target points, the quadrature rule weights and nodes (and hence the required values of the polynomials K_{nm}) are still determined entirely by a combination of the resulting star-point ξ_* (which is itself, in turn, determined via the first formula in (3.16)) and the distance d of the target point to the boundary given by (3.24)—see Fig. 5 for a depiction of this dependence. As a result, one can re-use the quadratures generated for a fixed set of star-points laying along the boundary $\partial\widehat{\mathcal{T}}_0$, selecting the star-point ξ_* that lays closest to the solution of (3.16). Utilizing a fixed-size (independent of the number of cells) list of quadrature nodes, weights, and Koornwinder values K_{nm} , the quadrature rules for near-singular targets can be efficiently computed by means of a simple lookup table and, again, the formula (4.7) (where, of course, the mapped Green function still needs to be re-evaluated, as before). The already-modest storage costs of these quadrature rules in such a scheme could be further limited by storing rules corresponding to star points on only one side of $\partial\widehat{\mathcal{T}}_0$.

4.2. Boxes and tensor-product Chebyshev systems: smooth quadratures and singular corrections

In this section, we describe smooth quadrature methods and efficient singular and near-singular corrections of these quadratures for source boxes $\mathcal{C}_{k+N_t} = \mathcal{B}_k$, $k = 1, \dots, N_b$ (see Equation (3.2)). Every box \mathcal{B}_k is mapped from a single reference box $\widehat{\mathcal{B}}_0 = [-1, 1]^2$, with mappings in the uniform grid context of this article taking the simple form

$$\mathbf{R}^k(\xi) = \mathbf{o}_k + h\xi, \quad \text{where } \xi = (\xi, \eta)^T, \quad \text{with } \xi, \eta \in [-1, 1]. \quad (4.8)$$

Smooth quadrature and interpolation. As is standard for such regular regions, we represent the local source distribution f using the truncated tensor-product Chebyshev series expansion

$$f(\mathbf{r}) = f(\mathbf{R}^k(\xi)) \approx \sum_{n=0}^{p-1} \sum_{m=0}^{p-1} f_{nm}^k T_{nm}(\xi, \eta), \quad \text{for } \mathbf{r} = \mathbf{R}^k(\xi) \in \mathcal{B}_k, \quad (4.9)$$

where T_{nm} is defined by the tensor-product Chebyshev polynomial

$$T_{nm}(\xi) = T_n(\xi) T_m(\eta), \quad (4.10)$$

with $T_n(t)$ denoting the Chebyshev polynomial of degree n defined for $-1 \leq t \leq 1$. We denote by $I_{\mathcal{B},p}$ the set of interpolation nodes for the truncated series representation (4.9), which are defined as tensor-product Chebyshev nodes, i.e. roots of the Chebyshev polynomial T_p . It is well-known that for smooth f this is a p -th order accurate expansion with series coefficients f_{nm}^k decaying rapidly as either of n, m increase, and, further, that the associated Clenshaw-Curtis smooth quadrature rule is suitable for targets \mathbf{r}_0 well-separated from \mathcal{B}_k ,

$$\int_{\mathcal{B}_k} G(\mathbf{r}, \mathbf{r}_0) f(\mathbf{r}) d\mathbf{r} \approx \sum_{l=1}^{p^2} G(\mathbf{r}_l, \mathbf{r}_0) f(\mathbf{r}_l) w_l,$$

and, further, that is compatible with fast summation methods. Mirroring the upsampling procedure described in Section 3.2.1 the method uses a smooth quadrature of order q , $q \geq p$, for an interpolant (4.9) of order p so that source

approximation error is dominant, with the Chebyshev expansion (4.9) with associated interpolation nodes $I_{\mathcal{B},p}$ upsampled to the interpolation nodes $I_{\mathcal{B},q}$.

Singular and near-singular corrections. For singular and near-singular evaluation, similar to the case for triangular regions, we construct a linear map from function values f_{nm}^k in the truncated expansion (4.9) to corrected quadratures $\mathcal{V}_k(\mathbf{r}_0)$. The singular corrections are generated using the same Poincare-based singular integration technique used for triangular regions, applied to each element of the smooth basis in (4.9). Similarly to (4.5) we find

$$\mathcal{V}_k[f](\mathbf{r}_0) \approx \sum_{n=0}^{p-1} \sum_{m=0}^{p-1} f_{nm}^k \widehat{\mathcal{V}}_k[T_{nm}](\xi_0^k), \quad (4.11)$$

where $\widehat{\mathcal{V}}_k[\rho]$ is defined by Equation (3.3) and where from the mapping (4.8) we have $\xi_0^k = (\mathbf{r}_0 - \mathbf{o}_k)/h$. Crucially, it then follows from translational invariance $G(\mathbf{r}, \mathbf{r}_0) = G(\mathbf{r} - \mathbf{r}_0)$ that

$$\widehat{\mathcal{V}}_k[T_{nm}](\xi_0^k) = \int_{\mathcal{B}_0} G(h(\xi - \xi_0^k)) T_{nm}(\xi) h^2 d\xi, \quad (4.12)$$

where we emphasize that the expression for $\widehat{\mathcal{V}}_k$ is independent of k .

It follows from (4.12) and the fact that the reference-space interpolation and smooth quadrature nodes $I_{\mathcal{B},p}$ for \mathcal{B}_k are identical for all boxes, that the singular corrections can be computed once and re-used for all boxes as a lookup table. A similar argument leads to lookup tables for near-singular target points that arise at interpolation nodes of nearby boxes: since the grid is structured these need only be computed once (this remains true even with adaptivity, merely the size of the still-finite look-up table expands). Finally, for near-singular targets laying in triangles that lay in close proximity to boxes, corrections can be easily generated with the added efficiency (since, in this context the map for every box is the same) that the values of the Newton potential in the near-field can be re-used across boxes and indeed interpolation to machine precision of the volume potential is possible for generating near-singular corrections.

5. Numerical results

This section presents demonstrations of the character of the proposed numerical methods, applied to the Poisson and modified Helmholtz (Yukawa) equations. Together with the Stokes equations (see Fig. 1 for a solution of this equation with the proposed methodology) these equations represent the major classes of constant-coefficient elliptic PDEs to which our methods are applicable, even though, as noted in Section 6, our methods possess broader applicability. We first show the entirely routine and expected convergence of the associated homogeneous problem, then consider test cases that involve use of the volume potential, demonstrating both convergence and the asymptotic costs of the method with the use of the FMM.

As a preliminary test, we demonstrate the expected convergence of the numerical solver for the integral equations arising in the augmented boundary value problem Equation (2.3). The boundary integral equations are discretized using spectral Kress quadratures, the resulting linear systems are solved using GMRES with a relative residual tolerance of 10^{-15} , and, for evaluation of the layer potentials spectrally-accurate schemes [6] are used that yield high accuracy in the numerical solution arbitrarily close to the boundary. Fig. 7 validates the expected exponential convergence of the numerical solution as the total number of quadrature nodes increases. In the remainder of the numerical results, a sufficient number of boundary integral quadrature nodes are selected so that the error from the homogeneous component of the solver does not dominate.

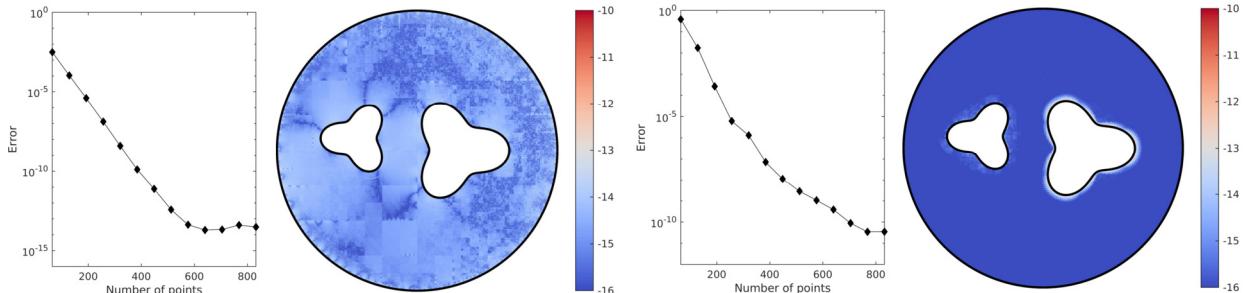


Fig. 7. Laplace and modified Helmholtz Dirichlet boundary value problem (BVP) test problems, with solution (and hence Dirichlet boundary values) given by point sources (for Helmholtz, $\lambda = 10$) at the center of each inclusion. Left (resp. Right): convergence plot of the collocation scheme for the Laplace (resp. Helmholtz) BVP with respect to the total number N of collocation points; error plots correspond to $N = 832$.

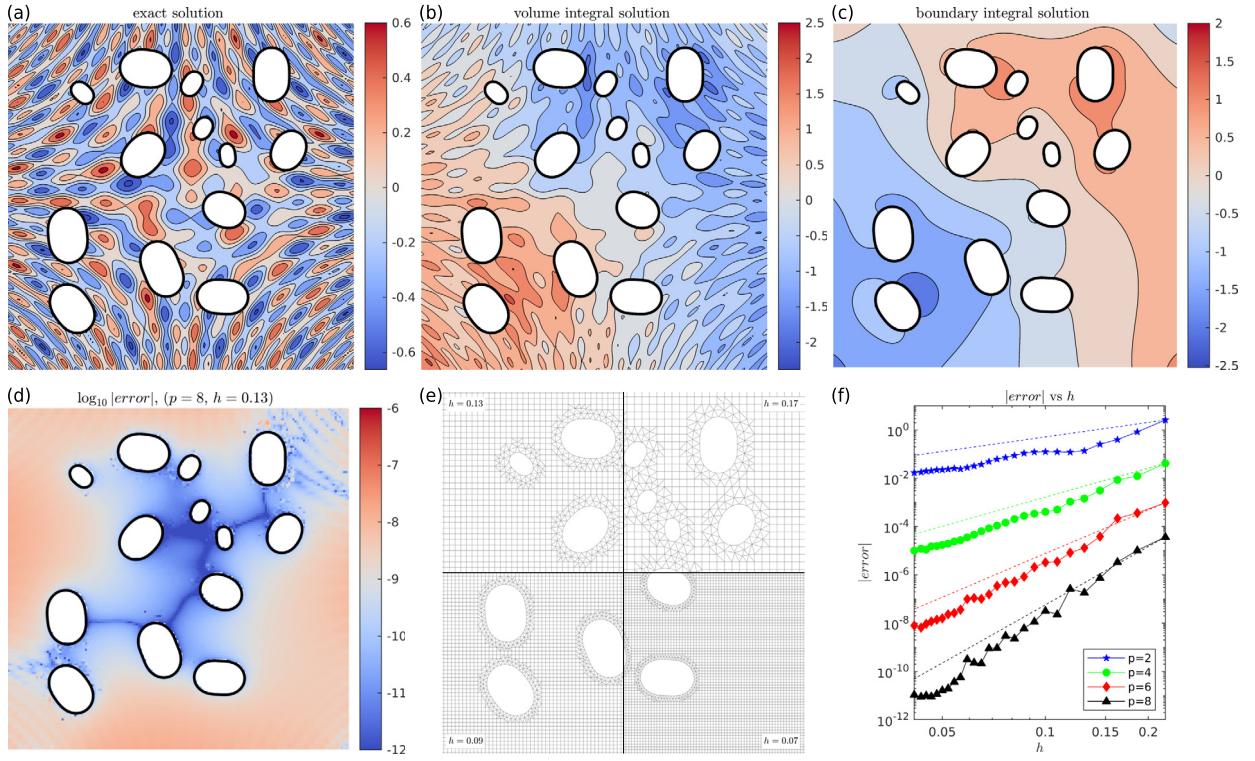


Fig. 8. Poisson problem simulation; solution and errors displayed in the subregion $[-4.2, 4.2]^2$ of a polydisperse domain. The three top subplots (a-c) depict, respectively, the contour lines for the associated solution u to Equation (2.1), the particular solution u_p arising from Equation (1.1), and the solution u_H to the boundary value problem (2.3). Subplot (d) is a typical plot of the base-10 logarithm of the error at points in a uniform evaluation grid, here corresponding to the solver run with $p = 8$ and $h = 0.13$. Subplot (e) displays four computational meshes at various levels of h -refinement. Subplot (f) demonstrates the convergence (with respect to gridsize h) of the numerical solution produced by the order $p = 2, 4, 6$ and 8 versions of the scheme; errors are measured on the same 100×100 uniform target grid over $[-6.5, 6.5]^2$ in all simulations, and dashed lines depict the expected order of convergence of $\mathcal{O}(h^p)$.

Table 2

Total number of triangular cells N_t and regular boxes N_b at various levels of h -refinement ($h = 13/N$) for the polydisperse test problem depicted in Fig. 8: the interior of a circle of radius 6.5 units and exterior to the elliptical inclusions. The last two rows demonstrate the linear and quadratic growth of N_t and N_b , respectively, as $h \rightarrow 0$; the ratios D_1 and D_2 are normalization constants $D_1 = 0.1857 \cdot 2329$ and $D_2 = 0.1857^2 \cdot 2476$.

N	70	100	130	160	190	220	250	280	310
h	0.1857	0.1300	0.1000	0.0813	0.0684	0.0591	0.0520	0.0464	0.0419
N_t	2329	3293	4230	5205	5698	6643	8049	8531	8965
N_b	2476	5774	10408	16358	23618	32227	42162	53403	65932
hN_t/D_1	—	0.99	0.98	0.98	0.90	0.90	0.97	0.92	0.87
$h^2 N_b/D_2$	—	1.14	1.22	1.27	1.29	1.32	1.34	1.35	1.36

Remark 7. Several Poisson examples in this section concern the inhomogeneity (the inhomogeneity being modified in an obvious manner for Helmholtz problems) given by

$$f(x, y) = 6 \sin(6x) + 8 \cos(8(y + \frac{1}{10})) + 4(x^2 + y^2) \sin(4xy) + 3 \cos(3x) \sin(3y) \quad (5.1)$$

with associated solution

$$u(x, y) = \frac{1}{6} \sin(6x) + \frac{1}{8} \cos(8(y + \frac{1}{10})) + \frac{1}{4} \sin(4xy) + \frac{1}{6} \cos(3x) \sin(3y). \quad (5.2)$$

We emphasize that while $f(x, y)$ is defined and is smooth for all (x, y) , only function values at interior points $(x, y) \in \Omega$ are used in the solution process.

We consider a Dirichlet problem for the Poisson equation in the presence of a polydisperse system Ω of inclusions, in the interior of a circle of radius 6.5 units. The solution and its numerical components are displayed in the region $[-4.2, 4.2]^2$ in Fig. 8 (the solution in all of $x^2 + y^2 < 6.5^2$ is not displayed, but the error and therefore the convergence plot in panel (f)

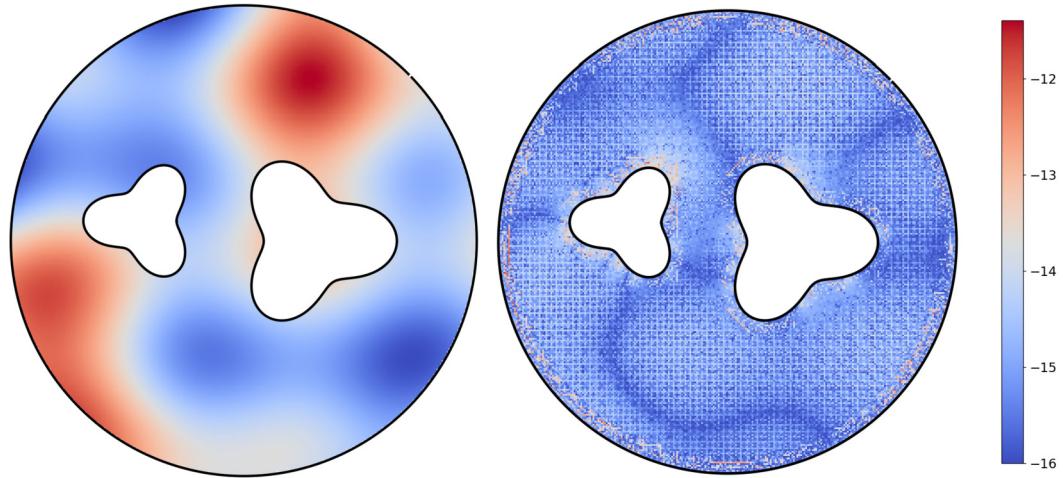


Fig. 9. Modified Helmholtz ($\lambda^2 = 100$) example solution with maximum value 0.56 (left) and error with color bar indicating the base-10 logarithm of the error (right); solver parameters used were $p = 10$, $h = 2.2/N$ with $N = 90$ (resulting in $N_b = 3633$ and $N_t = 1994$). The maximum error over approximately 51000 uniformly-spaced target points is $\|\mathbf{e}\|_{\ell^\infty} = 3.9 \cdot 10^{-12}$, while the discrete L^2 error is $\|\mathbf{e}\|_{\ell^2} = 6.9 \cdot 10^{-14}$.

is computed over the entire solution domain). The Poisson solution is tested on a 100×100 uniform grid of target points that lay inside Ω ; the maximum error is plotted in panel (f) of Fig. 8. This experimental setup can be seen as a somewhat more challenging test case for the volume potential scheme since it results in targets that can lay arbitrarily close to cell boundaries. Nevertheless we see convergence rates consistent with expectations from approximation theory down to a level on the order of 10^{-12} . In Fig. 8 panel (d) we observe that the error grows as x and/or y increase which is expected in view of the locations of sharper gradients of $f(x, y)$. Table 2 assembles relevant mesh data for this problem as the discretization is refined and demonstrates the expected linear (resp. quadratic) growth of the number N_t (resp. N_b) of boundary fitted cells (resp. regular boxes).

It is also of interest to solve the inhomogeneous modified Helmholtz equation (cf. Section 2 with $L = -\Delta + \lambda^2$), a problem with diverse applications, e.g. in wave scattering and elliptic time marching. The Green function of this operator shares the same singular kernel behavior as the classical Newton potential, and so the same singular and near-singular quadrature rules used there apply. In Fig. 9 we show the error in the proposed method applied to the modified Helmholtz equation with manufactured solution given in (5.2) over a region bounded by a circle of radius 1. The method presented is quite successful for low to moderate real values of λ with no modifications. However, both the accuracy of smooth quadrature rules and efficacy of the FMM are strained as λ is increased; a more complete study of the optimal relation between λ , the smooth quadrature rule and associated upsampling rate q , and the near-field selection region will be presented at a later date.

We turn next to performance demonstrations of the proposed method, which are summarized in Table 3. This experiment tests the solution of the Dirichlet Poisson problem in the region depicted in Fig. 7 (which features an enclosing circle of radius 3.5 units), with source function and associated solution given by Equation (5.1) and for volume target points at all volume interpolation nodes (i.e. for each cell C_k the appropriate (mapped) points in either of $I_{\mathcal{T},p}$ and $I_{\mathcal{B},p}$), and it allows us to make several observations and conclusions about the character of the solver. These experiments serve first to demonstrate that the costs associated with the sparse correction step of the method are, as desired, negligible in comparison to the FMM step of the method, which can be verified with the data in columns marked "%fmm". The timings clearly confirm the expected linear scaling of the method. The experiments also serve to demonstrate the costs and error quantities associated with producing the solution at all interpolation nodes in the domain decomposition scheme; the values of the solution at such points suffice to efficiently and accurately produce volume potential values at arbitrary targets $\mathbf{r}_0 \in \Omega$, but more crucially they can be easily seen as the values required for the solution of various nonlinear and time-dependent PDEs, by e.g. iteration or time-stepping. The error $\|\mathbf{e}_1\|_{\ell^\infty}$ in the second-to-last column of Table 3 is the maximum error in the solution u of the problem (2.1) evaluated across all ndof volume quadrature points and is consistent with the expected 5-th order convergence for the $p = 5$ scheme. Finally, the experiment demonstrates the expected linear (resp. quadratic) dependence of the number N_t (resp. N_b) of triangular regions (resp. boxes) on the mesh gridsize h , which for this experiment is given by $h = 9/N$. Results are similar for the scheme with other orders, and results are not included for the sake of brevity. All timings were produced using a Python implementation on a single core (i.e. with multithreading explicitly disabled) of an Exherbo Linux workstation with an AMD Ryzen-5 5600X CPU. A detailed presentation of performant generation of singular and near-singular corrections as detailed in Section 4.1.1 will be reported at a later date.

It is also of interest to obtain the solution at arbitrary points in the domain via interpolation, a straightforward task when (as in the context of the test of the previous paragraph) evaluation points are precisely the interpolation nodes $I_{\mathcal{T},p}$ and $I_{\mathcal{B},p}$. The error $\|\mathbf{e}_2\|_{\ell^2}$ in the solution thus obtained by Koornwinder and Chebyshev interpolation is listed in the last

Table 3

Timings (in seconds) for an experiment with $p = 5$, grid size $h = 9/N$; the number of degrees of freedom and smooth quadrature sources can be determined (using $q = 2p$) from Equations (4.1) and (4.2), respectively, using the columns labeled N_t and N_b . The quantities T_{Γ}^{fmm} and T_{Γ}^{corr} relate to evaluation of $\mathcal{V}[f]$ on Γ at boundary integral collocation nodes (cf. (2.3b)) while T_{Ω}^{fmm} and T_{Ω}^{corr} relate to evaluation of u at all ndofs source function nodes in the domain; ‘fmm’ superscripts refer to the application of the smooth quadrature rule and ‘corr’ superscripts refer to the application of the sparse correction map. The time T_{Γ}^{bie} associated with solution of the boundary integral equations using GMRES, evaluation of the solution at all volumetric targets, and application of close corrections to the layer potential evaluations consistently required $T_{\Gamma}^{\text{bie}} = 1.4$. The %fmm figures are determined by the ratios of T_{Γ}^{fmm} to T_{Γ}^{corr} and of T_{Ω}^{fmm} to T_{Ω}^{corr} , and a percentage of 99.9 is displayed for any percentage above 99.9.

N	N_t	N_b	T_{Γ}^{fmm}	T_{Γ}^{corr}	%fmm	T_{Ω}^{fmm}	T_{Ω}^{corr}	%fmm	$\ \mathbf{e}_1\ _{\ell^{\infty}}$	$\ \mathbf{e}_2\ _{\ell^2}$
40	828	286	1.79e-01	1.1e-04	99.9	4.30e-01	2.44e-03	99.4	8.2e-05	1.3e-03
60	1147	946	3.05e-01	1.1e-04	99.9	9.54e-01	5.07e-03	99.5	1.1e-05	7.1e-05
80	1440	1922	6.44e-01	1.4e-04	99.9	1.53e+00	7.61e-03	99.5	2.6e-06	1.3e-05
100	1858	3227	8.25e-01	1.1e-04	99.9	2.76e+00	1.11e-02	99.6	5.8e-07	3.4e-06
120	2394	4851	1.11e+00	1.2e-04	99.9	3.65e+00	1.57e-02	99.6	1.6e-07	6.5e-07
140	2653	6843	1.72e+00	1.2e-04	99.9	4.57e+00	2.25e-02	99.5	5.3e-08	1.9e-07

column of Table 3, where $\|\mathbf{e}\|_{\ell^2} = \left(\sum_j h^2(\mathbf{e})_j^2 \right)^{1/2}$ denotes the discrete L^2 norm of the error \mathbf{e} , and where $h = 0.045$ is the gridsize of the uniform evaluation grid (with $\approx 17,000$ points in Ω) onto which solution values are interpolated.

6. Conclusions & future directions

We have presented a volume quadrature scheme that is used to solve inhomogeneous PDEs (either exterior or interior problems) on irregular (either simply- or multiply-connected) domains, and which is to our knowledge the first scheme for complex geometry that is of optimal asymptotic complexity in addition to being high-order accurate (including near to the boundary), and which, further, is capable of producing solutions to a wide variety of inhomogeneous elliptic PDEs, is efficient with respect to the number of degrees of freedom, is amenable to rapid generation of high-order accurate quadrature rules for singular and near-singular corrections, and is readily generalizable to non-PDE kernels and to three dimensional domains. It is worth mentioning that the present work solves inhomogeneous PDEs by direct evaluation of the true volume potential over Ω , in contrast to the extension-based methods that constitute most prior work. In certain circumstances this distinction is relevant because the true volume potential is desired in and of itself and extension-based methods are no longer helpful; we mention for example ongoing work in volume potentials that arise in fractional PDEs and in the physical problem of quantifying fluid mixing, as well as the potential-theoretic solution, via the Lippmann-Schwinger equation, of the problem of wave scattering by media with a spatially-variable (and potentially discontinuous) index of refractivity.

A few direct extensions are currently being pursued, among them adaptivity in the bulk region and coupling to meshing technologies such as `TriWild` which could enable further adaptivity in the region near to the boundary. Use of the methodology with non-PDE kernels is currently being investigated, as is application to nonlinear PDEs. In the context of PDEs with moving geometries, a principal concern is the rapid generation of new quadratures for cells of the irregular section of the grid which naturally changes; the approach to efficient pre-computation of singular corrections we described should ameliorate this cost in that context. Finally, this work presents a clear path to tackling the 3D volume potential problem since all elements of the proposed methodology generalize immediately to that setting, with quadratures and orthogonal polynomial systems on (mapped curvilinear) triangles translating to (mapped curvilinear) tetrahedrons.

Declaration of competing interest

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests: Shravan Veerapaneni reports financial support was provided by National Science Foundation.

Acknowledgements

We acknowledge support from NSF under grant DMS-2012424, and the Mcubed program at the University of Michigan (UM). Authors also acknowledge the computational resources and services provided by UM’s Advanced Research Computing. The authors are grateful to Zydrunas Gimbutas for helpful discussions.

Appendix A. Proofs of supporting lemmas

Proof of Lemma 1. Because \mathcal{S} is star-shaped with respect to the origin the function $\rho(\mathbf{tr})$ is defined for all $\mathbf{r} \in \partial\mathcal{S}$ and for all $t \in [0, 1]$, and as a consequence the right-hand integral in (3.14) is well-defined. Denoting $\mathbf{r} = (x, y)$ and, for $\mathbf{r} \in \partial\mathcal{S}$, $\boldsymbol{\tau} = (dx/ds, dy/ds)$ (which is well-defined at almost every $\mathbf{r} \in \partial\mathcal{S}$ since $\partial\mathcal{S}$ is piecewise smooth), we first have

$$\oint_{\partial\mathcal{S}} \left(\int_0^1 t \rho(\mathbf{tr}) dt \right) \mathbf{r} \times \boldsymbol{\tau} ds = \oint_{\partial\mathcal{S}} \left(\int_0^1 t \rho(\mathbf{tr}) dt \right) (x dy - y dx).$$

But then, since \mathcal{S} is a region with a piecewise-smooth boundary and the function $\int_0^1 t \rho(t\mathbf{r}) dt$ is continuously differentiable with respect to each of the variables x and y , we can apply Green's theorem [78, Thm. 10–43] to the vector field in the circulation integral to obtain

$$\oint_{\partial\mathcal{S}} \left(\int_0^1 t \rho(t\mathbf{r}) dt \right) \mathbf{r} \times \boldsymbol{\tau} \, ds = \int_{\mathcal{S}} \left(\frac{\partial}{\partial x} \left(x \int_0^1 t \rho(t\mathbf{r}) dt \right) + \frac{\partial}{\partial y} \left(y \int_0^1 t \rho(t\mathbf{r}) dt \right) \right) \, dA.$$

It is then a straightforward computation to verify that

$$\oint_{\partial\mathcal{S}} \left(\int_0^1 t \rho(t\mathbf{r}) dt \right) \mathbf{r} \times \boldsymbol{\tau} \, ds = \int_{\mathcal{S}} \left(\int_0^1 \left(\frac{d}{dt} t^2 \rho(t\mathbf{r}) \right) dt \right) \, dA = \int_{\mathcal{S}} \rho(\mathbf{r}) \, dA,$$

which completes the proof. \square

Proof of Corollary 1. The proof proceeds by isolating a ball containing the location of the singularity, showing the contribution due to this region is negligible, and applying Lemma 1 on a region excluding this ball. In order to proceed and utilize Lemma 1 we first change to a coordinate system centered at \mathbf{r}_* via the map $\mathbf{r} \mapsto \mathbf{r} + \mathbf{r}_*$, and write

$$\int_{\mathcal{S}} K(\mathbf{r} - \mathbf{r}_*) \rho(\mathbf{r}) \, dA = \int_{\mathcal{S}_*} K(\mathbf{r}) \rho(\mathbf{r} + \mathbf{r}_*) \, dA, \quad (\text{A.1})$$

where \mathcal{S}_* is an \mathbf{r}_* -translate of \mathcal{S} and, in particular, is star-shaped with respect to the origin.

It is useful to introduce a smooth cut-off function $\chi : [0, \infty) \rightarrow [0, 1]$, χ non-increasing and satisfying both $\chi(v) = 1$ in a neighborhood of $v = 0$ and $\chi(v) = 0$ for $v \geq 1$. In view of (A.1), using this function we have the identity, for every $\varepsilon > 0$,

$$\begin{aligned} \int_{\mathcal{S}} K(\mathbf{r} - \mathbf{r}_*) \rho(\mathbf{r}) \, dA &= \int_{\mathcal{S}_*} K(\mathbf{r}) (1 - \chi(|\mathbf{r}|/\varepsilon)) \rho(\mathbf{r} + \mathbf{r}_*) \, dA \\ &\quad + \int_{\mathcal{S}_*} K(\mathbf{r}) \chi(|\mathbf{r}|/\varepsilon) \rho(\mathbf{r} + \mathbf{r}_*) \, dA. \end{aligned} \quad (\text{A.2})$$

It is easy to see that since for sufficiently small $\varepsilon > 0$ the quantity $(1 - \chi(|\mathbf{r}|/\varepsilon))$ vanishes for \mathbf{r} in an ε -dependent neighborhood of the origin, the integrand of the first integral on the right-hand side of (A.2) is a smooth function of \mathbf{r} , and so by applying Lemma 1 to that integral we have

$$\begin{aligned} \int_{\mathcal{S}_*} K(\mathbf{r}) \rho(\mathbf{r} + \mathbf{r}_*) \, dA &= \oint_{\partial\mathcal{S}_*} \left(\int_0^1 t K(t\mathbf{r}) (1 - \chi(t|\mathbf{r}|/\varepsilon)) \rho(t\mathbf{r} + \mathbf{r}_*) \, dt \right) \mathbf{r} \times \boldsymbol{\tau} \, ds \\ &\quad + \int_{\mathcal{S}_*} K(\mathbf{r}) \chi(|\mathbf{r}|/\varepsilon) \rho(\mathbf{r} + \mathbf{r}_*) \, dA. \end{aligned} \quad (\text{A.3})$$

Denoting by $I_\varepsilon(\mathbf{r})$ the integrand of the outer integral above,

$$I_\varepsilon(\mathbf{r}) = (\mathbf{r} \times \boldsymbol{\tau}) \int_0^1 t K(t\mathbf{r}) (1 - \chi(t|\mathbf{r}|/\varepsilon)) \rho(t\mathbf{r} + \mathbf{r}_*) \, dt,$$

defining $\hat{\mathbf{r}}$ via $\mathbf{r} = |\mathbf{r}| \hat{\mathbf{r}}$, and estimating using the triangle inequality and the fact that $0 \leq \chi(v) \leq 1$ we find that the bound

$$|I_\varepsilon(\mathbf{r})| \leq |\mathbf{r}| |\hat{\mathbf{r}} \times \boldsymbol{\tau}| \int_0^1 t |K(t\mathbf{r})| \rho(t\mathbf{r} + \mathbf{r}_*) \, dt \leq C |\mathbf{r}| \int_0^1 t |K(t\mathbf{r})| \, dt$$

holds for every $\varepsilon > 0$ and for all $\mathbf{r} \in \mathcal{S}_*$, where $C = C(\mathcal{S}_*, \rho) > 0$ is a constant. Since K is weakly-singular, we further have that there exists a constant $\tilde{C} = \tilde{C}(\mathcal{S}_*, \rho, K) > 0$ such that, for all $\mathbf{r} \in \mathcal{S}_*$,

$$|I_\varepsilon(\mathbf{r})| \leq \tilde{C} |\mathbf{r}|^{1-\alpha} \int_0^1 t^{1-\alpha} \, dt, \quad (\text{A.4})$$

for some $\alpha < 2$, from which we conclude that $I_\varepsilon(\mathbf{r})$ is integrable over \mathcal{S}_* . Now, since firstly the upper bound for $I_\varepsilon(\mathbf{r})$ is integrable and since the integrand $I_\varepsilon(\mathbf{r})$ converges pointwise, and secondly the inner integral in $I_\varepsilon(\mathbf{r})$ also converges pointwise and is bounded above a function that is also integrable over the interval $[0, 1]$ (namely, by the integrand in the upper bound (A.4)), by two applications of the dominated convergence theorem in the $\varepsilon \rightarrow 0$ limit for the first integrand in the right-hand side of (A.3) we find

$$\lim_{\varepsilon \rightarrow 0} \int_{\mathcal{S}_*} K(\mathbf{r}) (1 - \chi(|\mathbf{r}|/\varepsilon)) \rho(\mathbf{r} + \mathbf{r}_*) dA = \oint_{\partial\mathcal{S}_*} \left(\int_0^1 t K(t\mathbf{r}) \rho(t\mathbf{r} + \mathbf{r}_*) dt \right) (\mathbf{r} \times \boldsymbol{\tau}) ds. \quad (\text{A.5})$$

But since K is a weakly-singular kernel function, i.e. it satisfies for some constant $D > 0$ the bound $|K(\mathbf{r})| \leq D|\mathbf{r}|^{-\alpha}$, $\alpha < 2$, we have that

$$\lim_{\varepsilon \rightarrow 0} \int_{B_\varepsilon(\mathbf{0})} |K(\mathbf{r})| dA = 0, \quad (\text{A.6})$$

so that by taking the limit as $\varepsilon \rightarrow 0$ in the right-hand side of (A.2) and using both (A.5) and (A.6), we have

$$\int_{\mathcal{S}} K(\mathbf{r} - \mathbf{r}_*) \rho(\mathbf{r}) dA = \oint_{\partial\mathcal{S}_*} \left(\int_0^1 t K(t\mathbf{r}) \rho(t\mathbf{r} + \mathbf{r}_*) dt \right) \mathbf{r} \times \boldsymbol{\tau} ds. \quad (\text{A.7})$$

The desired result follows by a change of coordinates $\mathbf{r} \mapsto \mathbf{r} - \mathbf{r}_*$. \square

References

- [1] G.C. Hsiao, W.L. Wendland, *Boundary Integral Equations*, Springer, 2008.
- [2] L. Greengard, V. Rokhlin, A fast algorithm for particle simulations, *J. Comput. Phys.* 73 (2) (1987) 325–348.
- [3] S. Hao, A.H. Barnett, P.-G. Martinsson, P. Young, High-order accurate methods for Nyström discretization of integral equations on smooth curves in the plane, *Adv. Comput. Math.* 40 (1) (2014) 245–272.
- [4] J. Helsing, R. Ojala, On the evaluation of layer potentials close to their sources, *J. Comput. Phys.* 227 (5) (2008) 2899–2921.
- [5] A. Klöckner, A. Barnett, L. Greengard, M. O’Neil, Quadrature by expansion: a new method for the evaluation of layer potentials, *J. Comput. Phys.* 252 (2013) 332–349.
- [6] A. Barnett, B. Wu, S. Veerapaneni, Spectrally accurate quadratures for evaluation of layer potentials close to the boundary for the 2D Stokes and Laplace equations, *SIAM J. Sci. Comput.* 37 (4) (2015) B519–B542.
- [7] A. Rahimian, A. Barnett, D. Zorin, Ubiquitous evaluation of layer potentials using quadrature by kernel-independent expansion, *BIT Numer. Math.* 58 (2) (2018) 423–456.
- [8] L. af Klinteberg, A.-K. Tornberg, Adaptive quadrature by expansion for layer potential evaluation in two dimensions, *SIAM J. Sci. Comput.* 40 (3) (2018) A1225–A1249.
- [9] B. Wu, H. Zhu, A. Barnett, S. Veerapaneni, Solution of Stokes flow in complex nonsmooth 2D geometries via a linear-scaling high-order adaptive integral equation scheme, *J. Comput. Phys.* 410 (2020) 109361.
- [10] C. Pérez-Arancibia, L.M. Faria, C. Turc, Harmonic density interpolation methods for high-order evaluation of Laplace layer potentials in 2D and 3D, *J. Comput. Phys.* 376 (2019) 411–434.
- [11] L.M. Faria, C. Pérez-Arancibia, M. Bonnet, General-purpose kernel regularization of boundary integral equations via density interpolation, *Comput. Methods Appl. Mech. Eng.* 378 (2021) 113703.
- [12] F. Ethridge, L. Greengard, A new fast-multipole accelerated Poisson solver in two dimensions, *SIAM J. Sci. Comput.* 23 (3) (2001) 741–760.
- [13] H. Langston, L. Greengard, D. Zorin, A free-space adaptive FMM-based PDE solver in three dimensions, *Commun. Appl. Math. Comput. Sci.* 6 (1) (2011) 79–122.
- [14] D. Malhotra, A. Gholami, G. Biros, A volume integral equation Stokes solver for problems with variable coefficients, in: *SC’14: Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, IEEE, 2014, pp. 92–102.
- [15] D. Malhotra, G. Biros, PVFMM: a parallel kernel independent FMM for particle and volume potentials, *Commun. Comput. Phys.* 18 (3) (2015) 808–830.
- [16] L. Greengard, J.-Y. Lee, A direct adaptive Poisson solver of arbitrary order accuracy, *J. Comput. Phys.* 125 (2) (1996) 415–424.
- [17] M. Berger, A. Giuliani, A state redistribution algorithm for finite volume schemes on cut cell meshes, *J. Comput. Phys.* 428 (2021) 109820.
- [18] M.H. Langston, An Adaptive Fast Multipole Method-Based PDE Solver in Three Dimensions, PhD thesis, New York University, 2012.
- [19] A. Mayo, The fast solution of Poisson’s and the biharmonic equations on irregular regions, *SIAM J. Numer. Anal.* 21 (2) (1984) 285–299.
- [20] A. Mayo, The rapid evaluation of volume integrals of potential theory on general regions, *J. Comput. Phys.* 100 (2) (1992) 236–245.
- [21] A. McKenney, L. Greengard, A. Mayo, A fast Poisson solver for complex geometries, *J. Comput. Phys.* 118 (2) (1995) 348–355.
- [22] N.R. Rapaka, R. Samtaney, An efficient Poisson solver for complex embedded boundary domains using the multi-grid and fast multipole methods, *J. Comput. Phys.* 410 (2020) 109387.
- [23] G. Biros, L. Ying, D. Zorin, A fast solver for the Stokes equations with distributed forces in complex geometries, *J. Comput. Phys.* 193 (1) (2004) 317–348.
- [24] T. Askham, A.J. Cerfon, An adaptive fast multipole accelerated Poisson solver for complex geometries, *J. Comput. Phys.* 344 (2017) 1–22.
- [25] O.P. Bruno, J. Paul, Two-dimensional Fourier continuation and applications, *SIAM J. Sci. Comput.* 44 (2) (2022) A964–A992.
- [26] F. Fryklund, E. Lehto, A.-K. Tornberg, Partition of unity extension of functions on complex domains, *J. Comput. Phys.* 375 (2018) 57–79.
- [27] L. af Klinteberg, T. Askham, M.C. Kropinski, A fast integral equation method for the two-dimensional Navier-Stokes equations, *J. Comput. Phys.* 409 (2020) 109353.
- [28] F. Fryklund, M.C.A. Kropinski, A.-K. Tornberg, An integral equation-based numerical method for the forced heat equation on complex domains, *Adv. Comput. Math.* 46 (5) (2020) 1–36.
- [29] G. Of, O. Steinbach, P. Uthaler, Fast evaluation of volume potentials in boundary element methods, *SIAM J. Sci. Comput.* 32 (2) (2010) 585–602.
- [30] B. Vioreanu, V. Rokhlin, Spectra of multiplication operators as a numerical tool, *SIAM J. Sci. Comput.* 36 (1) (2014) A267–A288.

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

[32] S. Mousavi, H. Xiao, N. Sukumar, Generalized Gaussian quadrature rules on arbitrary polygons, *Int. J. Numer. Methods Eng.* 82 (1) (2010) 99–113.

[33] D. Dunavant, High degree efficient symmetrical Gaussian quadrature rules for the triangle, *Int. J. Numer. Methods Eng.* 21 (6) (1985) 1129–1148.

[34] G. Cowper, Gaussian quadrature formulas for triangles, *Int. J. Numer. Methods Eng.* 7 (3) (1973) 405–408.

[35] R. Cools, P. Rabinowitz, Monomial cubature rules since “Stroud”: a compilation, *J. Comput. Appl. Math.* 48 (3) (1993) 309–326.

[36] J.N. Lyness, R. Cools, A Survey of Numerical Cubature over Triangles, *Proceedings of Symposia in Applied Mathematics*, vol. 48, American Mathematical Society, 1994, pp. 127–150.

[37] A.H. Stroud, *Approximate Calculation of Multiple Integrals*, Prentice-Hall, Inc., Englewood Cliffs, NJ, 1971.

[38] R. Cools, Constructing cubature formulae: the science behind the art, *Acta Numer.* 6 (1997) 1–54.

[39] J. Strain, Locally corrected multidimensional quadrature rules for singular functions, *SIAM J. Sci. Comput.* 16 (4) (1995) 992–1017.

[40] J.C. Aguilar, Y. Chen, High-order corrected trapezoidal quadrature rules for functions with a logarithmic singularity in 2-D, *Comput. Math. Appl.* 44 (8–9) (2002) 1031–1039.

[41] J. Bremer, Z. Gimbutas, A Nyström method for weakly singular integral operators on surfaces, *J. Comput. Phys.* 231 (14) (2012) 4885–4903.

[42] J. Bremer, Z. Gimbutas, On the numerical evaluation of the singular integrals of scattering theory, *J. Comput. Phys.* 251 (2013) 327–343.

[43] A. Sommariva, M. Vianello, Product Gauss cubature over polygons based on Green’s integration formula, *BIT Numer. Math.* 47 (2) (2007) 441–453.

[44] Y. Sudhakar, J.M. De Almeida, W.A. Wall, An accurate, robust, and easy-to-implement method for integration over arbitrary polyhedra: application to embedded interface methods, *J. Comput. Phys.* 273 (2014) 393–415.

[45] D. Gundersen, K. Weiss, J.A. Evans, Spectral mesh-free quadrature for planar regions bounded by rational parametric curves, *Comput. Aided Des.* 130 (2021) 102944.

[46] B. Müller, F. Kummer, M. Oberlack, Highly accurate surface and volume integration on implicit domains by means of moment-fitting, *Int. J. Numer. Methods Eng.* 96 (8) (2013) 512–528.

[47] R. Saye, High-order quadrature methods for implicitly defined surfaces and volumes in hyperrectangles, *SIAM J. Sci. Comput.* 37 (2) (2015) A993–A1019.

[48] A.H. Barnett, Efficient high-order accurate Fresnel diffraction via areal quadrature and the nonuniform fast Fourier transform, *J. Astron. Telesc. Instrum. Syst.* 7 (2) (2021) 021211.

[49] S.N. Fata, Treatment of domain integrals in boundary element methods, *Appl. Numer. Math.* 62 (6) (2012) 720–735.

[50] Q. Wang, W. Zhou, Y. Cheng, G. Ma, X. Chang, A line integration method for the treatment of 3D domain integrals and accelerated by the fast multipole method in the BEM, *Comput. Mech.* 59 (4) (2017) 611–624.

[51] Q. Wang, W. Zhou, Y. Cheng, G. Ma, X. Chang, Q. Huang, The boundary element method with a fast multipole accelerated integration technique for 3D elastostatic problems with arbitrary body forces, *J. Sci. Comput.* 71 (3) (2017) 1238–1264.

[52] P.W. Partridge, C.A. Brebbia, L.C. Wrobel, *Dual Reciprocity Boundary Element Method*, Springer Science & Business Media, 2012.

[53] X.-W. Gao, The radial integration method for evaluation of domain integrals with boundary-only discretization, *Eng. Anal. Bound. Elem.* 26 (10) (2002) 905–916.

[54] X.-W. Gao, Evaluation of regular and singular domain integrals with boundary-only discretization—theory and Fortran code, *J. Comput. Appl. Math.* 175 (2) (2005) 265–290.

[55] H. Zhu, Fast, High-Order Accurate Integral Equation Methods and Application to PDE-Constrained Optimization, PhD thesis, University of Michigan, 2021.

[56] H. Zhu, S. Veerapaneni, High-order close evaluation of Laplace layer potentials: a differential geometric approach, *SIAM J. Sci. Comput.* 44 (3) (2022) A1381–A1404.

[57] K.E. Atkinson, The numerical evaluation of particular solutions for Poisson’s equation, *IMA J. Numer. Anal.* 5 (3) (July 1985) 319–338.

[58] R. Kress, *Linear Integral Equations*, vol. 82, Springer, 1989.

[59] J. Sifuentes, Z. Gimbutas, L. Greengard, Randomized methods for rank-deficient linear systems, *Electron. Trans. Numer. Anal.* 44 (2015) 177–189.

[60] C. Geuzaine, J.-F. Remacle, Gmsh: a 3-D finite element mesh generator with built-in pre- and post-processing facilities, *Int. J. Numer. Methods Eng.* 79 (11) (2009) 1309–1331.

[61] Y. Hu, T. Schneider, X. Gao, Q. Zhou, A. Jacobson, D. Zorin, D. Panozzo, TriWild: robust triangulation with curve constraints, *ACM Trans. Graph.* 38 (4) (2019) 1–15.

[62] L. Greengard, M. O’Neil, M. Rachh, F. Vico, Fast multipole methods for the evaluation of layer potentials with locally-corrected quadratures, *J. Comput. Phys.* 10 (2021) 100092.

[63] C.H. Edwards, *Advanced Calculus of Several Variables*, Courier Corporation, 2012.

[64] W.J. Gordon, C.A. Hall, Construction of curvilinear co-ordinate systems and applications to mesh generation, *Int. J. Numer. Methods Eng.* 7 (4) (1973) 461–477.

[65] W.J. Gordon, C.A. Hall, Transfinite element methods: blending-function interpolation over arbitrary curved element domains, *Numer. Math.* 21 (2) (1973) 109–129.

[66] B. Szabó, I. Babuška, *Finite Element Analysis*, John Wiley & Sons, 1991.

[67] M. Abramowitz, I.A. Stegun, *Handbook of Mathematical Functions with Formulas, Graphs, and Mathematical Tables*, vol. 55, US Government Printing Office, 1972.

[68] T. Koornwinder, Two-variable analogues of the classical orthogonal polynomials, in: *Theory and Application of Special Functions*, Elsevier, 1975, pp. 435–495.

[69] K.C. Chung, T.H. Yao, On lattices admitting unique Lagrange interpolations, *SIAM J. Numer. Anal.* 14 (4) (1977) 735–743.

[70] M. Spivak, *Calculus on Manifolds: A Modern Approach to Classical Theorems of Advanced Calculus*, CRC Press, 2018.

[71] J.C. Maxwell, On approximate multiple integration between limits of summation, *Proc. Camb. Philos. Soc.* 3 (1877) 39–47.

[72] B.K. Alpert, Hybrid Gauss-trapezoidal quadrature rules, *SIAM J. Sci. Comput.* 20 (5) (1999) 1551–1584.

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

[74] P. Kolm, V. Rokhlin, Numerical quadratures for singular and hypersingular integrals, *Comput. Math. Appl.* 41.3–4 (2001) 327–352.

[75] V. Rokhlin, End-point corrected trapezoidal quadrature rules for singular functions, *Comput. Math. Appl.* 20 (7) (1990) 51–62.

[76] O.P. Bruno, E. Garza, A Chebyshev-based rectangular-polar integral solver for scattering by geometries described by non-overlapping patches, *J. Comput. Phys.* 421 (2020) 109740.

[77] M. Wala, A. Klöckner, Optimization of fast algorithms for global quadrature by expansion using target-specific expansions, *J. Comput. Phys.* 403 (2020) 108976.

[78] T.M. Apostol, *Mathematical Analysis: A Modern Approach to Advanced Calculus*, Addison-Wesley Publishing Company, 1957.