ADAPTIVE QUADRATURE BY EXPANSION FOR LAYER POTENTIAL EVALUATION IN TWO DIMENSIONS*

LUDVIG AF KLINTEBERG† AND ANNA-KARIN TORNBERG†

Abstract. When solving partial differential equations using boundary integral equation methods, accurate evaluation of singular and nearly singular integrals in layer potentials is crucial. A recent scheme for this is quadrature by expansion (QBX), which solves the problem by locally approximating the potential using a local expansion centered at some distance from the source boundary. In this paper we introduce an extension of the QBX scheme in two dimensions (2D) denoted AQBX—adaptive quadrature by expansion—which combines QBX with an algorithm for automated selection of parameters, based on a target error tolerance. A key component in this algorithm is the ability to accurately estimate the numerical errors in the coefficients of the expansion. Combining previous results for flat panels with a procedure for taking the panel shape into account, we derive such error estimates for arbitrarily shaped boundaries in 2D that are discretized using panel-based Gauss–Legendre quadrature. Applying our scheme to numerical solutions of Dirichlet problems for the Laplace and Helmholtz equations, and also for solving these equations, we find that the scheme is able to satisfy a given target tolerance to within an order of magnitude, making it useful for practical applications. This represents a significant simplification over the original QBX algorithm, in which choosing a good set of parameters can be hard.

Key words. boundary integral equation, adaptive, quadrature, nearly singular, error estimate

AMS subject classifications. 65R20, 65D30, 65D32, 65G99

DOI. 10.1137/17M1121615

1. Introduction. Integral equation methods are a class of numerical methods that are based on the reformulation of an elliptic partial differential equation (PDE) as a boundary integral equation. When applicable, this solution approach has several attractive features. Among these are high-order discretization methods, well-conditioned linear systems after discretization, a reduced number of unknowns compared to volume methods, and straightforward handling of moving boundaries.

A suitable starting point for our discussion on integral equation methods is the representation of the solution u using layer potentials, which are evaluated by integrating the PDE's fundamental solution G and a layer density σ over the domain boundary $\partial\Omega$. We represent u as a linear combination of the double layer potential D and the single layer potential S,

(1)

$$u(z) = D\sigma(z) + \alpha S\sigma(z) = \int_{\partial\Omega} \frac{\partial G(z,w)}{\partial n_w}\, \sigma(w)\, \mathrm{d}s_w + \alpha \int_{\partial\Omega} G(z,w)\sigma(w)\, \mathrm{d}s_w, \quad z \in \Omega,$$

where n_w denotes the unit normal pointing into Ω. For a Dirichlet problem with boundary condition $u = f$, considering the limit $\Omega \ni z \to \partial\Omega$ of (1) leads to a second

---

†Department of Mathematics / Swedish e-Science Research Centre (SeRC), KTH Royal Institute of Technology, 100 44 Stockholm, Sweden (ludvigak@kth.se, akto@kth.se).

kind integral equation in σ,

$$(2) \qquad \left(\frac{1}{2}I + D + \alpha S\right)\sigma(z) = f(z), \quad z \in \partial\Omega.$$

This is a generic form for a Dirichlet problem, and we will later define the forms for the interior Laplace and the exterior Helmholtz equations. The constant α is selected such that the integral equation has a unique solution and is well-conditioned [7]. Nyström discretization of this equation using a suitable quadrature rule generates a dense linear system that can be solved rapidly by exploiting the fact that the off-diagonal blocks are of low rank. Solution methods include the fast multipole method (FMM) [12] and fast direct methods [11, 17, 19].

The main difficulty in the procedure outlined above lies in finding a quadrature rule that can evaluate the layer potentials $S\sigma(z)$ and $D\sigma(z)$ when the target point z is close to or on the boundary $\partial\Omega$. In this case, the integrals of the layer potentials are nearly singular or singular, requiring specialized quadrature methods. For two-dimensional problems, several efficient such methods are available; see the summaries [13] and [14]. Some of the two-dimensional methods can be extended to axisymmetric surfaces in three dimensions (3D) [15]. However, for general surfaces in 3D the available methods are not as mature as those in two dimensions (2D). We refer to [23] for a recent summary of the current state of the art.

1.1. QBX. In this paper, we will focus on the relatively recent method of quadrature by expansion (QBX) [5, 18]. The method provides a way of evaluating both singular and nearly singular integrals by representing the layer potential as a local expansion, centered at a point some distance away from the boundary. While originally proposed for Helmholtz in 2D, it can be generalized to other PDEs in 2D and 3D; see [1, 3, 24] for its successful use in several applications. A strength of QBX is that it uses the same type of local expansions as the FMM, which allows the two methods to be combined into a fast method for evaluating layer potentials at arbitrary locations. This is a topic of ongoing research [22]. As with most methods, however, QBX solves one problem and introduces another. While the problem of evaluating layer potentials on or very close to the boundary is solved, it is replaced with the new problem which is how to efficiently compute the local expansion of the layer potential. In particular, computing the expansion coefficients entails evaluating a series of integrals with increasing order of near singularity. Effectively one has traded a hard problem for several easier problems. Nevertheless, QBX is still a competitive method for these problems, especially since it has a solid analytical groundwork [9].

One of the difficulties of QBX is that of parameter selection. The convergence of the local expansion is governed by the order p at which it is truncated, and by the distance r between the boundary and the expansion center. The expansion coefficients are computed using a quadrature rule for smooth integrands, and for them to be accurate it is necessary to upsample the boundary points by some factor κ. These three parameters together affect the two competing errors of QBX: the truncation error and the coefficient error (often referred to as the quadrature error). The truncation error increases with r and decreases with p, while the coefficient error increases with p and decreases with κ and r (see [1, Fig. 3] for an example). Together r, p, and κ form a large parameter space, and how to best set these parameters is not clear. Instead, experimentation must be used to determine good parameter ranges for a specific application. This is in itself not unfeasible, but it would be preferable to reduce the number of free parameters.

1.2. Contribution. In this paper we propose a scheme for adaptively setting the order p and upsampling factor κ at the time of computation, such that the error is maintained below a target tolerance. The key ingredient for this scheme to be successful is the ability to accurately estimate the magnitude of the coefficient error, which is the quadrature error in the expansion coefficients. Such estimates were derived in [2] for simple geometries in 2D, namely, flat Gauss–Legendre panels and the trapezoidal rule on the unit circle. Here we extend these estimates, by locally using a polynomial to represent the mapping between a flat panel and a panel of general shape. This greatly increases the accuracy of the estimates, and allows us to build an adaptive QBX scheme based on them. We here restrict ourselves to analyzing Gauss–Legendre panel quadrature, but the methodology could equally well be applied to a discretization using the trapezoidal rule.

Taking into account the mapping of the parametrization when analyzing nearly singular quadrature errors is by itself not new; a discussion on nearly singular quadrature errors similar to ours can be found in [5]. Our main contribution in this regard is the construction of an explicit representation of the mapping, which we then invert in order to compute an error estimate.

This paper is organized as follows: In section 2 we introduce the general structure of our scheme, and derive the details for the Laplace double layer potential and the Helmholtz combined field potential. In section 3 we show how to compute the coefficient error estimates necessary for the scheme to be useful. In section 4 we briefly discuss how to combine the scheme with a fast method. In section 5 we show a selection of numerical results that illustrate the performance of our method.

2. Foundations of AQBX. In this section we begin by introducing the foundations of our method using a generic notation, before we give the specific details for the Laplace and Helmholtz equations. We start from a given layer potential representation,

$$(3) \qquad u(z) = \int_{\partial\Omega} G(z,w)\sigma(w)\,\mathrm{d}s_w.$$

To evaluate this using adaptive QBX (AQBX), we first need to split the fundamental solution using a suitable addition theorem,

$$(4) \qquad G(z,w) = \sum_{m=0}^{\infty} A_m^r(w,z_0) \cdot B_m^r(z,z_0),$$

where A_m^r and B_m^r are either vectors or scalars, depending on the fundamental solution. For a given r, the above addition theorem is valid for

$$(5) \qquad |z - z_0| < r \le |w - z_0|.$$

For the specific formulas for Laplace and Helmholtz, see (24) and (35). We assume that (4) is normalized such that the following holds for $|z - z_0| \le r$ and integer $m \ge 0$:

$$(6) \qquad \begin{aligned} \left|B_m^r(z,z_0)\right| &\le 1, \\ \max_{m,z}\left|B_m^r(z,z_0)\right| &= 1. \end{aligned}$$

Here, and throughout the paper, we let $|\cdot|$ denote the ℓ^2-norm for vectors, and the complex modulus for scalars. If the second condition holds for each m, i.e.,

$\max_z |B_m^r(z, z_0)| = 1$, then we get slightly sharper error estimates in Algorithm 1. This is, however, not necessary, and indeed only holds for the Laplace formulation used in this paper. The addition theorem allows us to pick an expansion center z_0, determine r as

$$(7) \qquad r = \min_{w \in \partial\Omega} |w - z_0|,$$

and evaluate the layer potential as a local expansion in the neighborhood of z_0,

$$(8) \qquad u(z) = \sum_{m=0}^{\infty} a_m \cdot B_m^r(z, z_0), \quad |z - z_0| \leq r,$$

where

$$(9) \qquad a_m = \int_{\partial\Omega} A_m^r(w, z_0)\sigma(w)\, \mathrm{d}s_w.$$

The fact that (8) also holds at the equality $|z - z_0| = r$ was shown in [9] for the Laplace and Helmholtz equations, but can be generalized to other kernels, e.g., the Stokes equations [1]. This allows us to also evaluate the expansion on $\partial\Omega$, at the single point which is closest to z_0. In fact, in [5] it is shown that the expansion even converges in a disc about z_0 with radius $R > r$, as long as the density σ is analytic inside that disc.

In a computational scheme the local expansion is truncated at a maximum order p, and the coefficients, which we now denote $\tilde{a}_m$, are computed using a suitable quadrature rule. This gives us the QBX approximation of the layer potential,

$$(10) \qquad u_p(z) = \sum_{m=0}^{p} \tilde{a}_m \cdot B_m^r(z, z_0).$$

The error in this approximation can, assuming exact arithmetic, be separated into a truncation error e_T and a coefficient error e_C [9],

$$(11) \qquad u(z) - u_p(z) = \underbrace{u(z) - \sum_{m=0}^{p} a_m \cdot B_m^r(z, z_0)}_{e_T} + \underbrace{\sum_{m=0}^{p} (a_m - \tilde{a}_m) \cdot B_m^r(z, z_0)}_{e_C}.$$

2.1. Truncation error. The truncation error e_T of the scheme arises because we limit the local expansion to $p + 1$ terms. In our normalized form (6) it can be bounded by

$$(12) \qquad |e_T| = \left| \sum_{m=p+1}^{\infty} a_m \cdot B_m^r(z, z_0) \right| \leq \sum_{m=p+1}^{\infty} |a_m|.$$

It was shown in [9] that the truncation error for several kernels satisfies

$$(13) \qquad e_T \leq M(p, r) r^{p+1} \|u\|_{\mathcal{C}^{p+1}(|z-z_0| \leq r)}$$

for some positive $M(p, r)$. In our experience, the truncation error typically decays exponentially in p, with a rate that is proportional to r. Finding a usable a priori

estimate for the error is hard, since it depends on both the local geometry of $\partial\Omega$ and the regularity of the density σ in a nontrivial way. However, assuming that the expansion coefficients decay exponentially, a good estimate for the truncation error is

$$(14) \qquad |e_T| \approx \left|a_{p+1} \cdot B_{p+1}^r(z, z_0)\right| \leq \left|a_{p+1}\right|.$$

In practice we only have the coefficients up to a_p. Therefore, we can define a useful (and usually conservative) a posteriori estimate as

$$(15) \qquad |e_T| \approx \left|a_p \cdot B_p^r(z, z_0)\right| \leq \left|a_p\right|.$$

2.2. Coefficient error. The coefficient error e_C in (11) is a result of the numerical evaluation of the coefficient integrals (9) for $m = 0, \ldots, p$, which for a given boundary $\partial\Omega$ is evaluated using some n-point quadrature rule. Assuming that the density σ and the boundary $\partial\Omega$ are well-resolved by that quadrature (which is a prerequisite for the underlying Nyström discretization), the main source of the error is the near singularity in A_m^r when evaluated at z_0. The order of this near singularity typically grows with m, and to counter this the density must be upsampled (by interpolation) to a grid which is fine enough to resolve A_m^r. The amount of upsampling required can be determined through $E_C(n, m)$, which is an accurate a priori estimate of the coefficient error in term m when using n quadrature nodes,

$$(16) \qquad E_C(n, m) \approx |a_m - \tilde{a}_m| \geq \left|(a_m - \tilde{a}_m) \cdot B_m^r(z, z_0)\right|.$$

We will in section 3 show how to derive such error estimates for the cases of the Laplace and Helmholtz equations.

Remark 1 (resolution error). Here we only discuss errors in the coefficients $\tilde{a}_m$ due to near singularities in the integrals (9). However, there is also a lower bound on the accuracy of the coefficients, imposed by how well the underlying grid resolves the boundary $\partial\Omega$ and the layer density σ. For a given panel, this error could be estimated by analyzing a modal expansion of the grid point coordinates and density values, such as the Legendre polynomial expansion used in section 3. This is briefly explored in section 5.3, though a more in-depth analysis is beyond the scope of the present paper.

2.3. The AQBX scheme. Let us now assume that we have a discretization of $\partial\Omega$ characterized by n, which for a global quadrature denotes the total number of points, and for a panel-based quadrature denotes the number of points on each panel. Denoting that quadrature Q_n, we define a combined interpolation and quadrature operator $\mathrm{Q}_{\kappa n}$, which computes the quadrature by first upsampling the density to κn points (for simplicity we assume κ integer). Furthermore, we assume that the error when computing a coefficient a_m is well estimated by a function $E_C(\kappa n, m)$. Then the adaptive algorithm for evaluating $u(z)$ in the neighborhood of z_0 to a tolerance ϵ can be summarized using Algorithms 1 and 2. Note that if Algorithm 1 has produced $p + 1$ coefficients, then Algorithm 2 is guaranteed to terminate within $p + 1$ iterations, since $|\delta_p| \leq |a_p| < \epsilon$. For more conservative termination criteria, one can use $\max(|a_{m-1}|, |a_m|) < \epsilon$ and $\max(|\delta_{m-1}|, |\delta_m|) < \epsilon$ to take into account any even/odd behavior in the expansion.

2.4. Specific applications. We will now proceed with formulating AQBX for two different applications: the Laplace double layer potential and the Helmholtz combined field potential.

A1230   LUDVIG AF KLINTEBERG AND ANNA-KARIN TORNBERG

Algorithm 1 Compute expansion coefficients at z_0 to tolerance ϵ.

---
  function EXPANSION COEFFICIENTS(z_0, ϵ)
    $\kappa \leftarrow 1$, $m \leftarrow 0$
    repeat
      while $E_C(\kappa n, m) > \epsilon$ **do**   ▷ Check e_C estimate
        $\kappa \leftarrow \kappa + 1$ ▷ Increase upsampling rate
      end while
      $a_m \leftarrow Q_{\kappa n}[A_m^r(\cdot, z_0)\sigma(\cdot)]$ ▷ Compute coefficient
      $m \leftarrow m + 1$
    until $|a_m| < \epsilon$ ▷ Break when e_T estimate below tolerance
    return $\{a_m\}_{m=0}^p$
  end function

Algorithm 2 Evaluate $u(z)$ to tolerance ϵ using expansion at z_0.

---
  function EVALUATE EXPANSION$(z, z_0, \{a_m\}_{m=0}^p, \epsilon)$
    $u \leftarrow 0$, $m \leftarrow 0$
    repeat
      $\delta_m \leftarrow a_m \cdot B_m^r(z, z_0)$ ▷ Evaluate mth expansion term
      $u \leftarrow u + \delta_m$
      $m \leftarrow m + 1$
    until $|\delta_m| < \epsilon$ ▷ Break when e_T estimate below tolerance
    return u
  end function

2.4.1. Laplace equation. We first consider the Laplace Dirichlet problem in a domain Ω bounded by a boundary $\partial\Omega$,

$$\Delta u = 0 \quad \text{in } \Omega, \tag{17}$$
$$u = f \quad \text{on } \partial\Omega. \tag{18}$$

The fundamental solution to this PDE is

$$\phi(z, w) = \frac{1}{2\pi} \log|z - w|. \tag{19}$$

The interior Dirichlet problem can be represented using the double layer potential,

$$u(z) = D\sigma(z) = \int_{\partial\Omega} \frac{\partial\phi(z, w)}{\partial n_w} \sigma(w)\, ds_w. \tag{20}$$

In complex notation this can be compactly represented as

$$u(z) = \operatorname{Re} v(z), \tag{21}$$
$$v(z) = \frac{1}{2\pi} \int_{\partial\Omega} \frac{n_w}{z - w} \sigma(w)\, ds_w. \tag{22}$$

On $\partial\Omega$ the density σ satisfies the integral equation

$$\left(\frac{1}{2}I + D\right)\sigma = f. \tag{23}$$

It should be noted that this particular layer potential actually has a smooth limit on the boundary, so no special quadrature is needed for solving the integral equation, unless different parts of the boundary are close to each other. Special quadrature is, however, still needed for evaluating the solution close to the boundary, once σ has been computed.

Starting from the Taylor expansion of the Cauchy kernel,

$$(24) \qquad \frac{-1}{z-w} = \sum_{m=0}^{\infty} \frac{(z-z_0)^m}{(w-z_0)^{m+1}},$$

it is straightforward to derive the AQBX formulation (4) for $v(z)$,

$$(25) \qquad A_m^r(w, z_0) = -\frac{r^m n_w}{2\pi(w-z_0)^{m+1}},$$

$$(26) \qquad B_m^r(z, z_0) = \frac{(z-z_0)^m}{r^m}.$$

Note that, by these definitions together with (9), the real part of the coefficient a_0 will simply hold the value of the double layer potential evaluated at z_0, i.e., $\operatorname{Re} a_0 = u(z_0)$.

2.4.2. Helmholtz equation. We now consider the Helmholtz Dirichlet problem in an unbounded domain Ω exterior to a boundary $\partial\Omega$, which for a wavenumber k is stated as

$$(27) \qquad \Delta u + k^2 u = 0 \quad \text{in } \Omega,$$
$$(28) \qquad u = f \quad \text{on } \partial\Omega.$$

The solution must satisfy the Sommerfeld radiation condition for $r = |z|$,

$$(29) \qquad \lim_{r\to\infty} r^{1/2}\left(\frac{\partial u}{\partial r} - iku\right) = 0,$$

which gives a fundamental solution that is essentially the zeroth-order Hankel function of the first kind,

$$(30) \qquad \phi_k(z,w) = \frac{i}{4} H_0^{(1)}(k|z-w|).$$

It is possible to represent the solution using the combined field integral representation,[1]

$$(31) \quad u(z) = \int_{\partial\Omega}\left(\frac{\partial\phi_k(z,w)}{\partial n_w} - \frac{ik}{2}\phi_k(z,w)\right)\sigma(w)\,\mathrm{d}s_w = \left(D_k\sigma - \frac{ik}{2}S_k\sigma\right)(z).$$

Here the double layer kernel is

$$(32) \qquad \frac{\partial\phi_k(z,w)}{\partial n_w} = \frac{ik}{4}H_1^{(1)}(k|z-w|)\frac{(z-w)\cdot n_w}{|z-w|}$$

$$(33) \qquad = \frac{ik}{4}H_1^{(1)}(k|z-w|)|z-w|\operatorname{Re}\left[\frac{n_w}{z-w}\right].$$

---

[1]The general form of the representation is $D_k - i\eta S_k$. We have here set $\eta = k/2$, following [14].

On the boundary we get the integral equation

$$(34) \qquad \left(\frac{1}{2}I + D_k - \frac{ik}{2}S_k \right) \sigma = f.$$

To formulate AQBX for the combined field representation, we start from the Graf addition theorem [20, section 10.23(ii)],

$$(35) \qquad \frac{i}{4}H_0^{(1)}(k|z - w|) = \sum_{m=-\infty}^{\infty} \frac{i}{4}H_m^{(1)}(kr_w)e^{-im\theta_w}J_m(kr_z)e^{im\theta_z}, \quad r_z < r_w.$$

Here (r_w, θ_w) and (r_z, θ_z) are the polar coordinates of $w - z_0$ and $z - z_0$,

$$(36) \qquad \begin{aligned} r_w &= |w - z_0|, & r_z &= |z - z_0|, \\ e^{-im\theta_w} &= \frac{|w - z_0|^m}{(w - z_0)^m}, & e^{im\theta_z} &= \frac{(z - z_0)^m}{|z - z_0|^m}. \end{aligned}$$

We can thus form an expansion for the kernel of the combined field representation (31) as

$$(37) \qquad \frac{\partial \phi_k(z, w)}{\partial n_w} - \frac{ik}{2}\phi_k(z, w) = \sum_{m=-\infty}^{\infty} c_m(w, z_0)J_m(kr_z)e^{im\theta_z},$$

$$(38) \qquad c_m(w, z_0) = d_m(w, z_0) - \frac{ik}{2}s_m(w, z_0).$$

Here s_m is an immediate result of (35),

$$(39) \qquad s_m(w, z_0) = \frac{i}{4}H_m^{(1)}(kr_w)e^{-im\theta_w},$$

while d_m is obtained by differentiation of (39). A compact form for d_m was derived in [5],

$$(40) \qquad d_m(w, z_0) = \frac{ik}{8} \left(H_{m-1}^{(1)}(kr_w)e^{-i(m-1)\theta_w}\bar{n}_w - H_{m+1}^{(1)}(kr_w)e^{-i(m+1)\theta_w}n_w \right).$$

The Bessel functions J_m decay rapidly with m. It is, however, easy to experimentally verify that a normalization satisfying (6) is obtained by using the first term of the power series for $J_m(kr)$ [20, section 10.4, section 10.8],

$$(41) \qquad J_{\pm m}(kr) = (\pm 1)^m \frac{1}{m!}\left(\frac{kr}{2}\right)^m + \mathcal{O}\left(\frac{1}{(m+1)!}\left(\frac{kr}{2}\right)^{m+2}\right), \quad m \geq 0.$$

We have defined our general AQBX formulation (4) for indices $m \geq 0$. To fit the expansion (37) into this we let A_m^r and B_m^r be vector valued for $m > 0$, representing the terms with indices $\pm m$ in (37). The AQBX formulation for the combined field representation is then given by

$$(42) \qquad A_0^r(w, z_0) = c_0(w, z_0),$$
$$(43) \qquad B_0^r(z, z_0) = J_0(kr_z)$$

and, for $m > 0$,

$$(44) \qquad A_m^r(w, z_0) = \frac{\sqrt{2}}{m!} \left(\frac{kr}{2} \right)^m \left(c_m(w, z_0), \; c_{-m}(w, z_0) \right),$$

$$(45) \qquad B_m^r(z, z_0) = \frac{m!}{\sqrt{2}} \left(\frac{2}{kr} \right)^m \left(J_m(kr_z) e^{im\theta_z}, \; J_{-m}(kr_z) e^{-im\theta_z} \right),$$

such that the coefficients $a_m \in \mathbb{C}^2$ for $m > 0$.

We again note that the coefficient a_0 (9) will hold the value of the potential at z_0, since (trivially) $s_0 = \phi_k$ in (30) and (through (33))

$$(46) \qquad d_0(w, z_0) = -\frac{ik}{4} H_1^{(1)}(kr_w) \operatorname{Re} \left[e^{-i\theta_w} n_w \right] = \frac{\partial \phi_k(z_0, w)}{\partial n_w}.$$

3. Coefficient errors. In this section we derive a central piece of AQBX: the coefficient error estimate E_C required in Algorithm 1. We will consider the layer potentials of section 2.4 combined with a panel-based quadrature, where the boundary curve is subdivided into smaller segments, each of which is discretized using an n-point Gauss–Legendre quadrature. This is sometimes referred to as a composite Gauss–Legendre quadrature.

We begin by considering the error in the contribution to a coefficient from a single panel; the total error can then be computed as the sum of errors from all adjacent panels. For this, let Γ be an open curve (i.e., a panel) parametrized by an analytic function $\gamma(t) \in \mathbb{C}$, $t \in [-1, 1]$, with the normal defined as $n(t) = i\gamma'(t)/|\gamma'(t)|$, and γ oriented such that n points into the domain Ω. For simplicity we denote by $\sigma(t)$ the pullback of σ under γ, i.e., $\sigma(t) = \sigma(\gamma(t))$.

3.1. Laplace coefficient error. Beginning with QBX for the Laplace double layer potential (section 2.4.1), the expansion coefficients are computed as

$$(47) \qquad a_m = -\frac{r^m}{2\pi} \int_\Gamma \frac{\sigma(w) n_w}{(w - z_0)^{m+1}} \, \mathrm{d}s_w$$

$$(48) \qquad \qquad = -\frac{ir^m}{2\pi} \int_{-1}^1 \frac{\sigma(t)}{(\gamma(t) - z_0)^{m+1}} \gamma'(t) \, \mathrm{d}t.$$

The coefficient errors are introduced when this integral is evaluated using a discrete quadrature rule. In a boundary integral equation context, we assume a panel Γ such that γ and σ are well-resolved by an n-point Gauss–Legendre quadrature rule. We will here focus on the standard choice $n = 16$.

For our discussion on quadrature errors, we introduce the following notation: Let I denote the integral over $[-1, 1]$, and let Q_n denote the Gauss–Legendre approximation of that integral,

$$(49) \qquad \mathrm{I}[f] = \int_{-1}^1 f(x) \, \mathrm{d}x,$$

$$(50) \qquad \mathrm{Q}_n[f] = \sum_{i=1}^n f(x_i) w_i.$$

The quadrature error R_n is then defined as

$$(51) \qquad \mathrm{R}_n[f] = \mathrm{I}[f] - \mathrm{Q}_n[f],$$

and I, Q_n, and R_n are all linear functionals on $C(-1, 1)$.

If the above assumptions on the resolution hold, then we can expect the quadrature error to be small for the integrand (48), provided that z_0 is far away from Γ. If on the other hand z_0 is close to Γ, then the quadrature error will be dominated by the nearly singular quadrature error that arises when the integrand is evaluated close to its pole.

To estimate nearly singular quadrature errors, we can proceed in the same way as in [2]. The central property which we will use is the following: Let $f(t)$ be a function which is analytic on $[-1, 1]$ and everywhere inside a contour $\mathcal{C}$ enclosing $[-1, 1]$, except at a finite number of poles $\{t_j\}_{j=0}^N$ enclosed by $\mathcal{C}$. If f is integrated on $[-1, 1]$ using the n-point Gauss–Legendre rule, then the quadrature error (51) is given by [8]

$$
(52) \qquad R_n[f] = \frac{1}{2\pi i} \int_{\mathcal{C}} f(t) k_n(t)\, \mathrm{d}t - \sum_{j=0}^{N} \mathrm{Res}\left[f(t)k_n(t), t_j\right],
$$

where k_n is the characteristic remainder function,

$$
(53) \qquad k_n(t) = \frac{1}{P_n(t)} \int_{-1}^{1} \frac{P_n(s)}{t-s}\, \mathrm{d}s.
$$

While not available in closed form, it can in the limit $n \to \infty$ be shown to satisfy [6, App.]

$$
(54) \qquad k_n(t) = \frac{2\pi}{(t \pm \sqrt{t^2-1})^{2n+1}} \left(1 + \mathcal{O}\left(1/n\right)\right),
$$

with the sign in the denominator given by the sign of $\mathrm{Re}\, t$. This result, though asymptotic, is accurate already for small n [2]. Since $n \geq 1$, we can bound the remainder function as

$$
(55) \qquad \left|k_n(t)\right| \leq \frac{2\pi C(t)}{|t \pm \sqrt{t^2-1}|^{2n+1}}
$$

for some $C(t) > 0$. If t lies on the Bernstein ellipse B_ρ, defined as the ellipse with foci ± 1 where the semimajor and semiminor axes sum to $\rho > 1$, then $|t \pm \sqrt{t^2-1}| = \rho$ and

$$
(56) \qquad \left|k_n(t)\right| \leq \frac{2\pi C_\rho}{\rho^{2n+1}},
$$

where $C_\rho = \max_{t \in B_\rho} C(t)$. If $|f| \leq M_\rho$ on some Bernstein ellipse B_ρ, and we let the contour integral in (52) follow that ellipse, then the integral can be bounded as

$$
(57) \qquad \left|\frac{1}{2\pi i} \int_{B_\rho} f(t) k_n(t)\, \mathrm{d}t\right| \leq \frac{C_\rho M_\rho}{\rho^{2n+1}} \int_{B_\rho} |\,\mathrm{d}t| \leq \frac{4 C_\rho M_\rho}{\rho^{2n}}.
$$

The last above step is motivated by the geometry of a Bernstein ellipse; the ratio of the circumference of B_ρ has its maximum as $\rho \to 1$, in which case the ellipse collapses onto $[-1, 1]$. The contribution from the contour integral in (52) vanishes completely if $f(t)k_n(t)t \to 0$ as $|t| \to \infty$, and is generally small compared to the residue for poles close to $[-1, 1]$, as was noted in [2].

In the case where f has a simple pole t_0, we can bound the error as

$$(58) \qquad |R_n[f]| \leq \frac{2\pi C(t_0)}{|t_0 \pm \sqrt{t_0 - 1}|^{2n+1}} \lim_{t \to t_0} |(t - t_0)f(t_0)| + \frac{4C_\rho M_\rho}{\rho^{2n}}.$$

Now consider the more general case, when f has a single pole t_0 of order $m + 1$ (the case of several poles follows trivially). For the residue we can write

$$(59) \qquad \text{Res}\left[f(t)k_n(t), t_0\right] = \frac{1}{m!} \sum_{\ell=0}^{m} \binom{\ell}{m} k_n^{(m-\ell)}(t_0) \lim_{t \to t_0} \left[\frac{d^\ell}{dt^\ell} (t - t_0)^{m+1} f(t)\right],$$

since $k_n(t_0)$ is smooth for $t \notin [-1, 1]$. Defining

$$(60) \qquad \|f\|_{C^m(t_0)} := \max_{\ell \leq m} \lim_{t \to t_0} \left|\frac{d^\ell f}{dt^\ell}\right|,$$

the residue can be bounded as

$$(61) \qquad \left|\text{Res}\left[f(t)k_n(t), t_0\right]\right| \leq \frac{2^m}{m!} \|k_n\|_{C^m(t_0)} \|(t - t_0)^{m+1} f(t)\|_{C^m(t_0)}.$$

This allows us to bound the quadrature error as

$$(62) \qquad |R_n[f]| \leq \frac{2^m}{m!} \|k_n\|_{C^m(t_0)} \|(t - t_0)^{m+1} f(t)\|_{C^m(t_0)} + \frac{4C_\rho M_\rho}{\rho^{2n}}.$$

The above bound will overestimate the error by a large factor. In practice a good approximation of the error is achieved if we neglect the contribution from the contour, under the assumption that the pole is close to the interval, and assume that k_n varies much more rapidly with t than the other factors. Only keeping the term in (59) with the highest derivative in k_n, we get the error approximation

$$(63) \qquad R_n[f] \approx -\frac{1}{m!} k_n^{(m)}(t_0) \lim_{t \to t_0} \left((t - t_0)^{m+1} f(t)\right).$$

This, with the remainder function evaluated using (54), is the approximation that we will be using henceforth. The derivatives of k_n can (at least for small m) be derived analytically from (54) or, as shown in [2], approximated as

$$(64) \qquad k_n^{(m)}(t) \approx k_n(t) \left(\mp \frac{2n + 1}{\sqrt{t^2 - 1}}\right)^m.$$

Returning to the QBX coefficients of the Laplace double layer potential, the integral (47) clearly has a pole of order $m + 1$ at z_0. In parametrized form (48) the pole instead lies at the point $t_0 \in \mathbb{C}$, such that

$$(65) \qquad \gamma(t_0) - z_0 = 0.$$

Note that evaluating $\gamma(t_0)$ corresponds to evaluating an analytic continuation of $\gamma(t)$ in some neighborhood of $[-1, 1]$. Denoting by $\tilde{a}_m$ the approximation of a_m (as defined in (48)), and also assuming that there exists an analytic continuation of σ, then we have from (63) that the error can be approximated as

$$(66) \quad a_m - \tilde{a}_m = R_n\left[-\frac{ir^m \sigma \gamma'}{2\pi(\gamma - z_0)^{m+1}}\right] \approx \frac{ir^m}{2\pi m!} k_n^{(m)}(t_0) \lim_{t \to t_0} \frac{(t - t_0)^{m+1} \sigma(t) \gamma'(t)}{(\gamma(t) - z_0)^{m+1}}.$$

This we can simplify by noting that since $z_0 = \gamma(t_0)$,

$$(67) \qquad \lim_{t \to t_0} \frac{(t - t_0)^{m+1}}{(\gamma(t) - z_0)^{m+1}} = \frac{1}{(\gamma'(t_0))^{m+1}}.$$

Thus a coefficient error estimate $E_C(n, m) \approx |a_m - \tilde{a}_m|$ is given by combining (54), (64), (66), and (67),

$$(68) \qquad E_C(n, m) = \frac{r^m}{m!} \left| \frac{2n+1}{\gamma'(t_0)\sqrt{t_0^2 - 1}} \right|^m \frac{|\sigma(t_0)|}{|t_0 \pm \sqrt{t_0^2 - 1}|^{2n+1}}.$$

This result is essentially a generalization of [2, Thm. 1] to curved panels.

The estimate (68) holds well in the limit $n \to \infty$, but will lose accuracy with increasing m for a fixed n. We have in practice found that it starts losing its reliability around $m > n/2$, so that it then is safer to trigger an upsampling in Algorithm 1, instead of continuing to rely on the estimate.

Though the above derivation is rather straightforward, some more work is required if we want to evaluate (68) in practice, since that requires finding t_0 and being able to evaluate $\sigma(t_0)$ and $\gamma(t_0)$. For this, one would ideally like to have access to analytic expressions for σ and γ, but what we typically have is instead the values of the functions at the quadrature nodes. To be able to evaluate (68) using this pointwise data, we first need to construct continuations of γ and σ. Then we need to solve (65) using the continuation of γ, in order to find t_0 and finally evaluate the estimate. This may sound hard, but can actually be done in an efficient way using polynomial extrapolation.

Beginning with the continuation of γ, let t_i and w_i, $i = 1 \dots n$, be the nodes and weights of the n-point Gauss–Legendre quadrature. We can then let $\mathrm{P}_n[\gamma](t)$ be the polynomial of degree $n-1$ that interpolates $\gamma(t)$ at the nodes t_i. For this high-order ($n = 16$) interpolation to be accurate, we need to compute $\mathrm{P}_n[\gamma]$ in a way that is both well-conditioned and stable [27, Chap. 14]. The distribution of the Legendre points t_i ensures that our interpolation problem is well-conditioned. For stability, we use as our basis the Legendre polynomials $P_\ell(t)$,

$$(69) \qquad \mathrm{P}_n[\gamma](t) = \sum_{\ell=0}^{n-1} \hat{\gamma}_\ell P_\ell(t).$$

These are orthogonal on $[-1, 1]$, allowing us to explicitly compute the coefficients $\hat{\gamma}_\ell$ as

$$(70) \qquad \hat{\gamma}_\ell = \sum_{m=1}^{n} L_{\ell m} \gamma(t_m), \quad \ell = 0, \dots, n-1,$$

where, for a given n, L is a constant matrix with elements

$$(71) \qquad L_{\ell m} = \frac{2\ell + 1}{2} P_\ell(t_m) w_m.$$

To improve convergence in the following step, we assume that Γ has endpoints at -1 and 1; this can be achieved for any open curve by first applying a simple scaling and rotation to both Γ and z_0. Letting the interpolant $\mathrm{P}_n[\gamma](t)$ be the analytic continuation of $\gamma(t)$, we can now find t_0 in (65) by instead solving

$$(72) \qquad \mathrm{P}_n[\gamma](t_0) = z_0,$$

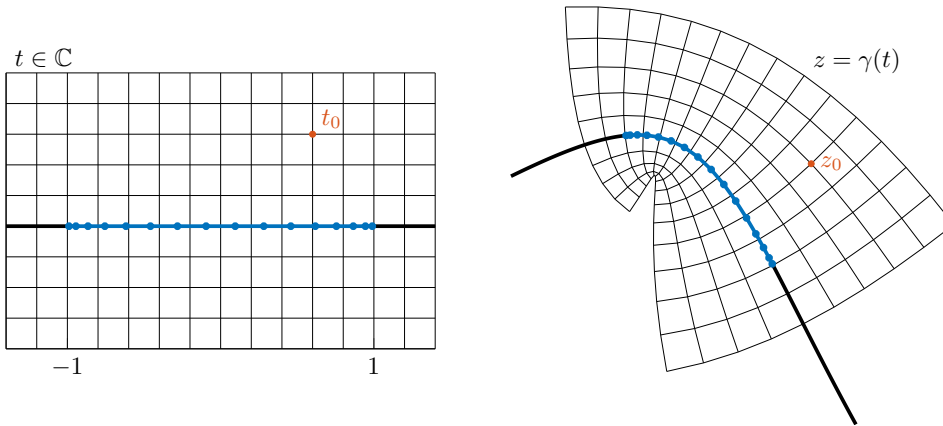Fig. 1. *Illustration of how the analytic continuation of γ maps the vicinity of $[-1,1]$ to the space surrounding Γ. Here Γ is a segment of the starfish domain seen in Figure 3. Polynomial interpolation of γ on a 16-point panel gives a locally very accurate approximation of the continuation: $\left|\gamma - \mathrm{P}_n[\gamma]\right| < 10^{-7}$ on the shown grid.*

using a numerical root-finding algorithm. This does, for the purposes of error estimation, work very well in the near neighborhood of Γ; see example in Figure 1. The Legendre polynomials $P_\ell(t)$ are evaluated using the standard recurrence relation [20, section 14.10.3], which holds also for complex t [20, section 14.21(iii)]. The polynomials grow exponentially in ℓ for $t \notin [-1,1]$, but we have not observed any stability issues related to this. In particular, the coefficients γ_ℓ of a well-resolved panel decay faster than the polynomials $P_\ell(t)$ grow for the range of t-values that are of interest to us.

Equation (72) can be solved efficiently using Newton's method, but the choice of starting guess can be important, especially near concave regions of the curve (e.g., below the curve of Figure 1). In such regions the inverse mapping $t_0 = \gamma^{-1}(z_0)$ is no longer single valued, and for our estimate to be accurate we want the solution t_0 that predicts the largest error, when the estimate is evaluated at that point. This "worst solution" can often be found by solving with $t_0 = \pm 1$ as starting guesses and comparing the results; this strategy was used for the results shown to the right in Figure 2. As an alternative, one can use $t_0 = z_0$ as a starting guess (given the above assumptions on the endpoints of Γ). This simpler strategy works well for practical purposes, and is what we use throughout the remainder of this paper.

Once t_0 is found, we are able to evaluate the coefficient error estimate as given in (68). The value of $\gamma'(t_0)$ can be found by differentiation of the interpolant,

$$(73) \qquad\qquad \gamma'(t_0) = P_n'[\gamma](t_0).$$

We can also evaluate $\sigma(t_0)$ through an interpolating polynomial,

$$(74) \qquad\qquad \sigma(t_0) = P_n[\sigma](t_0).$$

This works very well if σ is well-resolved on Γ, and can be used to obtain the fine-scale correspondence between error and estimate seen in Figure 3. A less expensive alternative is to use the max norm of σ on the interval/panel,

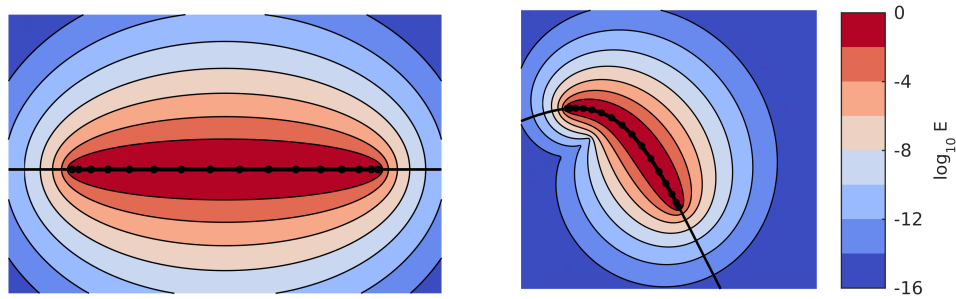$$(75) \qquad\qquad \sigma(t_0) \approx \|\sigma(t)\|_{L^\infty(-1,1)} = \|\sigma(z)\|_{L^\infty(\Gamma)}.$$

FIG. 2. *Quadrature errors when evaluating the integral* (48) *for* $m = 0$ *and* $\sigma = 1$ *on a flat and a curved panel. These panels correspond to those shown in Figure* 1. *Colored fields are actual error levels, contour lines are computed using the error estimate* (68) *with* t_0 *determined by inverting* $P_n[\gamma](t)$. *The ellipses to the left are the Bernstein ellipses.*
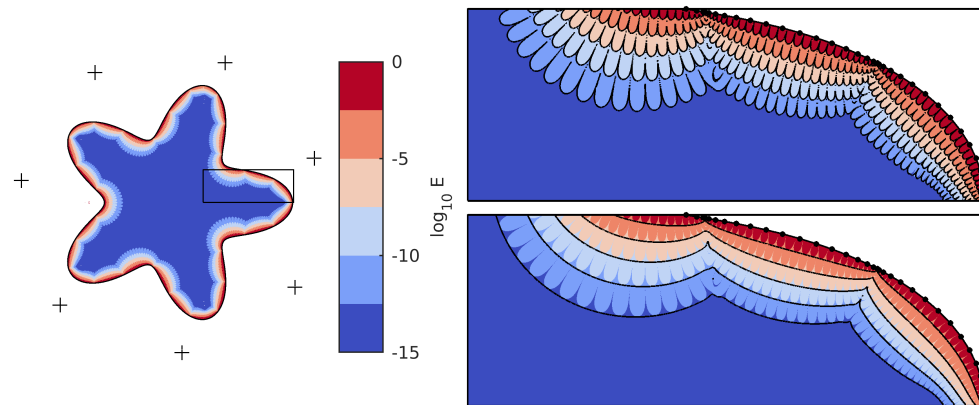

FIG. 3. *Error curves when evaluating the Laplace double layer potential directly using* 27 *panels of equal arc length, with* 16 *points on each panel. Colored fields represent the error compared to the exact solution. Black contour lines in the top right are computed using the estimate* (76) *with* (74), *and coincide almost perfectly with the actual error. Black contour lines in the bottom right are computed without the* $\text{Im}[\cdot]$ *in the estimate* (76), *which produces smooth contours that bound the error curves rather than lie on top of them. This would be enough in most practical applications.*

This works well in practice, as we mainly need to get the order of magnitude right, and also appears to be slightly more robust whenever σ is not fully resolved on Γ, especially for Helmholtz (see Figures 4 and 5). We use this approach in the results of section 5.

The above results can also be used for estimating the nearly singular quadrature error when evaluating the double layer potential (20) near Γ and, hence, to determine when AQBX needs to be used. To that end, we simply use the observation that $u(z_0) = \text{Re}\, a_0$. Denoting by $\tilde{u}(z_0)$ the potential evaluated using direct quadrature, we thus have from (66) that

$$(76) \qquad \left| u(z_0) - \tilde{u}(z_0) \right| \approx \frac{1}{2\pi} \left| \text{Im} \left[\sigma(t_0) k_n(t_0) \right] \right|.$$

This estimate, when evaluated using the above procedure, is very accurate. As an example, in Figure 3, we consider the Laplace Dirichlet problem in a starfish domain. The density is computed by solving (23) using the Nyström method and a right-hand

side given by a collection of point sources (marked by + in the figure). The solution is then evaluated inside the domain directly using the composite Gauss–Legendre quadrature, and compared to the reference solution given by the boundary condition. Comparing these results to those in [2, Fig. 7], which used a flat panel approximation, shows the importance of taking into account the inverse mapping of the target point.

3.2. Helmholtz coefficient error. We now turn to AQBX for the Helmholtz combined field potential (section 2.4.2). When evaluating the expansion coefficients, the Hankel functions in the integrands have a singularity as $r_w \to 0$, which to leading order behaves like [20, section 10.4, section 10.8]

$$(77) \qquad H^{(1)}_{\pm l}(kr_w) = -(\pm 1)^l \frac{2^l (l-1)!}{\pi} (kr_w)^{-l} + \mathcal{O}\left((kr_w)^{-(l-2)} \right).$$

The quadrature error due to near singularity is dominated by that from the highest-order pole, so for the purposes of error estimation it is suitable to only keep the highest-order Hankel function in the expression for (44), and to approximate that Hankel function using only the first term of (77). We can then approximate (39) as

$$(78) \qquad s_m(w, z_0) \approx 0,$$

and (40) as

$$(79) \qquad \begin{aligned} d_0(w, z_0) &= -\frac{ik}{4} H^{(1)}_1(kr_w) r_w \, \mathrm{Re}\left[\frac{n_w}{w - z_0} \right] \\ &\approx \frac{1}{2\pi} \, \mathrm{Re}\left[\frac{n_w}{w - z_0} \right] \end{aligned}$$

and, for $|m| > 0$ and using (36),

$$(80) \qquad d_m(w, z_0) \approx \frac{k}{8} H^{(1)}_{|m|+1}(kr_w) e^{-i(|m|+1)\theta_w} n_w$$

$$(81) \qquad \approx \frac{k}{8\pi} 2^{|m|+1} |m|! (kr_w)^{-(|m|+1)} e^{-i(|m|+1)\theta_w} n_w$$

$$(82) \qquad = \frac{2^{|m|} |m|!}{4\pi k^{|m|}} \frac{n_w}{(w - z_0)^{|m|+1}}.$$

Inserted into (42) and (44), this gives

$$(83) \qquad A^r_0(w, z_0) \approx \frac{1}{2\pi} \, \mathrm{Re}\left[\frac{n_w}{w - z_0} \right],$$

$$(84) \qquad A^r_m(w, z_0) \approx \sqrt{2}\,(1, 1) \frac{r^m}{4\pi} \frac{n_w}{(w - z_0)^{m+1}}, \quad m > 0.$$

Note that in the above series of simplifications (79)–(84) we have removed constant factors which are irrelevant to the magnitude of the error (i.e., -1 and i). Estimating the error of these simplified integrands on a panel with parametrization γ, it follows that

$$(85) \qquad |a_0 - \tilde{a}_0| \approx \left| \mathrm{Im}\, \mathrm{R}_n \left[\frac{\sigma \gamma'}{2\pi(\gamma - z_0)} \right] \right|,$$

$$(86) \qquad |a_m - \tilde{a}_m| \approx \left| \mathrm{R}_n \left[\frac{r^m \sigma \gamma'}{2\pi(\gamma - z_0)^{m+1}} \right] \right|, \quad m > 0.$$
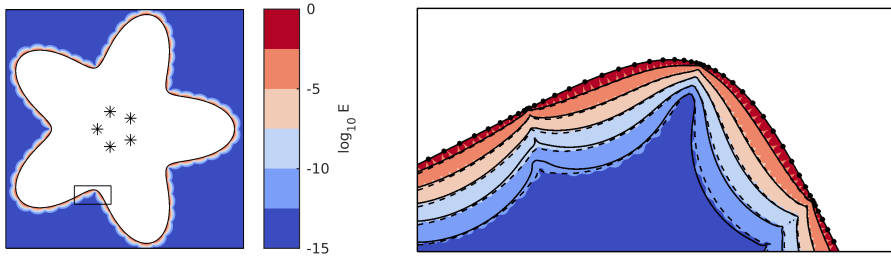
FIG. 4. *Error curves when evaluating the Helmholtz combined field potential directly using* 60 *panels of equal arc length, with* 16 *points on each panel and the point sources located on a circle of radius* 0.2*. The wavenumber is set to* $k = 4/h = 26.6$, h *being the length of the panels. Colored fields represent the error compared to the exact solution, black contours are computed using the estimate* (89)*. The solid lines correspond to* (89) *with* $\sigma(t_0)$ *evaluated using* (74)*, while the dashed lines use the approximation* (75)*.*

Note that (85) is bounded by (86) with $m = 0$, so it is sufficient to use only (86), if we can accept being on the conservative side for $m = 0$. Also note that the integrand in (86) is exactly the same as in the error estimate for the Laplace double layer potential (66). The conclusion is that, remarkably, the coefficient error estimate E_C for Helmholtz is identical to that previously derived for Laplace (68),

$$(87) \qquad E_C(n, m) = \frac{r^m \left|\sigma(t_0)\right|}{2\pi m! \left|\gamma'(t_0)\right|^m} \left|k_n^{(m)}(t_0)\right|.$$

This correspondence between Laplace and Helmholtz also holds for the quadrature error when evaluating the underlying potential itself. This can be seen from (85) and the observation from (46) that $a_0 = u(z_0)$,

$$(88) \qquad \left|u(z_0) - \tilde{u}(z_0)\right| \approx \frac{1}{2\pi} \left|\mathrm{Im}\left[\sigma(t_0)k_n(t_0)\right]\right|$$

$$(89) \qquad \qquad\qquad \leq \frac{1}{2\pi} \left|\sigma(t_0)k_n(t_0)\right|.$$

This estimate has been derived by simplifying the Helmholtz double layer kernel (79), and noting that the singularity to leading order is identical to that of the Laplace double layer. However, in experiments we can observe that the small-scale oscillations predicted by (88) (and which are clearly noticeable in the Laplace case) only appear for small wavenumbers, and it is therefore generally better to use (89). The reason for the disappearance of the oscillations is unknown to us, though an interesting feature is that they only seem to disappear when σ is the solution to (34). Setting $\sigma = 1$ and computing the error by comparing to a finer grid produces the oscillations also for large wavenumbers.

As a demonstration, in Figure 4 we repeat the experiment of Figure 3, but for the Helmholtz exterior Dirichlet problem. We set a number of point sources (marked $*$) inside a starfish domain and solve the integral equation (34) using the discretization scheme of [14]. The correspondence between error and estimate is still very good, though not as excellent as in the Laplace case. Figure 5 shows that the estimate appears to also work for a source point very close to the boundary, when the accuracy of the solution σ has started to deteriorate. Here, when σ is no longer well-resolved, one can clearly see that evaluating $\sigma(t_0)$ using the panel max norm (75) is more stable than using polynomial extrapolation (74). Using (75) results in a useful error estimate
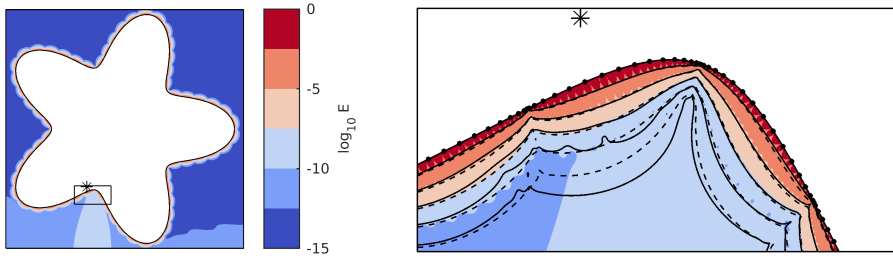
FIG. 5. *Here the discretization and visualization are the same as in Figure* 4, *but the source is now located at a distance* $h/3$ *away from the panel. The error is still well-estimated close to the boundary, where it is dominated by the nearly singular quadrature error, though evaluating* $\sigma(t_0)$ *using the polynomial extrapolation* (74) *appears unstable. Further away from the boundary the error is dominated by a lack of resolution, which our estimate will not capture.*

close to the boundary, where the nearly singular quadrature error dominates. Further away the error is due to σ not being accurate, something which our estimate does not take into account.

It is worth noting that the error estimate (89) is independent of the wavenumber k. This might come as a surprise, as one usually needs to increase the grid resolution with increasing wavenumber. However, here we only take into account the nearly singular quadrature error, under the assumption that far-field interactions are well-resolved. The result simply reflects the fact that the singularity in the kernel is independent of k.

4. Local AQBX. Since QBX is a special quadrature scheme for target points that are close to or on the boundary $\partial\Omega$, it makes sense to only use QBX for those parts of the boundary that are close to a given target point. This is known as "local QBX" [21] (as opposed to "global QBX"), and can be particularly straightforward to combine with a fast method. For panel-based quadrature on a simple curve this is easy to implement; only the panels that are near a given expansion center are used in the local expansion. Selecting panels to include can be done using an error estimate of the layer potential, such as (76) in combination with a tolerance, or by simply including a fixed number of neighboring panels (this works well if all panels are of equal length). When evaluating the potential, the contribution from the near panels is computed through the local expansion, while the contribution from the remaining panels is computed directly using the underlying Gauss–Legendre quadrature.

Let the boundary be composed of a set of panels Γ_i,

$$(90) \qquad \partial\Omega = \bigcup_i \Gamma_i.$$

We can then denote by $\mathcal{N}$ the near panels that are included in the local expansion, and by $\mathcal{F}$ the far panels that are evaluated directly. Note that this division must be made such that the panels in $\mathcal{F}$ are well-separated from all target points at which the expansion will be evaluated. If we write the layer potential as

$$(91) \qquad u^{\partial\Omega}(z) = \int_{\partial\Omega} G(z,w)\sigma(w)\,\mathrm{d}s_w,$$

then the numerical approximation of u using local QBX can be written as

$$(92) \qquad \tilde{u}^{\partial\Omega}(z) = u^{\mathcal{N}}_{\mathrm{QBX}}(z) + u^{\mathcal{F}}_{\mathrm{direct}}(z).$$

To combine this with a fast method that directly evaluates the interactions between

all source points (such as the FMM), one can simply subtract the direct contribution from the near panels,

$$\tilde{u}^{\partial\Omega}(z) = u_{\text{fast}}^{\partial\Omega}(z) - u_{\text{direct}}^{\mathcal{N}}(z) + u_{\text{QBX}}^{\mathcal{N}}(z). \tag{93}$$

The last two terms in this expression can together be viewed as a correction term to the direct quadrature. Since this correction only has local support, the local QBX scheme is FMM compatible, in the sense introduced in [13]. This "black box" way of using the fast method could potentially introduce cancellation errors, when the direct contribution from the near panels is added and subtracted, though we have not observed any such problems in practice. The alternative is to modify the FMM to ignore those local contributions in the first place, which is nontrivial for complex geometries.

A subtle feature of local QBX is that the width of the segment $\mathcal{N}$ affects the convergence rate of the local expansion. This is due to artificially induced endpoint singularities at the interfaces between $\mathcal{N}$ and $\mathcal{F}$, and has been discussed to some extent in both [5] and [21]. The choice of width of $\mathcal{N}$ is therefore a balance; widening $\mathcal{N}$ means more points contributing to each expansion, while narrowing $\mathcal{N}$ gives a slower convergence of the expansion. Actually, not only is the convergence slower, the exponential convergence of the truncation error (13) also tends to be less regular. This in turn makes it harder for AQBX to correctly determine when to terminate. We have found that a good balance is struck by using the five panels that are closest to the expansion center.

5. Numerical experiments. We have implemented the above algorithms in MATLAB, and used the FMM as implemented in FMMLIB2D [10] for fast far-field evaluations. Timings will not be reported here, as our code is a proof of concept rather than a production implementation. We will in the following numerical experiments only report on the Helmholtz problem, as that is the more challenging one. Carrying out the same experiments for the Laplace problem just results in similar, though slightly better, results.

In our numerical experiments we will mainly use AQBX as a postprocessing tool, meaning that we use it to evaluate the layer potential given by a known density. This allows us to study the behavior of AQBX in isolation, without taking into account the method used to compute the density. However, in section 5.4 we show that AQBX can be also used for solving the integral equation.

For our experiments we set up the reference problem shown in Figure 6: The Helmholtz problem in the domain exterior to the starfish curve $\gamma(t) = (1+0.3\cos(10\pi t)) \times e^{-2\pi it}$, $t \in [0,1]$, with Dirichlet boundary condition given by the potential from five point sources in the interior domain. The point sources are randomly positioned on a circle with radius 0.2 centered at the origin, with strengths that are randomly drawn and then normalized such that $\|u\|_\infty = 1$ on $\partial\Omega$. This way all errors reported below are both relative and absolute.

We discretize the boundary using 200 Gauss–Legendre panels of order 16 and equal arc length h, and position one expansion center at a distance $r = h/4$ in the normal direction from each point on the boundary. The wavenumber is set to $k = 2/h = 44.36$. The density σ is computed by solving the integral equation (34) using the Nyström method of [14], which gives us a solution that has a relative error of approximately 10^{-14} when using direct quadrature away from the boundary (measured on a circle of radius 2). Increasing the number of panels or decreasing k does not significantly improve the accuracy of the solution, nor the accuracy of the QBX evaluation.
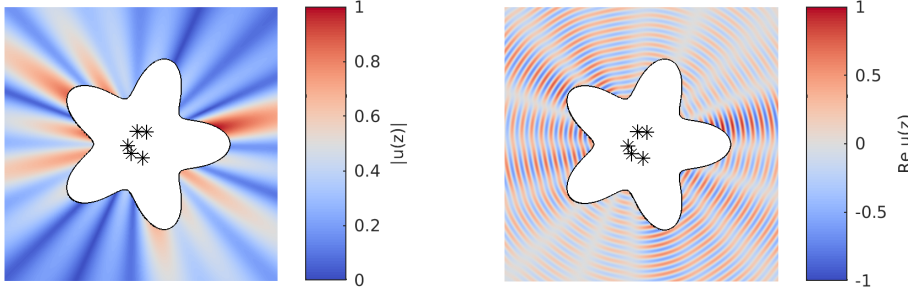
FIG. 6. *Solution to Helmholtz equation* ($k = 44.36$) *given by five point sources located inside a starfish domain.*
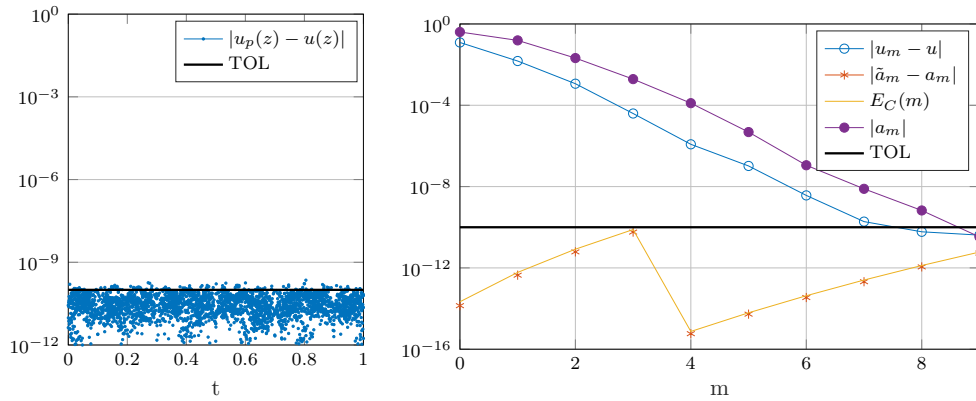

FIG. 7. *Results when evaluating the solution to the Helmholtz equation using AQBX with tolerance set to* 10^{-10}, *marked as thick black line. Left: Distribution of error along* $\partial\Omega$, *which shows that the error stays close to the tolerance, though it is not strictly met. Right: Example showing the behavior of AQBX at a single expansion center, when evaluating at the closest boundary point. The error decays exponentially with expansion order, at the same rate as the coefficients* a_m. *At the same time the coefficient error* $|a_m - \tilde{a}_m|$ *is growing, but is well estimated by the estimate* $E_C(m)$ *(86). Note the jump in coefficient error between* $m = 3$ *and* $m = 4$, *where the grid is upsampled to maintain the error below tolerance.*

Once we have computed σ, we can evaluate the layer potential using AQBX, and compare the result to the exact solution given by the potential from the five point sources. The error in AQBX has, in our tests, always been largest when evaluating the layer potential on the boundary (where the integral is singular), so we mainly report the errors as measured there.

5.1. Performance of the algorithm. We first perform a few experiments that illustrate how AQBX works. Figure 7 shows the error along the entire boundary when evaluating the solution using AQBX and a tolerance of 10^{-10}. One can clearly see how the magnitude of the coefficients a_m provides a good overestimate of the truncation error, while the coefficient error is closely tracked by the estimate E_C. Figures 8 and 9 show example results from when AQBX is used for evaluating the potential in the domain, where the integral is nearly singular. It can be seen that AQBX is only activated at the points where it is needed, and that the potential is then evaluated to the desired accuracy at those points.

The lowest error that we can achieve using QBX (both adaptive and direct) for this problem is around 10^{-12}, which can also be seen in Table 1. This presents a loss
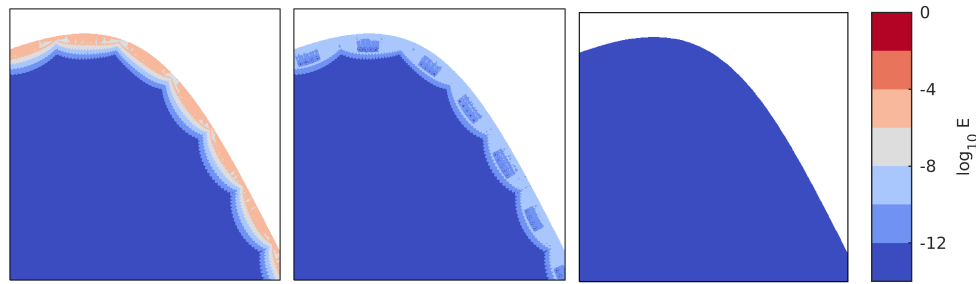
FIG. 8. *Errors when evaluating the potential of Figure 6 on a* 500×500 *grid, in the region highlighted in Figure 4. The AQBX tolerance is set to* 10^{-4}, 10^{-8}, *and* 10^{-12}, *from left to right, and the region where AQBX is activated is determined using the error estimate for the potential* (89).
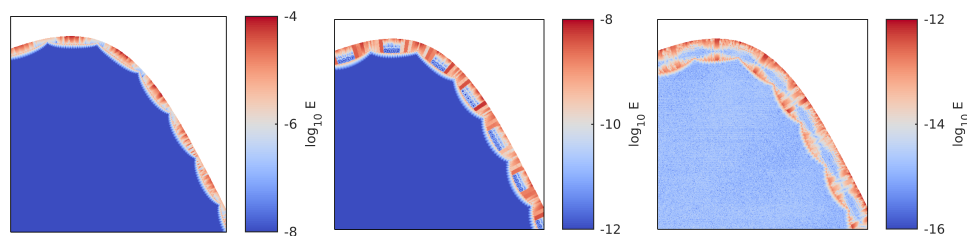

FIG. 9. *Here the same data are plotted as in Figure 8, but with compressed color bars. This makes it possible to see the structure of the remaining errors for each tolerance.*

TABLE 1
Results for varying tolerance and $r/h = 1/4$, *comparing AQBX and direct QBX on our reference problem. Error is measured on* $\partial\Omega$. *Reported parameters are for AQBX an average over all expansion centers (avg), and for direct QBX the optimal fixed values used at all centers (opt). Note that the smallest error that we can obtain is around* 10^{-12}. *This limitation holds also when using direct QBX.*

Tol. ϵ	Eval. error $\|u_p - u\|_\infty$	Exp. order p avg (opt)	Upsamp. rate κ avg (opt)	Work W avg (opt)	Speedup (avg) $W_{\mathrm{QBX}}/W_{\mathrm{AQBX}}$
10^{-4}	$1.4 \cdot 10^{-4}$	5.6 (4)	1.1 (2)	6.0 (8)	1.3
10^{-6}	$1.7 \cdot 10^{-6}$	7.0 (5)	1.5 (2)	10.4 (10)	1.0
10^{-8}	$1.5 \cdot 10^{-8}$	8.8 (7)	1.9 (3)	17.0 (21)	1.2
10^{-10}	$2.2 \cdot 10^{-10}$	10.3 (9)	2.3 (3)	23.2 (27)	1.2
10^{-12}	$2.0 \cdot 10^{-12}$	12.2 (11)	2.6 (4)	32.2 (44)	1.4
10^{-13}	$1.1 \cdot 10^{-12}$	13.3 (11)	2.8 (4)	37.6 (44)	1.2

of two digits of accuracy compared to the error of 10^{-14} achieved using the direct quadrature away from the boundary. While we cannot immediately explain this loss, it is reminiscent of the results in [18, Table 1]. There, the error when using QBX for evaluating the double layer is two orders of magnitudes larger than for the single layer.

5.2. Comparison to direct QBX. We believe that the main benefit of using AQBX rather than a direct QBX implementation is that the parameter choice is greatly simplified; given an expansion distance r and a tolerance ϵ, the upsampling rate κ and expansion order p are set on the fly as needed at each expansion center. A second benefit is that setting κ and p on the fly can save some work, compared to using fixed values everywhere. In an attempt to quantify this, we now introduce a measure of the work (W) needed to form a local expansion, in terms of source evaluations per

Dist. r/h	Eval. error $\|u_p - u\|_\infty$	Exp. order p avg (opt)	Upsamp. rate κ avg (opt)	Work W avg (opt)	Speedup (avg) $W_{\mathrm{QBX}}/W_{\mathrm{AQBX}}$
0.10	$1.7 \cdot 10^{-10}$	8.0 (6)	4.1 (7)	32.8 (42)	1.3
0.25	$2.2 \cdot 10^{-10}$	10.3 (9)	2.3 (3)	23.2 (27)	1.2
0.50	$6.2 \cdot 10^{-9}$	13.8 (12)	1.7 (2)	23.0 (24)	1.0
0.75	$5.4 \cdot 10^{-10}$	17.7 (20)	1.5 (2)	27.1 (40)	1.5
1.00	$1.1 \cdot 10^{-9}$	22.0 (30)	1.8 (2)	40.3 (60)	1.5

original source point. If direct QBX is used with order p and a fixed upsampling rate κ, then the work is given by

$$(94) \qquad W_{\mathrm{QBX}} = p\kappa.$$

If AQBX is used to compute p coefficients, with upsampling rate κ_m used to evaluate the mth coefficient, then the work is given by

$$(95) \qquad W_{\mathrm{AQBX}} = \sum_{m=1}^{p} \kappa_m.$$

As a comparison, in Tables 1 and 2 we measure the work when computing all expansion coefficients in our reference problem using AQBX, and compare that to the work needed if p and κ are fixed everywhere to the minimum values required to achieve the same accuracy as AQBX. These fixed values are tuned by hand to the optimal values for this specific problems, but our algorithm still gives a slight speedup in our definition of work. More importantly, our algorithm is in most cases able to keep the error at the desired order of magnitude without any manual intervention. Hand-tuning parameters is, on the other hand, strictly unfeasible in real applications, and will generally result in an unnecessarily conservative choice of parameters.

Our measure of work does not take into account the extra effort needed to evaluate the estimate of the AQBX scheme. The reported speedup should therefore be viewed as an upper limit compared to the optimal parameter set, and as an indicator that automatic parameter selection does not necessarily have to be more expensive than using fixed parameters. To minimize the overhead of the scheme one can evaluate the error estimates (68) and (87) recursively, and the multiple levels of upsampling can for each panel be computed and reused as needed using a caching algorithm.

5.3. Source points close to the boundary. In order to push the limits of the algorithm, we run a sequence of successively harder test cases. In each case, the Dirichlet boundary condition is given by the potential from a source point that lies at a distance d from the boundary, with d becoming successively smaller. Geometry and results are shown in Figure 10. Setting the AQBX tolerance to 10^{-12}, which is the lowest that we can achieve, we see that the algorithm starts losing accuracy once the source point is within a distance $d = h$ from the boundary. Meanwhile, the error when using direct quadrature, measured away from the boundary, does not start growing until $d = h/2$.

The computation of the QBX coefficients relies on upsampling the density σ on each panel using polynomial interpolation. If the density is not well represented by its interpolant, then that will limit the accuracy of the QBX coefficients. One way of approximating the accuracy of this interpolation is by considering the relative magnitude of the highest-order coefficient in the Legendre expansion (69) of the density on
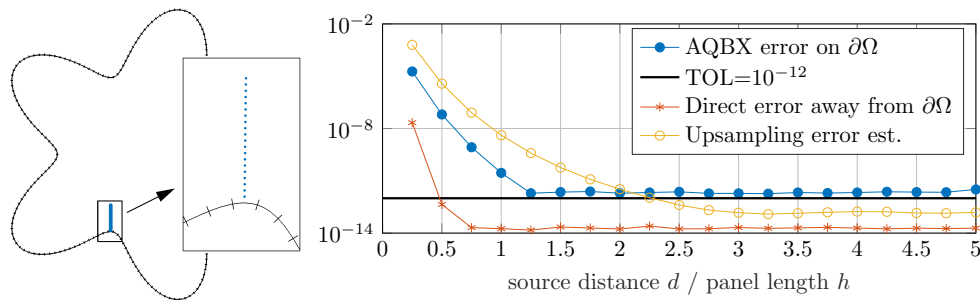
FIG. 10. *Results when using the same discretization as in previous cases (200 panels, $k = 2/h = 44.36$), but with Dirichlet boundary conditions given by a point source at a distance d from the boundary. We run 20 different cases, with the point source in each case given by one of the blue dots in the left graphic. The source strength is for each case normalized such that $\|u\|_\infty = 1$ on $\partial\Omega$. AQBX is evaluated using a tolerance of 10^{-12}. The direct error is the maximum relative error in the solution when evaluated using direct quadrature on a circle of radius 2 (the outer radius of the starfish is 1.3). The upsampling error estimate is the maximum of $|\hat\sigma_{15}|/\|\sigma\|_\infty$ over all panels.*

each panel, $|\hat\sigma_{15}|/\|\sigma\|_\infty$. We take the maximum of this measure over all panels, and denote it the "upsampling error estimate," shown in Figure 10. While not a strictly defined error measure, this quantity gives us some indication of how well the density is resolved on the grid. It is clear from the figure that this error estimate starts growing for source points closer than $3h$ from the boundary. Additionally, it seems that the AQBX error grows at the same rate for $d \leq h$, though we cannot say for certain that this is the mechanism governing the AQBX error.

The results for this particular test case are encouraging, since AQBX does not exhibit any adverse behavior due to the nearby singularities. Instead, the error increase appears to be due to a lack of resolution.

5.4. Solving the integral equation. The above tests indicate that AQBX works well for evaluating layer potentials, both close to and on the boundary. Since the method is accurate on the boundary, it can also be used to solve the underlying integral equation (2). AQBX is then used to evaluate the left-hand side of the discretized integral equation (i.e., the matrix-vector product), and a solution is found iteratively using GMRES [25]. The details of this are for direct QBX discussed to some extent in [18]. There, they recommend that the principal value integral of the double layer is computed using an average of QBX expansions on both sides of the boundary. This was further studied in [23], where they found that two-sided expansions had better convergence properties. We follow this recommendation here.

Table 3 shows the results when solving our test problem for a range of tolerances. We find that the tolerance set for the GMRES iterations is matched by both the error in the solution u and the smoothness of the density σ. However, we find that the AQBX tolerance must be set two orders of magnitude smaller than the GMRES tolerance, otherwise GMRES stagnates. We believe that this can partly be explained by the fact that AQBX does not strictly satisfy its given tolerance. Though beyond the scope of the present paper, it would be interesting to further study the interplay between the tolerances of AQBX and GMRES. In particular, we believe that AQBX could be successfully combined with the inexact GMRES method [26], which is designed to work with a matrix-vector product that has been deliberately made inexact.

6. Conclusions. We have in this paper formulated a scheme for AQBX, which allows for the evaluation of singular and nearly singular layer potential integrals on a

Table 3

Results when solving the problem of Figure 6 using AQBX, with 200 panels. The iteration count is the number of GMRES iterations needed to converge, and the reference value is the count when using the method of [14]. The error in the solution is measured on a circle of radius 2, and the upsampling error estimate is reported as the maximum over all panels.

GMRES tol. ϵ_{GMRES}	AQBX tol. ϵ_{AQBX}	Iter. count AQBX (ref)	Rel. error $\frac{\|u_p - u\|_\infty}{\|u\|_\infty}$	Upsamp. err. est. $\max \frac{|\hat{\sigma}_{15}|}{\|\sigma\|_\infty}$
10^{-2}	10^{-4}	5 (5)	$9.0 \cdot 10^{-03}$	$7.1 \cdot 10^{-03}$
10^{-4}	10^{-6}	11 (11)	$7.1 \cdot 10^{-05}$	$2.7 \cdot 10^{-05}$
10^{-6}	10^{-8}	17 (17)	$5.6 \cdot 10^{-07}$	$3.5 \cdot 10^{-07}$
10^{-8}	10^{-10}	22 (22)	$9.0 \cdot 10^{-09}$	$4.6 \cdot 10^{-09}$
10^{-10}	10^{-12}	28 (28)	$6.4 \cdot 10^{-11}$	$7.8 \cdot 10^{-11}$
10^{-12}	10^{-14}	34 (33)	$3.9 \cdot 10^{-13}$	$7.9 \cdot 10^{-13}$

curve discretized using composite Gauss–Legendre quadrature. The scheme automatically sets parameters in order to satisfy a given tolerance. This is a simplification compared to the original QBX scheme, which has a large parameter space. Given a target tolerance, the only free parameter here is the expansion radius r. Since the remaining parameters are set on the fly, varying r will mainly affect the speed of the algorithm. The optimal value for r with respect to speed will be implementation dependent, though values of W in Table 2 suggest that r/h in the range 0.25–0.50 is a good choice.

The key component of our scheme is the ability to accurately estimate the magnitude of the quadrature errors in the QBX coefficients. To do this we have built on the results of [2], where such estimates were reported for a flat panel. We have extended these to curved panels by taking into account the mapping between a flat panel and a curved panel. This mapping can be locally constructed using only the locations of the quadrature nodes on each panel, and therefore requires no additional analytical information. A side benefit of our estimates is that the nearly singular quadrature error of the underlying layer potential can be accurately estimated, which provides an excellent criterion for when to activate special quadrature also when other methods are used [16, 4]. This could also prove useful for QBX schemes where the expansion is formed by multiple layer potential evaluations in a neighborhood of the expansion center [3, 23].

The focus on target accuracy is, in our experience, uncommon in the context of singular and nearly singular quadrature. We do, however, believe that this is important if integral equation methods are to be used in large-scale simulations, where the focus is on achieving a target accuracy at the lowest possible computational cost.

While several excellent special quadrature methods exist in 2D, methods in 3D have not yet reached the same maturity. The QBX method has been successfully used on simple geometries in 3D [1], while development for more general use is ongoing. If accurate parameter selection is important in 2D, it is absolutely essential in 3D, as costs are higher and the impact of suboptimal parameter choices more severe. The principles of the present 2 dimensional scheme can be extended to 3D. Estimates for the coefficient errors in 3D were developed in [2] for the special case of spheroidal surfaces. Developing estimates for general surfaces in 3D is a topic of ongoing work, and results will be reported at a later date.

REFERENCES

[1] L. af Klinteberg and A.-K. Tornberg, *A fast integral equation method for solid particles in viscous flow using quadrature by expansion*, J. Comput. Phys., 326 (2016), pp. 420–445, https://doi.org/10.1016/j.jcp.2016.09.006.

[2] L. af Klinteberg and A.-K. Tornberg, *Error estimation for quadrature by expansion in layer potential evaluation*, Adv. Comput. Math., 43 (2017), pp. 195–234, https://doi.org/10.1007/s10444-016-9484-x.

[3] T. Askham and A. Cerfon, *An adaptive fast multipole accelerated Poisson solver for complex geometries*, J. Comput. Phys., 344 (2017), pp. 1–22, https://doi.org/10.1016/j.jcp.2017.04.063.

[4] A. Barnett, B. Wu, and 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 (2015), pp. B519–B542, https://doi.org/10.1137/140990826.

[5] A. H. Barnett, *Evaluation of layer potentials close to the boundary for Laplace and Helmholtz problems on analytic planar domains*, SIAM J. Sci. Comput., 36 (2014), pp. A427–A451, https://doi.org/10.1137/120900253.

[6] W. Barrett, *Convergence properties of Gaussian quadrature formulae*, Comput. J., 3 (1961), pp. 272–277, https://doi.org/10.1093/comjnl/3.4.272.

[7] T. Betcke, S. Chandler-Wilde, I. Graham, S. Langdon, and M. Lindner, *Condition number estimates for combined potential integral operators in acoustics and their boundary element discretisation*, Numer. Methods Partial Differential Equations, 27 (2011), pp. 31–69, https://doi.org/10.1002/num.20643.

[8] J. D. Donaldson and D. Elliott, *A unified approach to quadrature rules with asymptotic estimates of their remainders*, SIAM J. Numer. Anal., 9 (1972), pp. 573–602, https://doi.org/10.1137/0709051.

[9] C. L. Epstein, L. Greengard, and A. Klöckner, *On the convergence of local expansions of layer potentials*, SIAM J. Numer. Anal., 51 (2013), pp. 2660–2679, https://doi.org/10.1137/120902859.

[10] Z. Gimbutas and L. Greengard, *FMMLIB2D*, http://www.cims.nyu.edu/cmcl/fmm2dlib/fmm2dlib.html (2012).

[11] L. Greengard, D. Gueyffier, P.-G. Martinsson, and V. Rokhlin, *Fast direct solvers for integral equations in complex three-dimensional domains*, Acta Numer., 18 (2009), pp. 243–275, https://doi.org/10.1017/S0962492906410011.

[12] L. Greengard and V. Rokhlin, *A new version of the fast multipole method for the Laplace equation in three dimensions*, Acta Numer., 6 (1997), pp. 229–269, https://doi.org/10.1017/S0962492900002725.

[13] S. Hao, A. H. Barnett, P. G. Martinsson, and P. Young, *High-order accurate methods for Nyström discretization of integral equations on smooth curves in the plane*, Adv. Comput. Math., 40 (2014), pp. 245–272, https://doi.org/10.1007/s10444-013-9306-3.

[14] J. Helsing and A. Holst, *Variants of an explicit kernel-split panel-based Nyström discretization scheme for Helmholtz boundary value problems*, Adv. Comput. Math., 41 (2015), pp. 691–708, https://doi.org/10.1007/s10444-014-9383-y.

[15] J. Helsing and A. Karlsson, *An explicit kernel-split panel-based Nyström scheme for integral equations on axially symmetric surfaces*, J. Comput. Phys., 272 (2014), pp. 686–703, https://doi.org/10.1016/j.jcp.2014.04.053.

[16] J. Helsing and R. Ojala, *On the evaluation of layer potentials close to their sources*, J. Comput. Phys., 227 (2008), pp. 2899–2921, https://doi.org/10.1016/j.jcp.2007.11.024.

[17] K. L. Ho and L. Greengard, *A fast direct solver for structured linear systems by recursive skeletonization*, SIAM J. Sci. Comput., 34 (2012), pp. A2507–A2532, https://doi.org/10.1137/120866683.

[18] A. Klöckner, A. Barnett, L. Greengard, and M. O'Neil, *Quadrature by expansion: A new method for the evaluation of layer potentials*, J. Comput. Phys., 252 (2013), pp. 332–349, https://doi.org/10.1016/j.jcp.2013.06.027.

[19] P. Martinsson and V. Rokhlin, *A fast direct solver for boundary integral equations in two dimensions*, J. Comput. Phys., 205 (2005), pp. 1–23, https://doi.org/10.1016/j.jcp.2004.10.033.

[20] NIST, *Digital Library of Mathematical Functions*, Release 1.0.16, http://dlmf.nist.gov/ (2017).

[21] M. Rachh, *Integral Equation Methods for Problems in Electrostatics, Elastostatics and Viscous Flow*, Ph.D. thesis, New York University, New York, 2015.

[22] M. Rachh, A. Klöckner, and M. O'Neil, *Fast algorithms for quadrature by expansion I: Globally valid expansions*, J. Comput. Phys., 345 (2017), pp. 706–731, https://doi.org/10.1016/j.jcp.2017.04.062.

[23] A. RAHIMIAN, A. BARNETT, AND D. ZORIN, *Ubiquitous evaluation of layer potentials using quadrature by kernel-independent expansion*, BIT, (2017), https://doi.org/10.1007/s10543-017-0689-2.

[24] L. RICKETSON, A. CERFON, M. RACHH, AND J. FREIDBERG, *Accurate derivative evaluation for any Grad-Shafranov solver*, J. Comput. Phys., 305 (2016), pp. 744–757, https://doi.org/10.1016/j.jcp.2015.11.015.

[25] Y. SAAD AND M. H. SCHULTZ, *GMRES: A generalized minimal residual algorithm for solving nonsymmetric linear systems*, SIAM J. Sci. Stat. Comput., 7 (1986), pp. 856–869, https://doi.org/10.1137/0907058.

[26] V. SIMONCINI AND D. B. SZYLD, *Theory of inexact Krylov subspace methods and applications to scientific computing*, SIAM J. Sci. Comput., 25 (2003), pp. 454–477, https://doi.org/10.1137/S1064827502406415.

[27] L. N. TREFETHEN, *Approximation Theory and Approximation Practice*, Other Titles Appl. Math. 128. SIAM, Philadelphia, 2012.