Quadrature by Expansion: A New Method for the Evaluation of Layer Potentials

Andreas Klöckner[a], Alexander Barnett[b], Leslie Greengard[a], Michael O'Neil[a]

[a] *Courant Institute of Mathematical Sciences, 251 Mercer Street, New York, NY 10012*
[b] *Department of Mathematics, Dartmouth College, Hanover, NH, 03755*

Abstract

Integral equation methods for the solution of partial differential equations, when coupled with suitable fast algorithms, yield geometrically flexible, asymptotically optimal and well-conditioned schemes in either interior or exterior domains. The practical application of these methods, however, requires the accurate evaluation of boundary integrals with singular, weakly singular or nearly singular kernels. Historically, these issues have been handled either by low-order product integration rules (computed semi-analytically), by singularity subtraction/cancellation, by kernel regularization and asymptotic analysis, or by the construction of special purpose "generalized Gaussian quadrature" rules. In this paper, we present a systematic, high-order approach that works for any singularity (including hypersingular kernels), based only on the assumption that the field induced by the integral operator is locally smooth when restricted to either the interior or the exterior. Discontinuities in the field across the boundary are permitted. The scheme, denoted QBX (quadrature by expansion), is easy to implement and compatible with fast hierarchical algorithms such as the fast multipole method. We include accuracy tests for a variety of integral operators in two dimensions on smooth and corner domains.

Key words: Layer Potentials, Singular Integrals, Quadrature, High-order accuracy, Integral equations

1. Introduction

One of the difficulties encountered in the practical application of integral equation methods lies in the need to evaluate integrals with singular or weakly singular kernels in complicated domains. For the sake of concreteness, we assume the computational task is to compute layer potentials such as the single and double layer potentials

$$S\sigma(x) := \int_\Gamma G(x,x')\sigma(x')\,dx' \tag{1}$$

$$D\mu(x) := \int_\Gamma \frac{\partial G}{\partial \hat{n}_{x'}}(x,x')\mu(x')\,dx' \tag{2}$$

for target points x on a closed, smooth contour $\Gamma \subset \mathbb{R}^2$, where G is the Green's function for an underlying elliptic PDE and $\hat{n}_{x'}$ denotes the outward unit normal at x'. In the case of the double layer potential, it is typically the principal value of $D\mu$ that is desired for $x \in \Gamma$.

In the present paper, we will restrict our attention to the Helmholtz equation

$$\Delta\phi + k^2\phi = 0,$$

for which

$$G(x,x') = \frac{i}{4}H_0^{(1)}(k|x-x'|), \tag{3}$$

Email address: kloeckner@cims.nyu.edu (Andreas Klöckner)

where $H_0^{(1)}$ denotes the Hankel function of the first kind of order 0. $H_0^{(1)}$ satisfies the Sommerfeld radiation condition

$$\lim_{r\to\infty} r^{1/2}\left(\frac{\partial}{\partial r} - ik\right) H_0^{(1)} = 0,$$

where $r = |x-x'|$ and $k \in \mathbb{C}$ with $\operatorname{Im} k \geq 0$. Using this Green's function, it is well-known that the kernels in (1), (2) are both logarithmically singular as operators acting on the boundary, and the quantities of interest are well-defined improper integrals. Off the boundary, the double layer potential must be treated with more care as the singularity is of the order $1/|x-x'|$ and the limiting value has a jump of $\mu(x')$ at the point $x' \in \Gamma$. In many applications, one would also like to compute the gradient of $S\sigma$ or $D\mu$, which involves principal value and finite-part (hypersingular) integrals.

When the target x is far from the boundary, the integrands in (1), (2) are smooth, and high-order quadratures can be obtained by standard methods. Difficulties are encountered only when x is either on or near the boundary. The problem of quadrature for singular or nearly singular integrals, of course, has a rich history and, we do not seek to review the literature here (see, for example, the texts [3, 7, 37]). The most common approach is probably product integration, that is to say exact integration of the kernel multiplied by a piecewise polynomial approximation of the density σ or μ on a piecewise smooth approximation of the boundary. Purely analytic rules, however, tend to be limited to a few singularities (such as $\log|x-x'|$ in 2D or $1/|x-x'|$ in 3D) and low order approximations of the density and boundary. To handle the kernels $H_0^{(1)}(k|x-x'|)$ or $e^{ik|x-x'|}/|x-x'|$, analytic rules are often combined with numerical quadratures through the method of *singularity subtraction*. More precisely, it is easy to verify that

$$\frac{i}{4}H_0^{(1)}(k|x-x'|) - \frac{1}{2\pi}\log|x-x'| \quad \text{and} \quad \frac{e^{ik|x-x'|}}{|x-x'|} - \frac{1}{|x-x'|}$$

are smoother functions than the original kernels themselves and somewhat easier to integrate numerically. More generally, if the kernel $G_1(x,x')$ can be integrated, say, on a flat surface by analytic means, then integrating $G_2(x,x') - G_1(x,x')$ is an easier task if G_2 and G_1 have the same leading order singularity [14, 17, 28].

Three other powerful approaches are (a) to design special purpose quadratures that integrate a specific class of singular functions with high-order accuracy [2, 9, 24, 29, 33, 36, 44–46], (b) to find a change of variables that removes the principal singularity [10, 14, 15, 19, 22, 27, 30, 35, 43, 47], and (c) to regularize the kernel so that smooth rules can be applied, followed by corrections through asymptotic analysis or Richardson extrapolation [5, 18, 23, 39, 43]. By contrast with singularity subtraction, methods of type (b) are sometimes referred to as using *singularity cancellation*. In the complex analytic (or harmonic) case, we should also note that some remarkable interpolatory methods have been developed by Helsing and Ojala [25] for off-surface evaluation.

The main purpose of the present paper is to introduce a rather different approach to the evaluation of layer potentials, based on the fact that the fields $S\sigma$ or $D\mu$ in (1), (2) are locally smooth functions when restricted to either the interior or the exterior, although they may be discontinuous across the boundary. The scheme, denoted QBX (quadrature by expansion), is easy to implement, high order accurate, and requires only a smooth underlying quadrature scheme. This underlying smooth rule may be global or composite/panel-based and adaptive. The method is also compatible with fast hierarchical algorithms. QBX is an extension of the work of Barnett and Nguyen [4], who address the near but off-surface evaluation problem.

The paper is organized as follows: In Section 2, we show how QBX follows naturally from considerations of potential theory. Section 3 describes the mathematical foundations of the method, and Section 4 demonstrates its numerical performance. A simple but complete description of the algorithm can be found in Section 3.3. Finally, we discuss additional details and potential extensions of the present work in Section 5.

2. Smooth, high-order quadrature

Let us assume, for the moment, that we are given a smooth, simply connected closed curve $\Gamma \subset \mathbb{R}^2$, with a parametrization $\Gamma = \{\gamma(t) : 0 \leq t < L\}$. We denote the interior of Γ by Ω^- and its exterior by Ω^+. We also assume that we have at our disposal an underlying quadrature rule capable of integrating smooth

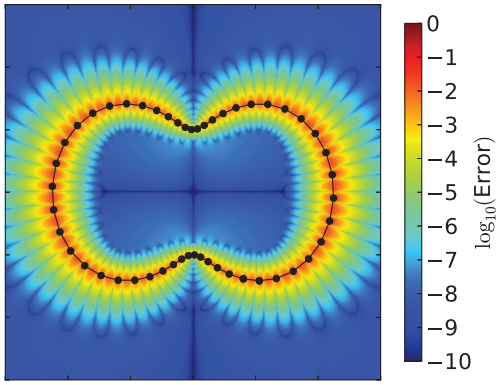

(a) Error in the potential using the trapezoidal rule with 50 quadrature points.

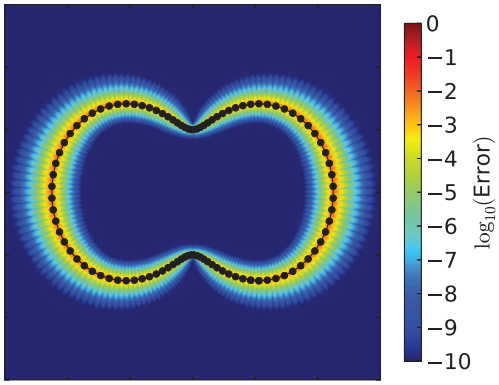

(b) Error in the potential using the trapezoidal rule with 100 quadrature points.

Figure 1. The potential $S\sigma$ is computed using the trapezoidal rule, a simple, high-order quadrature for smooth functions.

(non-singular) functions on Γ to high precision. In two dimensions, one option is the trapezoidal rule, since it is well-known to achieve superalgebraic convergence for smooth data on closed curves [14].

A very natural question at this point is the following: for a target location x *away* from Γ, how well does the trapezoidal rule compute $S\sigma(x)$ or $D\mu(x)$? Certainly, the integrands in (1) and (2) are not actually singular in this situation, so the real question is how close x can be to Γ before accuracy is lost. Before analyzing this error more carefully, let us carry out a simple computational experiment for the curve

$$\gamma(t) = \begin{pmatrix} \frac{3}{4}\cos(t-\pi/4)(1+\sin(2t)/2) \\ \sin(t-\pi/4)(1+\sin(2t)/2) \end{pmatrix},$$

with $0 \leq t < 2\pi$ for a Helmholtz parameter $k = 0.5$. Using either 50 nodes (Fig. 1(a)) or 100 nodes (Fig. 1(b)), we plot the error in both the interior and exterior of Γ. The colors in these figures indicate the absolute value of the pointwise error for the Helmholtz single-layer potential $S\sigma$ with $\sigma \equiv 1$.

These figures clearly suggest that the region in which the layer potential is inaccurate shrinks more or less in proportion to the grid spacing h (a fact well-known to practitioners of potential theory). To be a little more precise, let $T_N(S\sigma)$ denote the trapezoidal approximation of $S\sigma$ using N points, and let

$$E(x) = |S\sigma(x) - T_N(S\sigma)(x)|.$$

For a fixed ϵ, we define the "high-accuracy" region of the plane as the subset of $\mathbb{R}^2$ where $E(x) < \epsilon$. This will, in essence, be all of $\mathbb{R}^2$ with a tubular neighborhood of Γ removed. The extent of this tubular region depends on both h and ϵ [4].

The fact that h-refinement shrinks the region of inaccuracy, of course, is of no great value in evaluating layer potentials at points x on the curve Γ itself. For this, let us instead choose a point c off the surface with

$$c = x + 5h\hat{n}_{x'},$$

where $\hat{n}_x$ is the unit normal to Γ at x. From our initial experiment, it is reasonable to expect that c is in the "high-accuracy" region. Assuming $S\sigma$ is a smooth function in either the interior Ω^- or the exterior Ω^+, it is easy to see that

$$|T_N(S\sigma)(c) - S\sigma(x)| = |(T_N(S\sigma)(c) - S\sigma(c)) + (S\sigma(c) - S\sigma(x))| = O(\epsilon + h) \tag{4}$$

since $|c - x| = O(h)$, under the assumption that the trapezoidal rule is accurate to precision ϵ. In other words, the approximate value $T_N(S\sigma)(c)$ is a first-order accurate approximation of the on-surface value $S\sigma(x)$, within the error ϵ.

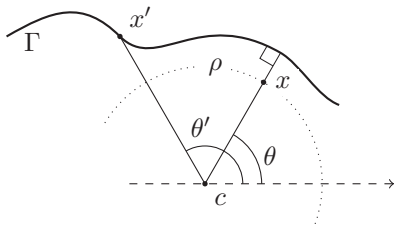

Figure 2. Geometric situation of Graf's addition theorem with sources along the curve Γ, as used in (5) and (6). Note that x will reside on Γ further on in the discussion.

Remarkably, it is straightforward to improve matters even further. Instead of evaluating $S\sigma(c)$, let us *expand* $S\sigma$ about c to order p. The classical separation of variables representation of a smooth solution to the homogeneous Helmholtz equation in a disk centered at c takes the form

$$\phi(x) = \sum_{l=-\infty}^{\infty} \alpha_l J_l(k\rho) e^{-il\theta} \tag{5}$$

where (ρ, θ) denote the polar coordinates of the target x with respect to the expansion center c, and J_l is the Bessel function of order l (see Fig. 2). For the single layer potential $S\sigma$, the coefficients α_l in the expansion (5) can be computed analytically:

$$\alpha_l = \frac{i}{4} \int_\Gamma H_l^{(1)}(k|x' - c|) e^{il\theta'} \sigma(x')\, dx', \qquad (l = -p, -p+1, \ldots, p) \tag{6}$$

where $(|x' - c|, \theta')$ denote the polar coordinates of the point x' with respect to c. These formulas follow immediately from Graf's addition theorem [1],

$$H_0^{(1)}(k|x - x'|) = \sum_{l=-\infty}^{\infty} H_l^{(1)}(k|x' - c|) e^{il\theta'} J_l(k|x - c|) e^{-il\theta}, \tag{7}$$

by interchanging the order of summation and integration. We note that Graf's addition theorem is generally applicable only if the target x is closer to the center than the source x':

$$|x - c| < |x' - c|. \tag{8}$$

The integral defining α_l is similar to that defining the original layer potential, except that $H_0^{(1)}(k|x' - c|)$ has been replaced with the more complicated but still smooth function $H_l^{(1)}(k|x' - c|)e^{il\theta'}$. Because of this smoothness, we evaluate α_l using the same trapezoidal rule $T_N(\alpha_l)$. When seeking to evaluate $S\sigma(x)$, however, we evaluate the local expansion (5) instead. Figure 3(a) shows the result of using an expansion of order $p = 3$ superimposed on the naive use of the trapezoidal rule T_N to compute the single layer potential directly. The circular 'cut-out' regions in this and the following figures indicate where local expansions were used to approximate $S\sigma$. Note that the error in $S\sigma$ computed by the local expansion is *far* smaller than the error of the naive computation throughout the circular region. In effect, the local expansion allows us to punch a disk-shaped hole into the tubular region of inaccuracy. This is precisely the idea underlying the close evaluation scheme of Barnett and Nguyen [4] for targets near, but not on, Γ.

In this paper, we take the approach one step further. Namely, we investigate the use of the expansion (5) in evaluating the layer potential $S\sigma(x)$ for x that actually lie on the curve Γ.

Formally, it is worth noting that the radius of convergence of the local expansion about c is $r = \min_{x' \in \Gamma} |x' - c|$ as Graf's addition theorem requires it. Thus, we are seeking to evaluate a local expansion *at* its radius of convergence where the accuracy is most difficult to analyze. This difficulty, however, stems from the use of the addition theorem for a singular field (the potential due to a point source $H_0^{(1)}$). Figure 3(a) shows that the expansion is, in fact, accurate: it provides about four digits of precision uniformly. Loosely speaking, accuracy follows from the fact that the field induced by the layer potential is (one-sided)

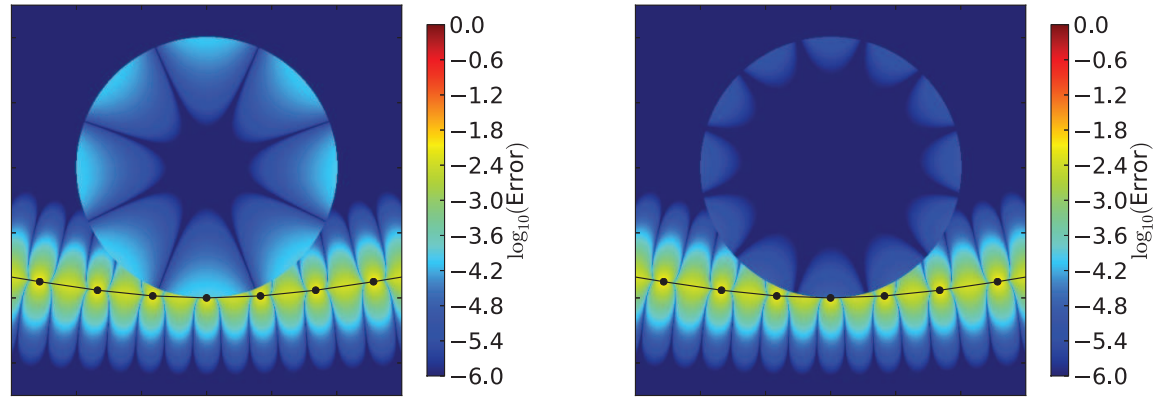

(a) $p = 3$, $N = 80$ quadrature nodes

(b) $p = 6$, $N = 80$ quadrature nodes

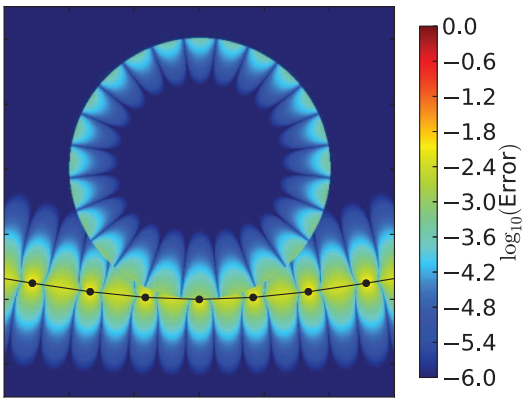

(c) $p = 12$, $N = 80$ quadrature nodes

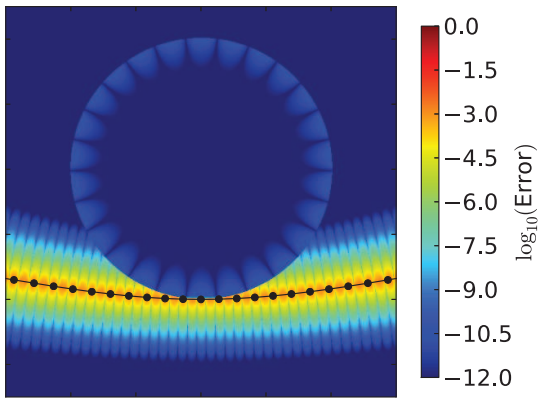

(d) $p = 12$, $N = 240$ quadrature nodes

Figure 3. The potential $S\sigma$ is computed using the trapezoidal rule Γ with either $N = 80$ or $N = 240$ points, except in a disk of radius $|c - x|$ centered at an off-surface point c that lies in the "high-accuracy" region of the trapezoidal rule (here, approximately $3h$ away from the curve). Only a portion of the boundary Γ is plotted, and x is the point where the disk and Γ are tangent. We plot the error in the disk using various expansion orders p and numbers of quadrature nodes.

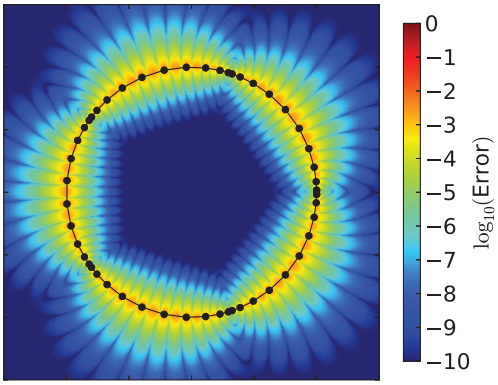

(a) Error in potential from (smooth) composite Gauss-Legendre quadrature, with 5 panels consisting of 10 quadrature nodes each.

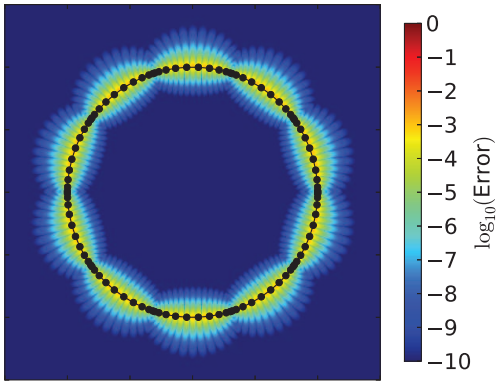

(b) Error in potential from (smooth) composite Gauss-Legendre quadrature, with 10 panels consisting of 10 quadrature nodes each.

Figure 4. The potential $S\sigma$ computed using composite $10th$ order Gauss-Legendre quadrature.

smooth in the interior and exterior domains Ω^- or Ω^+. The analytic issues here are somewhat involved and concern estimates on the decay of the coefficients α_l in terms of the smoothness of the curve $\gamma(t)$ and the density $\sigma(t)$. Those estimates are established in [16], and we will invoke them, as needed, below.

If using an expansion of order $p = 3$ provides an accurate value for $S\sigma(x)$, is it perhaps possible to obtain even more accuracy by further increasing p? Figure 3(b) shows the results of such an experiment. By setting $p = 6$, the accuracy of the potential in the vicinity of (and really also *on* Γ) increases from four to six digits. Further increasing the order to $p = 12$, however, causes a significant *loss* of accuracy (Fig. 3(c)). A consideration of the integrand in (6) shows why this occurs. As p is increased, both factors in the integrand $H_l^{(1)}(k|x' - c|)e^{il\theta'}$ increase in complexity: $e^{il\theta'}$ by becoming more oscillatory, and $H_l^{(1)}(k|x' - c|)$ by becoming more sharply peaked. This combined effect leads to the resolution of the underlying trapezoidal rule being exceeded. Thus, in the experiment of Figure 3(c), the coefficients α_l are both large and wrong. Fortunately, once identified, this issue is easy to resolve. Indeed, simply increasing the number of points in the trapezoidal rule compensates for the added complexity of the integrands involved in computing the higher order coefficients, and with $p = 12$ more than ten digits of accuracy are achieved (Fig. 3(d)).

The sequence of experiments described thus far suggest a path to the high-order accurate evaluation of layer potentials as operators on the boundary. It also highlights one aspect of the scheme that requires careful analysis, namely the interplay between h and p. The local grid spacing h must be chosen small enough so that the coefficients in the local expansion (5) are computed with the necessary precision. We have concentrated in our experiments on a single boundary point x. To evaluate $S\sigma$ everywhere on the boundary, we will simply introduce a large number of off-surface expansions centers whose corresponding disks cover a tubular neighborhood of Γ. We choose our expansion centers so that any desired target point is in the interior or on the boundary of one of these disks, enabling the application of QBX. If a target point is not in any of these disks, by definition it will be in the "high-accuracy" region associated with the original trapezoidal approximation. In the simplest approach, one may introduce an expansion center for each discretization node on the boundary. The procedure applied above to the single layer potential can be used just as well for the evaluation of integrals with hypersingular kernels. These are notoriously difficult using classical quadrature schemes.

None of the observations made above change substantially if we replace the trapezoidal rule with another high-order quadrature. Figure 4 presents the analog of Figure 1 for a circle discretized using composite Gauss-Legendre quadrature. Using ideas from Barnett and Nguyen [4], we believe that the error contours (ignoring nodal oscillations) are the conformal images of the Bernstein ellipses [14] associated with the integrand on each panel.

QBX as a regularization scheme

For readers familiar with multipole/partial wave expansions, the numerical results above may come as a surprise. After all, given a finite set of quadrature nodes (as in our computational examples), if the order of the local expansion were sufficiently high, it should converge to the field induced by a finite set of singular sources, namely the quadrature nodes. Instead, it is reproducing the continuous layer potential, even at the quadrature nodes themselves. There are two interpretations of this fact.

For some, it is most natural to understand this in terms of series approximations of smooth functions, as introduced above. For others, it is perhaps useful to interchange the order of summation and integration and write

$$S\sigma(x) \approx \int_\Gamma G_p(x,x')\sigma(x')\,dx',$$

where

$$G_p(x,x') = \sum_{l=-p}^{p} H_l^{(1)}(k|x'-c|)e^{il\theta'}J_l(k|x-c|)e^{-il\theta}, \tag{9}$$

for a target x on or near the boundary. That is, we can interpret the entire procedure as substituting the original Green's function G with G_p. It turns out that G_p is a surprisingly good filter. It regularizes the kernel in such a way that high-order accuracy is achieved without the need for additional correction. From this perspective, the need to decrease h with p is due to the fact that G_p itself is a more and more sharply peaked integrand as p increases.

3. Mathematical Foundations of QBX

3.1. Error analysis

We turn now to the principal result justifying the use of QBX as a quadrature scheme. We restrict our attention to composite Gauss-Legendre quadrature, but the proof is analogous for any smooth high-order rule. In what follows, h will be used to denote the *panel size* in the composite Gauss-Legendre grid, rather than the point spacing used previously in discussing the trapezoidal rule. We apologize for this abuse of notation.

Theorem 1. *Suppose that Γ is a smooth, bounded curve embedded in $\mathbb{R}^2$, that $B_r(c)$ is the ball of radius r about c, and that $\overline{B_r(c)}\cap\Gamma = \{x\}$. Let Γ be divided into M panels, each of length h and let q be a non-negative integer that defines the number of nodes of the smooth Gaussian quadrature used to compute the coefficients α_l^{QBX} according to the formula (6). For $0<\beta<1$, there are constants $C_{p,\beta}$ and $\tilde{C}_{2q,\beta}$ so that if σ lies in the Hölder space $\mathcal{C}^{2q,\beta}(\Gamma)$, then*

$$\left| S\sigma(x) - \sum_{l=-p}^{p} \alpha_l^{\mathrm{QBX}} J_l(k|x-c|)e^{il\theta_{cx}} \right| \le \left(\underbrace{C_{p,\beta}\, r^{p+1}\|\sigma\|_{\mathcal{C}^{p,\beta}(\Gamma)}}_{\text{Truncation error}} + \underbrace{\tilde{C}_{2q,\beta}\left(\frac{h}{4r}\right)^{2q}\|\sigma\|_{\mathcal{C}^{2q,\beta}(\Gamma)}}_{\text{Quadrature error}} \right). \tag{10}$$

Proof. We begin by writing

$$E = \left| \left(S\sigma(x) - \sum_{l=-p}^{p} \alpha_l J_l(k|x-c|)e^{il\theta_{cx}} \right) + \sum_{l=-p}^{p} \left(\alpha_l - \alpha_l^{\mathrm{QBX}}\right) J_l(k|x-c|)e^{il\theta_{cx}} \right|. \tag{11}$$

The first term stems from using a *truncated* p-term expansion in Bessel functions to approximate $S\sigma$, while the second term is the error that stems from the numerical approximation of the coefficients in the truncated series. It is shown in [16] that the first error is of the order $r^{p+1}\,\|\sigma\|_{\mathcal{C}^{p,\beta}(\Gamma)}$. For the second term, we note that on a curve segment Γ_i of length h, the standard estimate for q-point Gauss-Legendre quadrature is [14, (2.7.12)]

$$\left| \int_{\Gamma_i} H_l^{(1)}(k|x'-c|)e^{il\theta'}\sigma(x')\,dx' - \sum_{j=1}^{q} H_l^{(1)}(k|x_j-c|)e^{il\theta_j}\sigma(x_j)\,w_j \right| \le \frac{h^{2q+1}}{2q+1}\frac{(q!)^4}{(2q)!^3}\|D^{2q}(H_l^{(1)}(k|x'-c|)e^{il\theta'}\sigma(x'))\|_{\infty,\Gamma_i} \tag{12}$$

where D^n denotes the n^{th} derivative of the integrand with respect to the integration parameter along Γ_i. A straightforward combination of Stirling's approximation

$$\sqrt{2\pi}n^{n+\frac{1}{2}}e^{-n} < n! < 2\sqrt{\pi}n^{n+\frac{1}{2}}e^{-n},$$

summing over all panels, and bounds on the derivative allow us to write this term as

$$\left| \sum_{l=-p}^{p} \alpha_l J_l(k|x-c|)e^{il\theta_{cx}} - \sum_{l=-p}^{p} \alpha_l^{\text{QBX}} J_l(k|x-c|)e^{il\theta_{cx}} \right| \leq C_{q,\beta} \left(\frac{h}{4r}\right)^{2q}. \tag{13}$$

Combining the two estimates yields the desired result. □

There are several aspects of the preceding theorem that are worth noting.

- The two contributors to the error in the QBX approximation (10) are quite different. By placing an expansion center at a distance $r = O(h)$ away from the point $x \in \Gamma$, the analytic truncation error is of order h^{p+1}. This error goes to zero under quadrature mesh refinement. In order for the second component of the error to be small, however, we need $\frac{h}{4r} < 1$, so that $r > h/4$. A requirement of this type is essential, corresponding to the fact that if the expansion center is too close to the boundary relative to the discretization, accuracy will be lost. It is perhaps informative to set $r = h/2$, and write the error E from (11) in the form
$$E = O(\epsilon + h^{p+1}),$$
where $\epsilon = \left(\frac{1}{4}\right)^q$. Used in this manner, QBX is not classically convergent, but converges with controlled precision.

- If one wants to achieve a classically convergent scheme, it suffices to refine r more slowly than h (say, with $r = \sqrt{h}$). We prefer to keep the error components separate for the sake of clarity and because it permits additional tests of numerical consistency.

- Note that the estimate (10) explains the behavior of QBX discussed in Section 2, particularly the results shown in Fig. 3(c).

- For the sake of simplicity, Theorem 1 assumes that the curve Γ is divided into equal-sized segments. In practice, with an adaptive discretization of the curve, a slightly different version of the result is needed. Since the difficulty in the error analysis is entirely local and the estimates are similar to those obtained above, we omit the rather cumbersome analysis.

- In practice, one is often interested in evaluating the double-layer potential $D\mu$, or some derivative of $S\sigma$ or $D\mu$. Straightforward analysis shows that, for n derivatives of the Green's function, the error estimate in (10) is multiplied by a factor of r^{-n}. The use of QBX for such calculations is discussed in the next section.

3.2. Derivatives, jumps, and principal value integrals

Up to this point, we have focused on the calculation of the single-layer potential $S\sigma$. For the double layer $D\mu$ defined in (2), the scheme is only slightly different. The coefficients in (6) are simply replaced by

$$\alpha_l^D = \frac{i}{4} \int_\Gamma \frac{\partial}{\partial \hat{n}_{x'}} H_l^{(1)}(k|x'-c|)e^{il\theta'} \mu(x')\, dx'. \qquad (l = -p, -p+1, \ldots, p) \tag{14}$$

Because the scheme relies on a local expansion of the potential, subsequent derivatives of $S\sigma$ or $D\mu$ with respect to the *target* location x are particularly easy to obtain by analytic differentiation of the local (Bessel) expansion.

There is a complication which must be dealt with in evaluating operators other than the single layer potential (which is only weakly singular). QBX, by its construction, evaluates the one-sided limit of a layer

potential, with the side determined by the location of the expansion center. In practice, however, one might want to compute the integral

$$D\mu(x) := \int_\Gamma \frac{\partial G}{\partial \hat{n}_{x'}}(x, x')\mu(x')\, dx', \tag{15}$$

for $x \in \Gamma$. As an operator acting on the boundary, D has a continuous kernel and $D\mu$ is well-defined [37]. $D\mu$ is not, however, equal to its one-sided limit. Using the superscripts $+$ and $-$ to denote a point in the exterior or interior of Γ, respectively, the following jump relations are well-known [3, 7, 37]). For the single-layer potential,

$$S\sigma(x)|_\Gamma = \lim_{x^\pm \to x} S\sigma(x^\pm), \tag{16}$$

for its derivative,

$$\nabla_x S\sigma(x)|_\Gamma = \lim_{x^\pm \to x} \nabla_{x^\pm} S\sigma(x^\pm) \pm \frac{1}{2}\hat{n}\,\sigma(x), \tag{17}$$

where $\nabla_x S\sigma(x)$ is defined in the principal value sense, and for the double-layer potential,

$$D\mu(x)|_\Gamma = \lim_{x^\pm \to x} D\mu(x^\pm) \mp \frac{1}{2}\mu(x). \tag{18}$$

For higher derivatives with respect to the target location x, we have

$$\partial_{x_i}\partial_{x_j} S\sigma(x)|_\Gamma = \lim_{x^\pm \to x} \partial_{x_i^\pm}\partial_{x_j^\pm} S\sigma(x^\pm) \mp \frac{\kappa}{2}(-\delta_{i,j} + 2\hat{n}_i\hat{n}_j)\sigma \pm \frac{1}{2}(\hat{n}_i\hat{t}_j + \hat{t}_j\hat{n}_i)\frac{d\sigma}{ds}, \tag{19}$$

where $\hat{t}$ is the unit tangent, assumed to satisfy the identity

$$\hat{n} = \begin{pmatrix} \hat{t}_2 \\ -\hat{t}_1 \end{pmatrix}.$$

κ here is the curvature, δ is the Kronecker symbol and s is arc length. The expression $\partial_{x_i}\partial_{x_j} S\sigma(x)|_\Gamma$ is defined in the Hadamard finite-part sense.

Finally, in some settings, it is useful to consider derivatives of the double layer, tangentially oriented dipoles, and mixed source/target derivatives. For these, we have

$$\int_\Gamma v(x') \cdot \nabla_{x'} G(x, x')\sigma(x')\, ds\bigg|_\Gamma = \lim_{x^\pm \to x} \left(\int_\Gamma v(x') \cdot \nabla_{x'} G(x^\pm, x')\sigma(x')\, ds \right) \mp \frac{1}{2}(\hat{n} \cdot v(x))\sigma \tag{20}$$

$$\partial_{x_i} D\sigma|_\Gamma = \lim_{x^\pm \to x} (\partial_{x_i^\pm} D\sigma(x\pm)) \mp \frac{1}{2}\hat{t}_i \frac{d\sigma}{ds} \tag{21}$$

In the former expression, v is the direction in which the source derivative is to be taken. When it is tangentially oriented, $v(x') = \hat{t}(x')$, we denote the corresponding operator by R:

$$R\sigma(x) = \int_\Gamma \hat{t}(x') \cdot \nabla_{x'} G(x, x')\sigma(x')\, ds. \tag{22}$$

The jump relations described in (19)-(21) are not so well-known (see, for example, [34]).

In summary, if the one-sided limit is the quantity of interest, then QBX computes that directly and no post-processing work is required. If, however, the principal value integrals $D\mu(x)$, $\nabla_x S\sigma(x)$, or the finite-part integrals in (21) are desired, then additional steps are required. The simplest scheme involves subtracting the relevant quantity from the QBX-derived one-sided limit. This retains the expected order of accuracy. A second option is to compute *both* one-sided limits using QBX and average the quantities appropriately. That is, one can compute

$$D\mu(x) = \frac{1}{2}\left(\lim_{x^+ \to x} D\mu(x^+) + \lim_{x^- \to x} D\mu(x^-) \right) \tag{23}$$

by two applications of QBX.

There are two drawbacks to the latter approach and one advantage. First, it makes the scheme approximately twice as expensive as using a one-sided limit. Second, the limiting values obtained from the two sides of Γ can vary noticeably in their accuracy for a given choice of smooth rule and expansion order. As a result, the error in the principal value computed by averaging is dominated by the worse of the two limits. The advantage of using the two-sided limit is that the Nyström approximation of the operator $D\mu(x)$ is much better behaved spectrally. We discuss this issue in some detail in Section 3.5. When solving integral equations, we believe this advantage outweighs the other considerations.

3.3. Informal description of the algorithm

This section provides a complete description of the steps required to implement QBX. We assume that we are given a smooth curve Γ subdivided into M panels $\Gamma_1, \ldots, \Gamma_M$ of arc lengths $h_1, \ldots, h_M$, respectively and that $\hat{n}$ denotes the outward normal to Γ.

Set up parameters

1. Fix the desired accuracy ϵ.
2. Choose local expansion order p (so that $S\sigma$ will be computed to the order of accuracy $p+1$).
3. Choose q and r such that (10) is approximately satisfied to precision ϵ (assuming the underlying smooth rule is composite Gauss-Legendre quadrature). For points on panel m, a value of $r_m = h_m/2$ works well in practice. (See Section 3.1 for details.)

Compute one-sided limit

4. For each target point $x_j \in \Gamma_m \subset \Gamma$:
   (a) Fix the expansion center $c_j := x_j \mp (h_m/2)\hat{n}$, with $(-)$ corresponding to seeking the interior limit and $(+)$ corresponding to seeking the exterior limit.

   { *If c_j is too close to any other panel $n \neq m$, refine the quadrature ("source") grid (thereby shrinking h_m and moving c_j closer to Γ) until this is no longer case.* }
   (b) Compute the expansion coefficients. For example, for the single layer potential,

$$\alpha_{j,l} := \frac{i}{4} \int_\Gamma H_l^{(1)}(k|x' - c_j|) e^{il\theta'} \sigma(x')\, dx'$$

   for $l = -p, \ldots, p$ using the underlying qth order accurate rule. (See Fig. 2 for the definitions of θ, θ'.)
   (c) Evaluate the local expansion at c_i:

$$u_j := \sum_{l=-p}^{p} \alpha_{j,l} J_l(k|x - c|) e^{-il\theta}.$$

5. If the desired integral is a principal value or finite-part integral that has a jump condition, use the appropriate expression from Section 3.2 to subtract the appropriate term from the one-sided limit (or repeat the calculation with a center on the opposite side and average the two sided limits as in (23)).

A few observations are in order:

- In practice, the error from QBX is greater when the expansion center for a target point x lies on the concave side of the curve rather than the convex side. (See Fig. 1(a) for an illustration and [4] for analytic insight.)
- The algorithm contains a few nested loops, allowing for algorithmic variation. One can save storage, for example, by avoiding the allocation of memory to the expansion coefficients $\alpha_{j,l}$. Each "source" point can compute its contribution to u_j directly.
- The above algorithm directly *applies* the layer potential operator to a given density $\sigma(x')$. It is straightforward to modify the algorithm to compute and store all (or near neighbor) quadrature weights in a table, hence constructing the Nyström matrix approximating the integral operator (see Section 3.4.2).

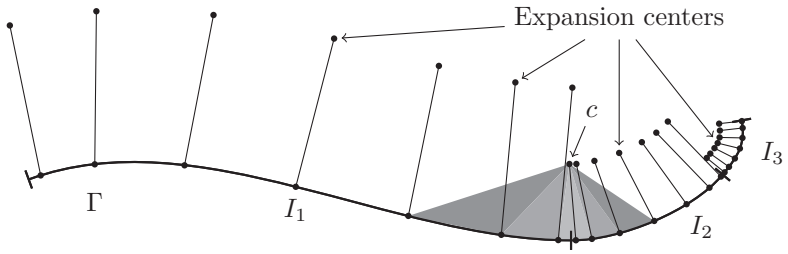

Figure 5. In some settings, one encounters source grids that are highly adaptive, with sudden changes in mesh spacing. This requires some control in QBX to avoid errors in the local expansion approximation (13). See Section 3.4.1 for discussion.

3.4. Remarks on grids

The QBX procedure does not require tight coupling between the "source" and "target" grids. By "source" grid, we mean the set of points on Γ where the density and Green's function are sampled in computing the local expansion coefficients using the underlying smooth quadrature rule. By "target" grid, we mean the set of points along Γ where we seek the value of the layer potential.

3.4.1. Adaptive boundary grids

In our discussion thus far, we have implicitly assumed the subintervals used to divide the boundary Γ are all of approximately the same length. Many applications, of course, are best addressed using some form of adaptive mesh refinement along the curve to resolve either complicated data or to discretize an unknown but complicated density.

The main issue for the application of QBX with such grids concerns the location of the expansion centers and the validity of the error estimate in (10). Fig. 5 illustrates the issue, under the assumption that centers are chosen using the rule $r \approx h_1/2, h_2/2, h_3/2$ on three successive panels I_1, I_2, I_3. Consider, now, the expansion center indicated by c, which is at a distance $r = h_2/2$ from Γ. The filled triangles in the figure illustrate the angles spanned by adjacent source quadrature nodes. The resolution provided in evaluating expansion coefficients at c by the sources on the panel I_1 is clearly much lower than that provided by the sources on I_2 or I_3. Moreover, the assumption that $r > h_1/4$, which is essential in (10) in order for the error to be small, is violated.

Fortunately, this problem is straightforward to address: one simply requires sufficient sampling on I_1 for the error estimate to guarantee high precision. There are several possible strategies in terms of implementation, and we list two here.

- If adjacent panels have substantially different lengths, interpolate the source density on the larger one to a fine grid that matches the resolution of the smaller one on the fly.
- In discretizing the boundary, require that no two adjacent panels differ in length by more than a factor of two and increase the number of points q per panel by a factor of 2.

We use the second (simpler) strategy for the data presented in Section 4.

3.4.2. Grids for solving integral equations

When solving integral equations with a Nyström method, a common grid is used for both sampling the unknown density and evaluating the resulting layer potential. This coincides with what we have referred to as the target grid, which should resolve the curve and the density to the desired precision. This grid may not be sufficiently fine to satisfy the requirements (10) and (8). Under those conditions, in the QBX procedure, one simply needs to interpolate the density from to a finer grid, which becomes what we have referred to as the source grid.

3.5. Spectral structure of operators approximated by QBX

In this section, we consider the spectral structure of the QBX-discretized layer potential operators. In addition to being of mathematical interest, the spectral structure also plays an important role in determining the performance of iterative methods such as GMRES [42] when used to solve integral equations.

We concentrate here on the double layer potential $D\mu(x)$ which plays an important role, for example, in solving the Dirichlet problem for the Helmholtz equation (at a non-resonant frequency k) in the interior Ω^- of Γ. Given Dirichlet data $f(x)$, representing ϕ as a double layer potential

$$\phi(x_0) = D\mu(x_0)$$

for $x_0 \in \Omega^-$, and using the jump relation (18), we obtain the equation

$$(-\frac{1}{2} + D)\mu(x) = f(x) \tag{24}$$

for $x \in \Gamma$.

On smooth boundaries, D is continuous and hence compact, with a discrete, bounded spectrum that has a unique accumulation point at zero [13]. In other words, it is a smoothing operator that damps out the high frequency modes in the density μ. We will show below that QBX is able to preserve all of these properties, most critically the spectral clustering at zero.

We assume we have a grid on Γ with M panels and q points per panel, and that we consider a density that lives in the space of piecewise $(q-1)$th order polynomials over the M panels. We assume the double layer potential is computed using QBX, with values output on the same grid, corresponding to a discrete $Mq \times Mq$ matrix.

Now, because the truncated p-term expansion represents locally smooth functions obeying a band limit related to p, high-frequency components of the density are either attenuated or aliased to lower frequencies as they transition from the source grid to the expansion. Empirically, attenuation is the dominant effect.

This behavior has several consequences. A beneficial feature is that QBX responds very benignly to potentially erroneous high-frequency data that may be present in the discretized densities or geometries. Also, some spectral features are reproduced with no further effort. For example, when applied to the single layer potential, the QBX-based one-sided limit faithfully reproduces the spectrum of the continuous operator, accumulating at zero.

Unfortunately, when computing the *double-layer potential* D using QBX based on the one-sided limit as

$$D_{h,\text{one-sided}}\ \mu(x) = \lim_{x^\pm \to x} D_h\mu(x^\pm) \mp \frac{1}{2}\mu(x)\,,$$

high frequency components are attenuated in $\lim_{x^\pm \to x} D_h\mu(x^\pm)$ but not, of course, in $\frac{1}{2}\mu(x)$. As a result, the spectrum of $D_{h,\text{one-sided}}$ does *not* accumulate at zero.

If one then solves the integral equation (24) iteratively, with $D_{h,\text{one-sided}}$ computed in this manner, then an iterative method will converge rapidly *up to* the level of discretization error, at which point it will stall. Since one does not know *a priori* exactly what the discretization error will be, this is rather inconvenient. Fortunately, computing D using the two-sided averaging approach (23) as discussed in Section 3.2 yields a discrete operator with a spectrum accumulating at zero, because both limits are filtered. Matching this feature of the continuous operator allows iterative linear solvers converge to rapidly to solutions having residuals near machine precision even *beyond* the level of discretization error. Using two-sided averaging as in (23) may not be the only way to achieve this spectral behavior, but it is particularly convenient.

The preceding discussion applies only to the compact case. For hypersingular (finite-part) integrals or Hilbert-Riesz type operators such as R in (22), there is no particular advantage in using the two-sided limit. Also, as discussed above, if jump conditions are not invoked, operators such as the (compact) single-layer potential can be represented faithfully by the one-sided procedure without difficulty. In the numerical results shown in Section 4, we have used two-sided averaging for all operators to which it applies. A more detailed discussion of the spectral properties of integral operators computed using QBX will be reported at a later date.

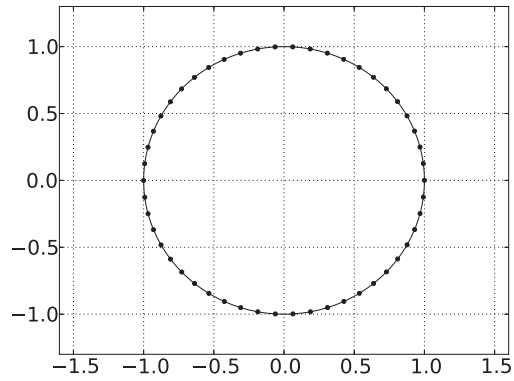

(a) A circle, decomposed into 50 panels, given by (25) with $\alpha = 1$.

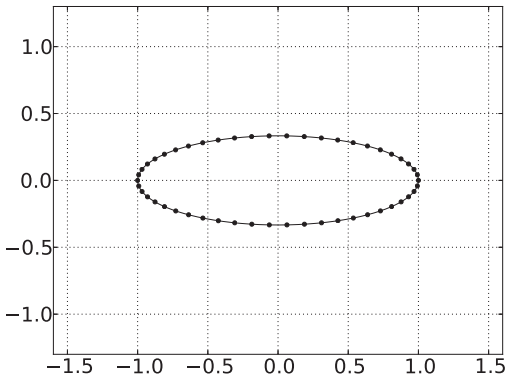

(b) An ellipse of aspect ratio 3:1 decomposed into 50 panels, given by (25) with $\alpha = 3$.

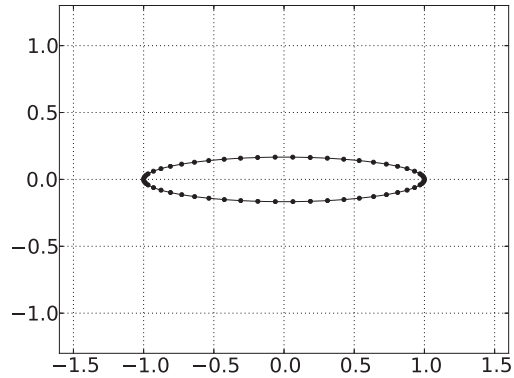

(c) An ellipse of aspect ratio 6:1 decomposed into 58 panels, given by (25) with $\alpha = 6$.

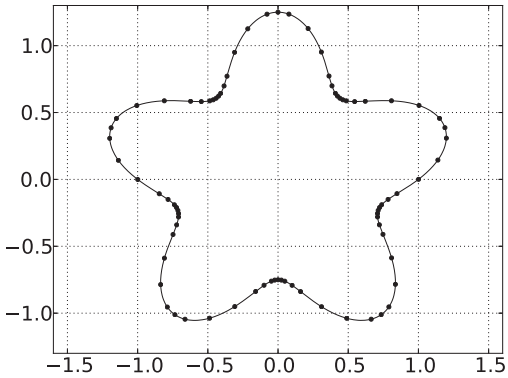

(d) A starfish-shaped curve decomposed into 80 panels, given by (26).

Figure 6. Test geometries and their panel subdivisions.

4. Numerical experiments

In this section, we illustrate the performance of QBX. We begin by describing some simple test geometries. We then present results for a variety of layer potentials, showing that high accuracy can be achieved even with modest-sized discretizations. Finally, we investigate the performance of QBX when used as part of an integral equation solver for a variety of Dirichlet and Neumann boundary value problems at various orders of accuracy.

For the sake of convenience, we will denote the normal derivatives of the single and double layer potentials by

$$S'\sigma(x) := \hat{n}(x) \cdot \nabla_x S\sigma(x),$$
$$D'\sigma(x) := \hat{n}(x) \cdot \nabla_x D\sigma(x).$$

When $x \in \Gamma$, the first is meant in the principal value sense and the second in the Hadamard finite-part sense.

4.1. Four test geometries

The four curves that we will use for our numerical tests are shown in Fig. 6. The ellipses of Figs. 6(a), 6(b), 6(c) are given by

$$\gamma(t) = \begin{pmatrix} \cos(2\pi t) \\ \frac{1}{\alpha}\sin(2\pi t) \end{pmatrix} \tag{25}$$

for $\alpha = 1$, 3, and 6, respectively, and the "starfish" of Figure 6(d) is given by

$$\gamma(t) = \left(1 + \frac{\sin(5 \cdot 2\pi t)}{4}\right) \begin{pmatrix} \cos(2\pi t) \\ \sin(2\pi t) \end{pmatrix}. \tag{26}$$

(a) Circle using $K = 50$ panels

Operator	L^2 Error	L^∞ Error
S	2.0e-15	4.2e-15
D	2.2e-13	1.9e-13
R	1.1e-12	6.7e-13
$\partial_x S$	9.7e-13	6.8e-13
$\partial_x D$	2.4e-11	1.5e-10
$\partial_y S$	9.8e-13	6.9e-13
$\partial_y D$	2.3e-11	1.7e-10
$\partial_x \partial_x S$	8.0e-10	3.5e-09
$\partial_x \partial_y S$	8.0e-10	3.5e-09
$\partial_y \partial_y S$	8.0e-10	3.5e-09

(b) 3-to-1 ellipse using $K = 50$ panels

Operator	L^2 Error	L^∞ Error
S	1.5e-15	2.6e-15
D	3.7e-13	4.5e-13
R	3.5e-12	4.9e-12
$\partial_x S$	3.6e-12	5.0e-12
$\partial_x D$	4.3e-11	3.6e-10
$\partial_y S$	1.2e-12	3.1e-12
$\partial_y D$	5.4e-11	1.5e-10
$\partial_x \partial_x S$	6.7e-10	4.5e-09
$\partial_x \partial_y S$	1.0e-09	7.5e-09
$\partial_y \partial_y S$	6.7e-10	4.5e-09

(c) 6-to-1 ellipse using $K = 58$ panels

Operator	L^2 Error	L^∞ Error
S	2.5e-14	1.6e-13
D	5.9e-13	1.4e-12
R	1.1e-10	2.6e-10
$\partial_x S$	1.1e-10	2.6e-10
$\partial_x D$	8.9e-08	9.9e-07
$\partial_y S$	2.6e-12	1.3e-11
$\partial_y D$	4.4e-09	2.0e-08
$\partial_x \partial_x S$	4.4e-09	1.8e-08
$\partial_x \partial_y S$	9.0e-08	9.9e-07
$\partial_y \partial_y S$	4.4e-09	1.8e-08

(d) Starfish using $K = 80$ panels

Operator	L^2 Error	L^∞ Error
S	1.4e-14	1.0e-14
D	6.8e-13	9.5e-13
R	7.1e-13	8.2e-13
$\partial_x S$	5.0e-12	9.8e-12
$\partial_x D$	8.2e-11	7.7e-10
$\partial_y S$	5.0e-12	9.9e-12
$\partial_y D$	8.5e-11	1.2e-09
$\partial_x \partial_x S$	9.5e-10	1.1e-08
$\partial_x \partial_y S$	9.5e-10	1.1e-08
$\partial_y \partial_y S$	9.5e-10	1.1e-08

Table 1. Relative error of QBX integral operators in the L^2 and L^∞ norms compared to a high-accuracy reference solution, using local expansions of order $p = 16$.

In each of these cases, $t \in [0, 1)$.

We decompose the curves $\gamma(t) = (x(t), y(t))$ into panels, with $x(t)$ and $y(t)$ represented by a 16-term Legendre polynomial expansion. We generate an initial subdivision that is equispaced in t. To ensure the accuracy of the expansion, we sample the curve at $64 = 4 \cdot 16$ points per panel and compute Legendre expansion coefficients by Gauss-Legendre quadrature.

Since we are integrating with respect to the parameter t rather than arc length, we first determine whether the curve is well-resolved by studying the spectral decay of the Legendre coefficients of $|\gamma'(t)| = \sqrt{\gamma_1'(t)^2 + \gamma_2'(t)^2}$, using the method of Klöckner et al. [31]. We then determine the L^2 energy contained in the tail of the series (i.e. in modes 16 and above in our case). If the estimated residual exceeds 10^{-11}, the panel is bisected. A panel is also bisected if its length (as computed by integrating $|\gamma'|$) is more than twice that of its neighboring panels, to avoid the issues described in Section 3.4.1.

A source oversampling factor (see Section 3.4) of 6 is used throughout, that is $q = 6 \cdot 16$. A factor of 2 is included to allow adjacent panels to differ in length by factors of two, and an additional factor of 3 is included to ensure that the second term in the estimate (10) is negligible. More judicious oversampling strategies will be considered at a later date.

4.2. Layer Potential Evaluation

Our first set of tests examines the ability of QBX to compute a range of standard and non-standard layer potential operators to high precision. We consider the operators S, $\partial_x S$, $\partial_y S$, $\partial_{xx} S$, $\partial_{xy} S$, $\partial_{yy} S$ (target derivatives of the single-layer potential) as well as D, $\partial_x D$, $\partial_y D$ (target derivatives of the double-layer potential). We also consider the layer potential induced by tangentially oriented dipoles (a source derivative in the tangential direction), which we denoted earlier by by R. R is the analog for Helmholtz potentials of the Hilbert transform in two dimensions or the Riesz transform in three dimensions.

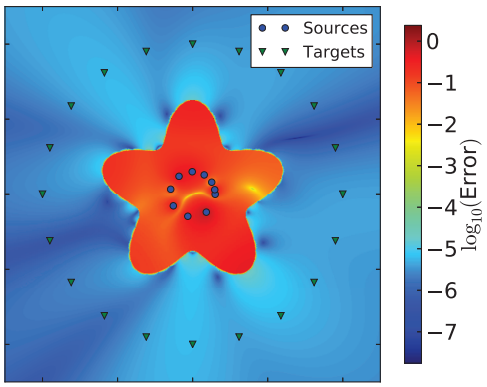

(a) Sources and targets for the test of the solution of an exterior boundary value problem.

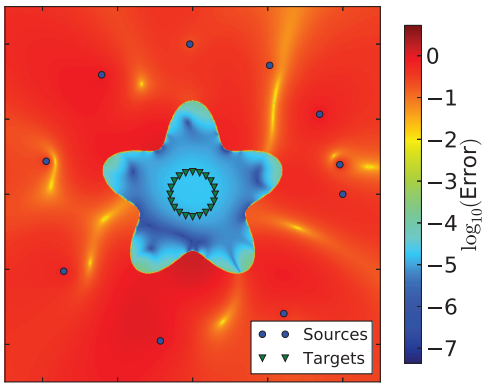

(b) Sources and targets for the test of the solution of an interior boundary value problem.

Figure 7. Setup of the integral equation test cases, shown with the 'starfish' geometry of Figure 6(d).

BC type	Side	k	p	$M = 30$	$M = 60$	$M = 90$	$M = 120$	EOC
Dirichlet	ext	1	1	5.3e-04 (7)	1.4e-04 (7)	6.3e-05 (7)	3.6e-05 (7)	**1.9**
			3	1.4e-06 (7)	9.5e-08 (7)	1.9e-08 (7)	6.2e-09 (7)	**3.9**
			5	3.4e-09 (7)	6.2e-11 (7)	5.7e-12 (7)	1.0e-12 (7)	**5.8**
		6	1	1.1e-02 (9)	3.0e-03 (9)	1.4e-03 (9)	7.8e-04 (9)	**1.9**
			3	2.8e-05 (9)	1.9e-06 (9)	3.8e-07 (9)	1.2e-07 (9)	**3.9**
			5	1.5e-08 (9)	2.3e-10 (9)	2.0e-11 (9)	3.7e-12 (9)	**6.0**
Neumann	ext	1	1	3.9e-02 (7)	2.0e-02 (8)	1.3e-02 (8)	1.0e-02 (7)	**1.0**
			3	1.6e-04 (7)	2.2e-05 (8)	6.7e-06 (8)	2.9e-06 (7)	**2.9**
			5	5.9e-07 (7)	2.1e-08 (8)	2.9e-09 (8)	7.2e-10 (7)	**4.8**
		6	1	2.1e-01 (12)	8.4e-02 (12)	5.2e-02 (12)	3.8e-02 (13)	**1.2**
			3	2.7e-04 (12)	4.6e-05 (12)	1.5e-05 (12)	6.5e-06 (14)	**2.7**
			5	3.9e-07 (12)	1.2e-08 (12)	1.6e-09 (14)	3.8e-10 (14)	**5.0**

Table 2. Convergence in the l^2 norm of the solution evaluated at a set of targets after solving a boundary value problem using an integral equation and QBX on the circle of Figure 6(a). GMRES iteration counts are shown in parentheses next to the error data. "EOC" is the empirical order of convergence, obtained by a log-least-squares fit of the shown l^2 errors.

We apply each of these operators to the density $\sigma(t) = \sin(10\pi t)$ and compare the computed result to a reference solution in the L^2 and L^∞ norms. The Helmholtz parameter was chosen as $k = 0.5$. The computations were carried out with local expansion order $p = 16$. We obtained our reference solution by using adaptive Gaussian quadrature with tolerance 10^{-12} in quadruple precision with target points at distances 10^{-6}, $10^{-6}/2$, and $10^{-6}/4$ from the curve along the normal on either side. We then computed one-sided limits v^+ and v^- on each side by third-order Richardson extrapolation. We computed the value $(v^+ + v^-)/2$ as the reference solution for principal value or finite-part on-surface integrals. Results are shown in Table 1, confirming that high accuracy is achievable with modest computational effort, as expected from a rapidly convergent scheme. We further note that operators involving derivatives with tangential components to the curve are either hypersingular or bounded (but not compact). Since differentiation is ill-conditioned, one should expect some loss of accuracy with successively higher derivatives.

4.3. Integral equation solvers

Our second and perhaps more important test examines the suitability of QBX in the context of solving integral equations of the second kind. For each interior domain, we define an exact solution as the field induced by a collection of singular point sources in the exterior. For each exterior domain, an exact solution is constructed using singular point sources in the interior. Given the exact solution, we compute either

BC type	Side	k	p	$M = 34$	$M = 60$	$M = 90$	$M = 120$	EOC
Dirichlet	int	1	1	2.7e-03 (15)	6.9e-04 (15)	3.1e-04 (15)	1.7e-04 (15)	**2.2**
			3	2.8e-06 (15)	1.8e-07 (15)	3.6e-08 (15)	1.1e-08 (15)	**4.4**
			5	5.0e-09 (15)	1.1e-10 (15)	1.0e-11 (15)	1.8e-12 (15)	**6.3**
		6	1	2.0e-02 (19)	5.2e-03 (19)	2.3e-03 (19)	1.3e-03 (19)	**2.2**
			3	4.7e-05 (19)	3.0e-06 (19)	6.2e-07 (19)	2.0e-07 (19)	**4.3**
			5	2.6e-07 (19)	6.7e-09 (19)	6.2e-10 (19)	1.1e-10 (19)	**6.1**
Neumann	int	1	1	2.5e-01 (18)	1.5e-01 (15)	9.8e-02 (14)	7.4e-02 (14)	**1.0**
			3	3.2e-03 (18)	3.7e-05 (13)	1.1e-05 (13)	4.9e-06 (13)	**5.1**
			5	9.3e-05 (16)	1.7e-08 (13)	2.2e-09 (13)	5.4e-10 (13)	**9.5**
		6	1	3.8e-01 (27)	1.6e-01 (20)	1.0e-01 (19)	7.5e-02 (19)	**1.3**
			3	1.2e-03 (26)	1.6e-04 (20)	4.7e-05 (21)	2.0e-05 (20)	**3.3**
			5	3.4e-05 (22)	1.8e-07 (21)	2.5e-08 (20)	6.2e-09 (20)	**6.8**

Table 3. Convergence in the l^2 norm of the solution evaluated at a set of targets after solving a boundary value problem using an integral equation and QBX on the 3-to-1 ellipse of Figure 6(b). GMRES iteration counts are shown in parentheses next to the error data. "EOC" is the empirical order of convergence, obtained by a log-least-squares fit of the shown l^2 errors.

BC type	Side	k	p	$M = 42$	$M = 68$	$M = 102$	$M = 137$	EOC
Dirichlet	ext	1	1	1.1e-03 (19)	2.7e-04 (19)	1.2e-04 (19)	6.7e-05 (19)	**2.3**
			3	3.3e-06 (19)	2.0e-07 (19)	4.0e-08 (19)	1.3e-08 (19)	**4.7**
			5	9.6e-09 (19)	3.7e-10 (19)	3.2e-11 (19)	4.8e-12 (19)	**6.4**
		6	1	1.0e-02 (21)	2.6e-03 (21)	1.2e-03 (22)	6.5e-04 (22)	**2.3**
			3	1.6e-05 (22)	1.0e-06 (22)	2.1e-07 (22)	6.7e-08 (22)	**4.6**
			5	8.1e-08 (22)	1.6e-09 (22)	1.4e-10 (22)	2.5e-11 (22)	**6.8**
Neumann	ext	1	1	3.9e-02 (23)	1.8e-02 (23)	1.2e-02 (23)	9.0e-03 (23)	**1.2**
			3	2.6e-04 (24)	2.9e-05 (22)	8.4e-06 (21)	3.5e-06 (21)	**3.6**
			5	1.3e-06 (23)	5.6e-08 (21)	7.9e-09 (20)	1.9e-09 (20)	**5.5**
		6	1	5.7e-02 (37)	3.2e-02 (36)	2.3e-02 (36)	1.7e-02 (35)	**1.0**
			3	4.0e-04 (36)	5.4e-05 (35)	1.6e-05 (33)	6.9e-06 (33)	**3.4**
			5	3.5e-06 (35)	1.3e-07 (32)	1.8e-08 (28)	4.4e-09 (27)	**5.6**

Table 4. Convergence in the l^2 norm of the solution evaluated at a set of targets after solving a boundary value problem using an integral equation and QBX on the 6-to-1 ellipse of Figure 6(c). GMRES iteration counts are shown in parentheses next to the error data. "EOC" is the empirical order of convergence, obtained by a log-least-squares fit of the shown l^2 errors.

BC type	Side	k	p	$M = 70$	$M = 105$	$M = 130$	EOC
Dirichlet	int	1	1	1.7e-04 (16)	7.8e-05 (16)	8.0e-05 (16)	**1.4**
			3	1.8e-06 (16)	4.4e-07 (16)	1.7e-07 (16)	**3.8**
			5	5.7e-08 (16)	5.6e-09 (16)	3.7e-09 (16)	**4.6**
		6	1	2.9e-02 (25)	1.3e-02 (25)	7.6e-03 (24)	**2.2**
			3	8.1e-05 (25)	1.8e-05 (25)	6.1e-06 (25)	**4.1**
			5	9.0e-07 (25)	9.7e-08 (25)	2.0e-08 (25)	**6.1**
Neumann	int	1	1	7.5e-02 (20)	5.4e-02 (20)	4.1e-02 (21)	**0.9**
			3	9.7e-04 (19)	3.2e-04 (19)	1.5e-04 (19)	**3.0**
			5	1.2e-05 (18)	2.0e-06 (18)	6.6e-07 (18)	**4.7**
		6	1	3.7e-01 (61)	3.2e-01 (61)	2.4e-01 (61)	**0.7**
			3	2.8e-03 (60)	9.6e-04 (60)	4.4e-04 (60)	**3.0**
			5	4.1e-05 (50)	6.5e-06 (50)	1.8e-06 (50)	**4.9**

Table 5. Convergence in the l^2 norm of the solution evaluated at a set of targets after solving a boundary value problem using an integral equation and QBX on the "starfish" geometry of Figure 6(d). GMRES iteration counts are shown in parentheses next to the error data. "EOC" is the empirical order of convergence, obtained by a log-least-squares fit of the shown l^2 errors.

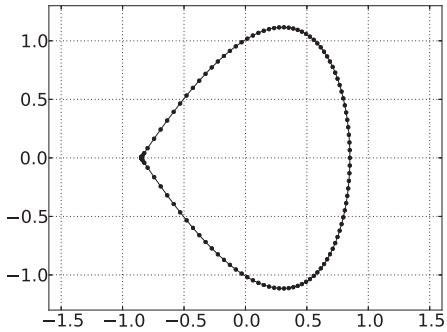

(a) Discretization of the singular 'teardrop' geometry, using 174 panels of order 16. The panels near the corner are dyadically refined until the smallest one has length 10^{-8}.

BC type	Side	k	p	$M = 80$	$M = 138$	EOC
Dirichlet	int	1	1	1.6e-03 (20)	2.1e-04 (21)	**3.7**
			3	5.9e-06 (20)	2.0e-07 (21)	**6.2**
			5	3.3e-08 (20)	3.3e-09 (21)	**4.2**
		6	1	6.9e-02 (38)	8.7e-03 (38)	**3.8**
			3	1.1e-04 (38)	3.8e-06 (38)	**6.1**
			5	1.0e-06 (38)	2.4e-08 (38)	**6.9**
	ext	1	1	3.4e-04 (19)	5.2e-05 (19)	**3.5**
			3	2.2e-06 (19)	8.2e-08 (19)	**6.0**
			5	1.3e-08 (19)	1.6e-09 (19)	**3.9**
		6	1	1.6e-02 (33)	1.9e-03 (33)	**3.9**
			3	4.5e-05 (33)	1.3e-06 (33)	**6.5**
			5	1.4e-07 (33)	1.3e-08 (33)	**4.4**

(b) Convergence in the l^2 norm of the solution evaluated at a set of targets after solving a boundary value problem using an integral equation and QBX on the 'teardrop' geometry of Figure 8(a).

Figure 8. Integral equation tests on a 'teardrop' geometry with a corner.

Dirichlet or Neumann data and solve the corresponding boundary value problem using an integral equation. We then test the accuracy of the solution at a set of target points. Figure 7 illustrates the geometry of our tests.

For the Dirichlet problem, we use the combined-field representation $u = -D_k\sigma + \alpha S_k\sigma$ [13], which leads to the second-kind equation

$$\mp\frac{1}{2}\sigma + \alpha S_k\sigma - D_k\sigma = f, \tag{27}$$

for σ, where f is the Dirichlet data obtained from the manufactured solution. Since the wave number is not large, we choose $\alpha = i$ throughout this section. The $(+)$ sign corresponds to the exterior problem and the $(-)$ sign to the interior problem. The subscript k in S_k and D_k is used to emphasize that the underlying Green's function is that for the Helmholtz equation with Helmholtz parameter k.

For the Neumann problem, we use a slight variation on the well-known combined-field representation [6, 11, 13, 38, 41]

$$u = S_k\sigma - \alpha D_k S_0\sigma,$$

where once again subscripts of k indicate the use of the Helmholtz kernel with parameter k, and a subscript of 0 indicates the use of a Laplace kernel. This representation leads to the second-kind integral equation

$$\mp\frac{1}{2}\sigma + S_k'\sigma - \alpha D_k' S_0\sigma = f \tag{28}$$

for σ, where f is the Neumann data obtained from the manufactured solution. Since QBX can integrate hypersingular kernels, we use (28) directly. One may also use the Calderón projection identity

$$D_0' S_0 = -I/4 + S_0' S_0' \tag{29}$$

[40] and some algebra to avoid hypersingular operators.

Given our boundary discretization, we assume the unknowns are point values of σ at the source nodes and enforce the integral equation at the same nodes, corresponding to a Nyström method. We use GMRES to solve (27) or (28) iteratively and QBX to carry out the matrix-vector products. We set the GMRES tolerance to 10^{-14} independent of the order of accuracy of the QBX-based quadrature.

Following the work of Bremer [8], we use as unknowns the density values multiplied by the square root of the corresponding quadrature weight. This has the effect that the discrete l^2 inner product approximates the continuous L^2 inner product and results in much improved conditioning, especially in the presence of widely varying panel sizes. (This is critical in geometries with corners, as discussed in the next section.)

After solving for σ in (27) or (28), we use the corresponding representation to evaluate the potential u at a number of target points in Ω. We then compare those values to the exact (manufactured) solution and compute the relative error.

Results for the geometries described in Section 4.1 are shown in Tables 2, 3, 4, and 5. We observe that, while slightly more erratic, the results for the Neumann operator exhibit the loss of one order of accuracy, as expected since we have used QBX for a hypersingular kernel.

Of particular note is the fact that, as predicted in Section 3.5, GMRES iteration reaches a residual of 10^{-14} with a modest number of iterations even for low order accurate discretization, in nearly all cases. This makes QBX-based solvers particularly robust.

4.3.1. Non-smooth geometries

In the derivation of QBX, we have assumed that the layer potential is locally smooth, so that an expansion in Bessel functions is rapidly convergent. Since many engineering problems involve geometries with corners (and therefore potentially non-smooth densities and unbounded layer potentials), it is of interest to study whether QBX can be used effectively for such problems as well.

Without knowing the precise singularity in the density, it is shown in [8, 26] that high-order polynomial approximation combined with high-order quadrature on a dyadically refined mesh yields high-order accuracy. Thus, the only question is whether QBX can evaluate layer potentials on such structures without excessive work. To this end, we consider a 'teardrop' shape with a single corner, described by the parametrization

$$\gamma(t) = 1.7\begin{pmatrix} \sin(\pi t) - 0.5 \\ \frac{1}{2}\cos(\pi t)(\pi t - \pi)\pi t \end{pmatrix}.$$

The curve and its discretization are shown in Figure 8(a). We have dyadically refined the boundary toward the corner until the smallest panel lengths are less than 10^{-8} on each side. Carrying out the same type of experiment as in the preceding section, we obtain the results in Table 8(b). The apparent drop in convergence order for $p = 5$ can be attributed to the error made in halting dyadic subdivision at $\epsilon = 10^{-8}$. These experiments demonstrate that there are no significant obstacles to using QBX in this context. The method converges here precisely because the layer potential is locally smooth on finer and finer length scales as one approaches the corner.

5. Generalizations and implementation issues

Our goal in this paper has been to present a new approach to quadrature, which we refer to as QBX ('quadrature by expansion'). While we have largely limited our attention to the Helmholtz equation in two dimensions, it should be clear that the overall approach is independent of dimension as well as the precise nature of the governing Green's function. In fact, the QBX approach is far *more* general, extending to kernels that are not directly connected to a partial differential equation. These extensions are discussed in [32]. The method gives rise to a number of important and interesting questions regarding efficiency, robustness, and automatic adaptivity.

As for implementation, the main issue we have ignored here is computational cost. As presented in Section 3.3, the asymptotic complexity of QBX is $O(NN_t)$, where N is the number of source points, N_t is the number of target points. Neglecting numerous opportunities for optimization, a straight implementation of the algorithm of Section 3.3 can apply a single-layer operator at order $p = 5$ to a density on 7680 source nodes with 1280 targets in 0.6 seconds using 16 cores of a 2.93 GHz Intel Nehalem machine. Fortunately, QBX can be accelerated using the fast multipole method (FMM) or any other hierarchical fast algorithm [12, 20]. The cost is then $O(N \log N + N_t \log N_t)$. This coupling, further cost savings, as well as extensions to three dimensions, are discussed in [21]. To give an indication of the achievable speedups, preliminary implementations show that the cost of an FMM-based QBX scheme for a layer potential is about twice that for a point-to-point FMM procedure. In particular, layer potentials with tens of thousands of discretization points are computed in seconds on a single CPU core. A variety of other optimizations are also possible: using direct evaluation for distant panel interactions and QBX for near neighbors only, adaptive oversampling to ensure accuracy of the local expansion coefficients with highly adaptive and irregular panel sizes, using the sample local expansion for several nearby target points, etc.

6. Conclusions

QBX permits the rapid, high-order accurate evaluation of layer potentials in a manner that is remarkably easy to implement. It is based on the fact that the induced potential is smooth in the exterior or interior domain. The scheme is equipped with a complete convergence theory. With minor modifications, QBX can evaluate layer potentials at off-surface points arbitrarily close to the boundary. Since it is an extension of the scheme developed in [4] for precisely that purpose, this is not a surprise. QBX presents an opportunity to develop a robust set of software tools for evaluating integral operators with singular or weakly singular kernels, with application to a broad range of large-scale simulations in physics and engineering.

Acknowledgments

The authors would like to thank Z. Gimbutas, C. Epstein, J.-Y. Lee, S. Jiang, S. Veerapaneni, M. Tygert, and T. Warburton for fruitful discussions. AK would also like to acknowledge the use of computing resources supplied by T. Warburton. The authors' work was supported through the AFOSR/NSSEFF Program Award FA9550-10-1-0180, NSF grant number DMS-0811005, and also under contract DEFG0288ER25053 by the Department of Energy.

References

[1] M. Abramowitz and I. Stegun. *Handbook of Mathematical Functions*. Dover, 1965.

[2] B. K. Alpert. Hybrid Gauss-Trapezoidal quadrature rules. *SIAM Journal on Scientific Computing*, 20 (5):1551–1584, Jan. 1999. doi: 10.1137/S1064827597325141.

[3] K. E. Atkinson. *The Numerical Solution of Integral Equations of the Second Kind*. Cambridge University Press, Cambridge, UK, 1997.

[4] A. Barnett and H. Nguyen. Evaluation of layer potentials close to the boundary for Laplace and Helmholtz problems on analytic planar domains. Technical report, Dartmouth College, 2012. in prep.

[5] J. T. Beale and M.-C. Lai. A method for computing nearly singular integrals. *SIAM Journal on Scientific Computing*, 38(6):1902–1925, 2001. doi: 10.1137/S0036142999362845.

[6] H. Brakhage and P. Werner. Über das Dirichletsche Außenraumproblem für die Helmholtzsche Schwingungsgleichung. *Archiv der Mathematik*, 16(1):325–329, 1965. doi: 10.1007/BF01220037.

[7] C. A. Brebbia, J. C. F. Telles, and L. C. Wrobel. *Boundary Element Techniques*. Springer, New York, 1984.

[8] J. Bremer. On the Nyström discretization of integral equations on planar curves with corners. *Applied and Computational Harmonic Analysis*, 32(1):45 – 64, 2012. doi: 10.1016/j.acha.2011.03.002.

[9] J. Bremer, Z. Gimbutas, and V. Rokhlin. A nonlinear optimization procedure for generalized Gaussian quadratures. *SIAM Journal on Scientific Computing*, 32:1761–1788, 2010. doi: 10.1137/080737046.

[10] O. P. Bruno and L. A. Kunyansky. A fast, high-order algorithm for the solution of surface scattering problems: basic implementation, tests, and applications. *Journal of Computational Physics*, 169:80–110, 2001. doi: 10.1006/jcph.2001.6714.

[11] O. P. Bruno, T. Elling, and C. Turc. Fast high-order algorithms and well-conditioned integral equations for high-frequency sound-hard scattering problems. *Int. Journal Numerical Methods in Engineering*, 2012. URL `http://filer.case.edu/cct21/Scattering_Neumann_BET.pdf`. To appear.

[12] H. Cheng, W. Crutchfield, Z. Gimbutas, L. Greengard, J. Huang, V. Rokhlin, N. Yarvin, and J. Zhao. Remarks on the implementation of the wideband FMM for the Helmholtz equation in two dimensions. *Contemporary Mathematics*, 408:99, 2006. doi: 10.1090/conm/408.

[13] D. Colton and R. Kress. *Inverse Acoustic and Electromagnetic Scattering Theory*. Springer, 2nd edition, Jan. 1998. ISBN 354062838X.

[14] P. J. Davis and P. Rabinowitz. *Methods of Numerical Integration*. Academic Press, San Diego, 1984.

[15] M. G. Duffy. Quadrature over a pyramid or cube of integrands with a singularity at a vertex. *SIAM Journal on Numerical Analysis*, 19(6):1260–1262, 1982. doi: 10.1137/0719090.

[16] C. Epstein, L. Greengard, and A. Klöckner. On the convergence of local expansions of layer potentials. Technical report, Courant Institute, 2012. in prep.

[17] L. Farina. Evaluation of single layer potentials over curved surfaces. *SIAM Journal on Scientific Computing*, 23(1):81–91, 2001. doi: 10.1137/S1064827599363393.

[18] J. Goodman, T. Y. Hou, and L. J. High-order and efficient methods for the vorticity formulation of the Euler equations. *Communications on Pure and Applied Mathematics*, 43:415–430, 1990. doi: 10.1002/cpa.3160430305.

[19] R. D. Graglia and G. Lombardi. Machine Precision Evaluation of Singular and Nearly Singular Potential Integrals by Use of Gauss Quadrature Formulas for Rational Functions. *IEEE Transactions on Antennas and Propagation*, 56(4):981–998, 2008. doi: 10.1109/TAP.2008.919181.

[20] L. Greengard and V. Rokhlin. A fast algorithm for particle simulations. *Journal of Computational Physics*, 73(2):325–348, 1987. doi: 10.1016/0021-9991(87)90140-9.

[21] L. Greengard, A. Klöckner, and Z. Gimbutas. Fast Algorithms for the Evaluation of Layer Potentials using 'Quadrature by Expansion' . Technical report, Courant Institute, 2012. in prep.

[22] W. Hackbusch and S. A. Sauter. On numerical cubatures of nearly singular surface integrals arising in BEM collocation. *Computing*, 52(2):139–159, 1994. doi: 10.1007/BF02238073.

[23] D. J. Haroldsen and D. I. Meiron. Numerical Calculation of Three-dimensional Interfacial Potential Flows using the Point Vortex Method. *Communications on Pure and Applied Mathematics*, 43:415–430, 1990. doi: 10.1137/S1064827596302060.

[24] J. Helsing. Integral equation methods for elliptic problems with boundary conditions of mixed type. *Journal of Computational Physics*, 228(23):8892–8907, Dec. 2009. ISSN 0021-9991. doi: 10.1016/j.jcp.2009.09.004.

[25] J. Helsing and R. Ojala. On the evaluation of layer potentials close to their sources. *Journal of Computational Physics*, 227:2899–2921, 2008. doi: 10.1016/j.jcp.2007.11.024.

[26] J. Helsing and R. Ojala. Corner singularities for elliptic problems: Integral equations, graded meshes, quadrature, and compressed inverse preconditioning. *Journal of Computational Physics*, 227(20):8820–8840, Oct. 2008. doi: 10.1016/j.jcp.2008.06.022.

[27] S. Jarvenpää, M. Taskinen, and P. Yla-Oijala. Singularity extraction technique for integral equation methods with higher order basis functions on plane triangles and tetrahedra. *International Journal for Numerical Methods in Engineering*, 58:1149–1165, 2003. doi: 10.1002/nme.810.

[28] C. G. L. Johnson and L. R. Scott. An Analysis of Quadrature Errors in Second-Kind Boundary Integral Methods. *SIAM Journal on Numerical Analysis*, 26(6):1356–1382, 1989. doi: 10.1137/0726079.

[29] S. Kapur and V. Rokhlin. High-Order Corrected Trapezoidal Quadrature Rules for Singular Functions. *SIAM Journal on Numerical Analysis*, 34(4):1331–1356, 1997. doi: 10.1137/S0036142995287847.

[30] M. A. Khayat and D. R. Wilton. Numerical Evaluation of Singular and Near-Singular Potential Integrals. *IEEE Transactions on Antennas and Propagation*, 53(10):3180–3190, 2005. doi: 10.1109/TAP.2005.856342.

[31] A. Klöckner, T. Warburton, and J. Hesthaven. Viscous Shock Capturing in a Time-Explicit Discontinuous Galerkin Method. *Mathematical Modelling of Natural Phenomena*, 6(03):57–83, 2011. doi: 10.1051/mmnp/20116303.

[32] A. Klöckner. Quadrature by Taylor Expansion: A high-order for method for general singular integrals. Technical report, Courant Institute, 2012. in prep.

[33] P. Kolm and V. Rokhlin. Numerical quadratures for singular and hypersingular integrals. *Computers and Mathematics with Applications*, 41(3–4):327–352, 2001. doi: 10.1016/S0898-1221(00)00277-7.

[34] P. Kolm, S. Jiang, and V. Rokhlin. Quadruple and octuple layer potentials in two dimensions I: Analytical apparatus. *Applied and Computational Harmonic Analysis*, 14(1):47–74, Jan. 2003. doi: 10.1016/S1063-5203(03)00004-6.

[35] R. Kress. Boundary integral equations in time-harmonic acoustic scattering. *Mathematical and Computer Modelling*, 15(3–5):229–243, 1991. doi: 10.1016/0895-7177(91)90068-I.

[36] R. Kress. On the numerical solution of a hypersingular integral equation in scattering theory. *Journal of Computational and Applied Mathematics*, 61(3):345–360, 1995. doi: 10.1016/0377-0427(94)00073-7.

[37] R. Kress. *Linear Integral Equations*. Applied Mathematical Sciences, vol. 82, Springer, 1999.

[38] R. Leis. Zur Dirichletschen Randwertaufgabe des Außenraumes der Schwingungsgleichung. *Mathematische Zeitschrift*, 90(3):205–211, 1965. doi: 10.1007/BF01119203.

[39] J. Lowengrub, M. Shelley, and B. Merriman. High-order and efficient methods for the vorticity formulation of the Euler equations. *SIAM Journal on Scientific Computing*, 14:1107–1142, 1993. doi: 10.1137/0914067.

[40] J. Nédélec. *Acoustic and Electromagnetic Equations*. Springer, 1 edition, Mar. 2001. ISBN 0387951555.

[41] O. I. Panič. On the solubility of exterior boundary-value problems for the wave equation and for a system of Maxwell's equations. *Uspehi Mat. Nauk*, 20(1 (121)):221–226, 1965.

[42] Y. Saad and M. H. Schultz. GMRES: A Generalized Minimal Residual Algorithm for Solving Nonsymmetric Linear Systems. *SIAM Journal on Scientific and Statistical Computing*, 7(3):856–869, July 1986. doi: 10.1137/0907058.

[43] C. Schwab and W. L. Wendland. On numerical cubatures of singular surface integrals in boundary element methods. *Numerische Mathematik*, 62:342–369, 1992. doi: 10.1007/BF01396234.

[44] A. Sidi and M. Israeli. Quadrature methods for periodic singular Fredholm integral equations. *Journal of Scientific Computing*, 3:201–231, 1988. doi: 10.1007/BF01061258.

[45] J. Strain. Locally-corrected multidimensional quadrature rules for singular functions. *SIAM Journal on Scientific Computing*, 16(4):992–1017, 1995. doi: 10.1137/0916058.

[46] N. Yarvin and V. Rokhlin. Generalized Gaussian Quadratures and Singular Value Decompositions of Integral Operators. *SIAM Journal on Scientific Computing*, 20(2):699–718, 1998. doi: 10.1137/S1064827596310779.

[47] L. Ying, G. Biros, and D. Zorin. A high-order 3D boundary integral equation solver for elliptic PDEs in smooth domains. *Journal of Computational Physics*, 219:247–275, 2006. doi: 10.1016/j.jcp.2006.03.021.