EVALUATION OF LAYER POTENTIALS CLOSE TO THE BOUNDARY FOR LAPLACE AND HELMHOLTZ PROBLEMS ON ANALYTIC PLANAR DOMAINS

ALEX H. BARNETT*

Abstract. Boundary integral equations are an efficient and accurate tool for the numerical solution of elliptic boundary value problems. The solution is expressed as a layer potential; however, the error in its evaluation grows large near the boundary if a fixed quadrature rule is used. Firstly, we analyze this error for Laplace's equation with analytic density and the global periodic trapezoid rule, and find an intimate connection to the complexification of the boundary parametrization. Our main result is then a simple and efficient scheme for accurate evaluation up to the boundary for single- and double-layer potentials for the Laplace and Helmholtz equations, using surrogate local expansions about centers placed near the boundary. The scheme—which also underlies the recent QBX Nyström quadrature—is asymptotically exponentially convergent (we prove this in the analytic Laplace case), requires no adaptivity, generalizes simply to three dimensions, and has $O(N)$ complexity when executed via a locally-corrected fast multipole sum. We give an example of high-frequency scattering from an obstacle with perimeter 700 wavelengths long, evaluating the solution at 2×10^5 points near the boundary with 11-digit accuracy in 30 seconds in MATLAB on a single CPU core.

Key words. potential theory, layer potential, integral equation, Laplace equation, Helmholtz equation, close evaluation

1. Introduction.

We are interested in solving boundary-value problems (BVPs) of the type

$$(\Delta + \omega^2)u = 0 \qquad \text{in } \Omega \tag{1.1}$$

$$u = f \qquad \text{on } \partial\Omega \tag{1.2}$$

where $\Omega \subset \mathbb{R}^2$ is either an interior or exterior domain with boundary curve $\partial\Omega$, and either $\omega = 0$ (Laplace equation) or $\omega > 0$ (Helmholtz equation). We will mostly use the above Dirichlet boundary condition in our examples, and note that Neumann and other types of boundary conditions can equally well benefit from our technique. The numerical solution of this type of BVP has numerous applications in electrostatics, equilibrium problems, and acoustic or electromagnetic wave scattering in the frequency domain. The case $f \equiv 0$ includes eigenvalue (cavity resonance) problems for the Laplacian.

The boundary integral approach [1, 26] has many advantages over conventional finite element or finite difference discretization of the domain: very few unknowns are needed since the problem is now of lower dimension, provable high-order accuracy is simple to achieve, and, in exterior domains, radiation conditions are automatically enforced without the use of artificial boundaries. Thus the approach is especially useful for wave scattering, including at high frequency [11, 10].

The integral equation approach exploits the known fundamental solution for the PDE,

$$\Phi(x,y) = \begin{cases} \frac{1}{2\pi}\log\frac{1}{|x-y|}, & \omega = 0, \\ \frac{i}{4}H_0^{(1)}(\omega|x-y|), & \omega > 0. \end{cases} \tag{1.3}$$

For the interior case, which is the simplest, the BVP (1.1)–(1.2) is converted to a boundary integral equation (BIE), which is of the Fredholm second kind,

$$(D - \tfrac{1}{2}I)\tau = f \ , \tag{1.4}$$

where τ is an unknown density function on $\partial\Omega$, I is the identity, $D : C(\partial\Omega) \to C(\partial\Omega)$ is the double-layer integral operator with kernel $k(x,y) = \partial\Phi(x,y)/\partial n(y)$, and $n(y)$ is the outward

*Department of Mathematics, Dartmouth College, Hanover, NH, 03755, USA

normal at $y \in \partial\Omega$. Numerical solution of (1.4), for instance via the Nyström method [26, Ch. 12] [1, Ch. 4] with N quadrature nodes on $\partial\Omega$, results in an approximation to τ sampled at these nodes, from which one may recover an approximation to τ on the whole of $\partial\Omega$ by interpolation (e.g. Nyström interpolation). Finally, one can evaluate the approximate BVP solution at any target point x in the domain as the double-layer potential

$$u(x) = \int_{\partial\Omega} \frac{\partial\Phi(x,y)}{\partial n(y)} \tau(y) ds_y \ , \qquad x \in \Omega \ . \tag{1.5}$$

It is convenient, and common practice, to reuse the *existing* N quadrature nodes underlying the Nyström method to approximate the integral (1.5)—in other words, to skip the interpolation step; we will call this the *native* evaluation scheme.[1] It is common wisdom that this gives an accurate solution when x is "far" from $\partial\Omega$, but a very inaccurate one close to $\partial\Omega$, even when the τ samples themselves are accurate. Fig. 1.1 (a), whose content will be familiar to anyone who has tested the accuracy of a BIE method, illustrates this: the error of evaluation grows to $O(1)$ as one nears $\partial\Omega$. (Also see [20, Fig. 2] or [23, Fig. 4] which show a similar story for a panel-based underlying quadrature.)

There are several ways to overcome this problem near the boundary. In order of increasing sophistication, these include: i) increase N in the Nyström method (although solving a linear system larger than one needs is clearly a waste of resources); ii) fix N, but then interpolate τ onto a finer fixed set of boundary nodes, enabling points closer to $\partial\Omega$ to be accurately evaluated [1, Fig. 7.4] (this is implemented in [2]—what is not discussed is that the number of finer nodes must grow without limit as $x \to \partial\Omega$); iii) use an adaptive quadrature scheme for (1.5) which is able to access the interpolant of τ [16] (although this achieves high accuracy for each point $x \in \Omega$, we have found it very slow [5, Sec. 5.1] because the adaptivity depends on x, with an arbitrarily large number of refinements needed as $x \to \partial\Omega$); iv) use various fixed-order methods based upon precomputed quadratures [28] or grids [27] for the Laplace case; or v) use high-order methods of Helsing–Ojala [20] for the Laplace case, which approach machine precision accuracy while maintaining efficiency within a fast multipole (FMM) accelerated scheme.

The method of [20] has recently been extended to the Helmholtz equation in two dimensions [19]. Here we present an alternative, simple, and efficient new method that addresses the close-to-boundary quadrature problem, with numerical effort independent of the distance of x from $\partial\Omega$. One advantage is that our scheme extends naturally to the three-dimensional case, unlike existing high-order two-dimensional schemes. Incidentally, our scheme equally well evaluates the potential *on the curve* $\partial\Omega$ (i.e. the limit of x approaching $\partial\Omega$ from one side), meaning that it can also be used to construct high-order Nyström quadratures to solve the BIE (1.4) itself; the resulting tool is called QBX [23].

Firstly, in section 2 we analyze (in Theorems 3 and 9) the evaluation error for Laplace double-layer and single-layer potentials, with the global periodic trapezoid rule on an analytic curve $\partial\Omega$ with analytic data—we are surprised not to find these results in the literature. This is crucial in order to determine the neighborhood of $\partial\Omega$ in which native evaluation is poor; it is within this "bad neighborhood" that the new close-evaluation scheme is used. We present the scheme in section 3: simply put, the idea is to interpolate τ to a fixed finer set of roughly $4N$ nodes, from which one computes (via the addition theorem) the coefficients of *local expansions* (i.e., Taylor expansions in the Laplace case, Fourier–Bessel expansions in the Helmholtz), around a set of expansion centers placed near (but not too near) $\partial\Omega$. It is these "surrogate" local expansions that are then evaluated at nearby desired target points x. This is reminiscent of the method of Schwab–Wendland [32] but with the major difference that expansion centers lie off

[1] In [23] this is called the "underlying" scheme, and in [20] "straight-up" quadrature.

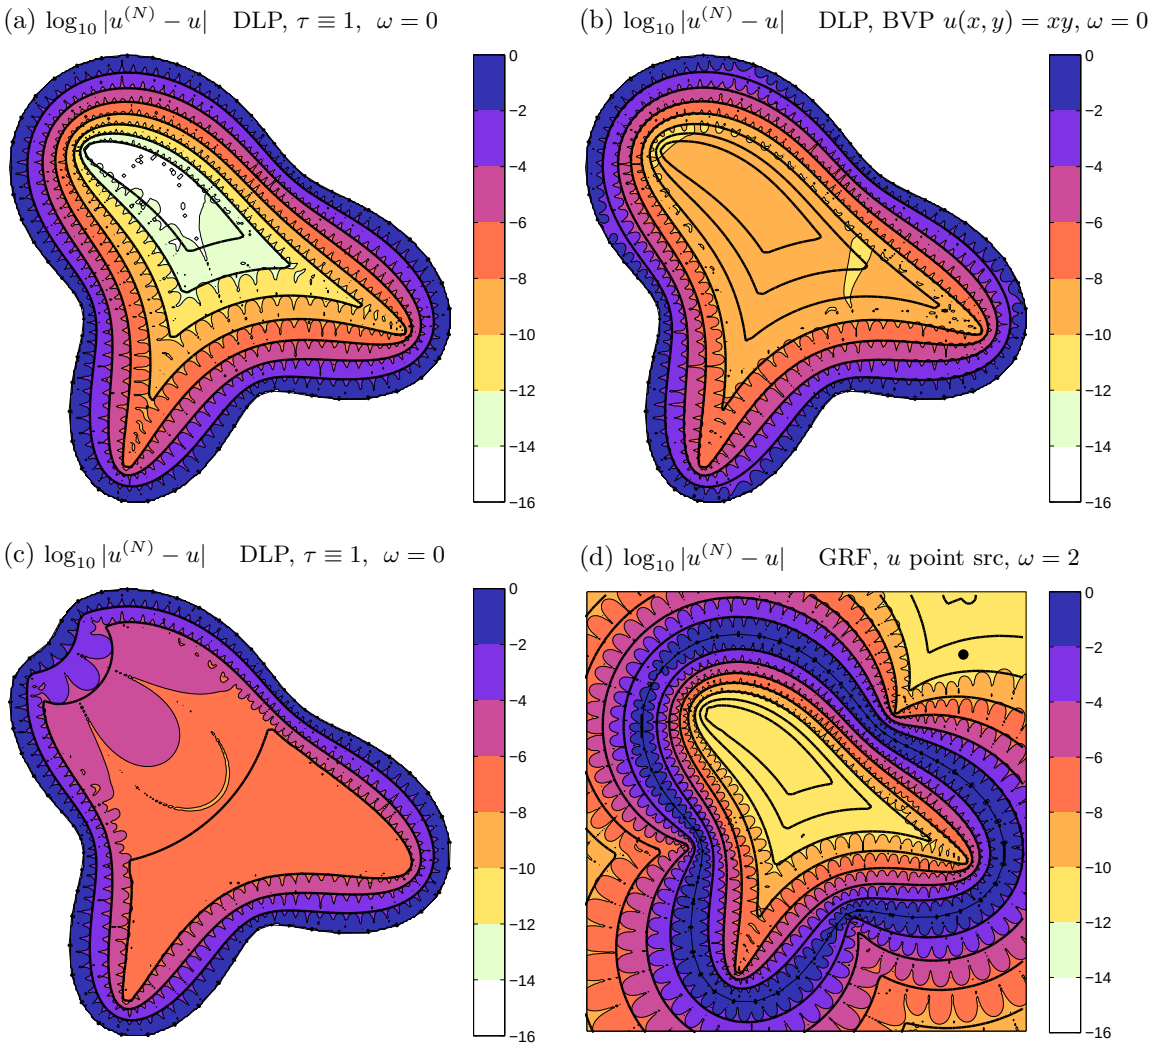

FIG. 1.1. *Evaluation error for Laplace and Helmholtz layer potentials using the native quadrature scheme described after* (1.5). *In all four cases there are* $N = 60$ *nodes. Error contours separated by a factor of* 10^2 *are shown (thin black lines), as are the contours (thick black lines) predicted by Theorem 3. In (a), (b), and (d) the interior domain* Ω *has boundary given by the polar function* $r(\theta) = 1 + 0.3\cos[3(\theta + 0.3\sin\theta)]$. *(a) is a double-layer with* $\tau \equiv 1$ *and* $\omega = 0$. *(b)* τ *is the Nyström solution for Dirichlet data corresponding to the potential* $u(x, y) = xy$, *with* $\omega = 0$. *(c) same as (a) but* $\partial\Omega$ *has a small Gaussian "bump" on its northwest side. (d) Test of Green's representation formula* (2.12), *both inside and outside, for* $\omega = 2$ *(roughly 0.8 wavelength across the diameter) with data due to an exterior point source at the dot shown.*

of, rather than on, $\partial\Omega$. In section 4.2 we show how to combine the new scheme with the FMM to achieve an overall $O(N)$ complexity for the evaluation of $O(N)$ target points lying in the bad neighborhood, and apply this to high-frequency scattering from a smooth but complicated obstacle 100 wavelengths across. Finally, we conclude and mention future directions in section 5.

2. Theory of Laplace layer potential evaluation error using the global trapezoid rule. Our goal in this section is to analyze rigorously the native (i.e., N-node) evaluation error for analytic single- and double-layer potentials for the Laplace equation ($\omega = 0$), on analytic curves. An example plot of such error varying over an interior domain is shown in Fig. 1.1(a).

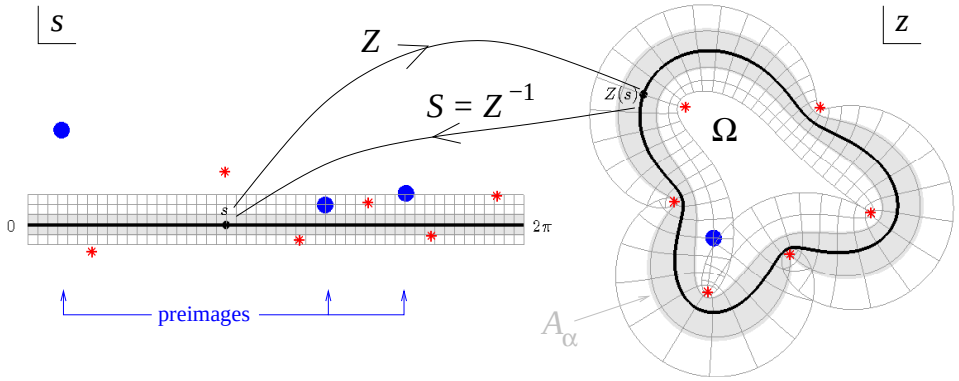

FIG. 2.1. *Parametrization of the boundary of the interior domain Ω from Fig. 1.1, as a map from s-plane (shown left with a square grid) to z-plane (shown right). One point on the real s axis and its image on the boundary are shown by small black dots. A single point in the interior and its three pre-images in the s-plane are shown by large blue dots. An annular neighborhood A_α in which Z is analytic and invertible, and its preimage, are shown in grey. The nearest six (branch-type) singularities of the Schwarz function of the domain, and their pre-images, are shown by $*$.*

2.1. Geometric preliminaries. We identify $\mathbb{R}^2$ with $\mathbb{C}$, and let the simple analytic closed curve $\partial\Omega$ define either a bounded interior, or unbounded exterior, open domain $\Omega \subset \mathbb{C}$. We need $Z : \mathbb{R} \to \mathbb{C}$ as an analytic 2π-periodic counter-clockwise parametrization of $\partial\Omega$, i.e. $Z([0, 2\pi)) = \partial\Omega$. This means $Z(s) = z_1(s) + iz_2(s)$, with z_1 and z_2 real analytic and 2π-periodic, and that Z may be continued as an analytic function in some neighborhood of the real axis. We assume that the *speed function* $|Z'(s)|$ is positive for all real s. These conditions means that Z is analytic and invertible in a strip $|\operatorname{Im} s| < \alpha$, for some $\alpha > 0$. The image of this strip under Z defines an annular (tubular) neighborhood of $\partial\Omega$ in which Z^{-1} is also analytic. Fig. 2.1 shows such a strip and its image, both shaded in grey. Also shown are the singularities that control its width: at these points $Z' = 0$, so that locally Z^{-1} takes the form of a (translated) square-root map, hence must have branch cut and cease to be single-valued. These points are also singularities of the so-called *Schwarz function* of the domain; see [13, Ch. 5, 6, 8] and [24]. Other types of Schwarz singularities are possible for domains with analytic boundaries (e.g. see [3]); however, for our analysis the type of singularity is irrelevant.

For $\alpha \in \mathbb{R}$, we will use the notation Γ_α to mean a translation of $\partial\Omega$ by α in the imaginary parameter direction,

$$\Gamma_\alpha \;:=\; Z(\{s = t + i\alpha : t \in \mathbb{R}\}) \;.$$

In particular, $\Gamma_0 = \partial\Omega$. Note that for all sufficiently small α, Γ_α is a Jordan curve, but that for larger $|\alpha|$, it will in general start to self-intersect. This is illustrated by the images of the grid-lines in Fig. 2.1. For $\alpha > 0$ we define A_α by the open annular neighborhood of $\partial\Omega$,

$$A_\alpha \;:=\; Z(\{s = t + ia : t, a \in \mathbb{R}, |a| < \alpha\}) \;,$$

i.e. the image of the strip $|\operatorname{Im} s| < \alpha$. Note that when $\Gamma_{-\alpha}$ and Γ_α do not intersect themselves or each other, then A_α is simply the open region lying between them.

2.2. Evaluation error in the double-layer case. Given a real-valued analytic density $\tau \in C(\partial\Omega)$, the Laplace double-layer potential (1.5) may be written $u = \operatorname{Re} v$, where v is the

function defined by the complex contour integral[2]

$$v(z) \;=\; \frac{-1}{2\pi i}\int_{\partial\Omega}\frac{\tau(y)}{y-z}dy\;, \qquad z\in\mathbb{C}\backslash\partial\Omega\;. \tag{2.1}$$

Thus v is analytic in Ω, and also in $\mathbb{C}\backslash\overline{\Omega}$. Since the evaluation error of u is bounded by that of v, we shall work with v from now on. Let $\tilde{\tau}$ be the pullback of τ under Z, i.e. $\tilde{\tau}(s)=\tau(Z(s))$ for all $s\in\mathbb{R}$. Rewriting (2.1) in terms of the parameter gives,

$$v(z) \;=\; \frac{-1}{2\pi i}\int_0^{2\pi}\frac{\tilde{\tau}(s)}{Z(s)-z}Z'(s)ds\;, \qquad z\in\mathbb{C}\backslash\partial\Omega\;. \tag{2.2}$$

For quadrature of (2.2) we now choose the global periodic trapezoid rule along the real s axis, introducing nodes $2\pi j/N$, $j=1,\ldots,N$, and equal weights $2\pi/N$, thus

$$v^{(N)}(z) \;:=\; \frac{-1}{iN}\sum_{j=1}^{N}\frac{\tilde{\tau}(2\pi j/N)}{Z(2\pi j/N)-z}Z'(2\pi j/N)\;. \tag{2.3}$$

The integrand in (2.2) is analytic, implying exponential convergence of (2.3) by the following classical theorem [12] (see e.g. [26, Thm. 12.6]).

THEOREM 1 (Davis). *Let f be 2π-periodic and analytic in the strip $|\operatorname{Im} s|\le\alpha$ for some $\alpha>0$, and let $|f|\le F$ in this strip. Then the quadrature error of the periodic trapezoid rule,*

$$E_N \;:=\; \frac{2\pi}{N}\sum_{j=1}^{N}f(2\pi j/N) \;-\int_0^{2\pi}f(s)ds\;, \tag{2.4}$$

obeys the bound

$$|E_N| \;\le\; \frac{4\pi F}{e^{\alpha N}-1}\;. \tag{2.5}$$

However, to achieve (rather than merely approach) the correct convergence rate, we will need the following generalization (similar to that of Hunter [22]):

LEMMA 2. *Let f be 2π-periodic and meromorphic in the strip $|\operatorname{Im} s|\le\alpha$ for some $\alpha>0$, with only one simple pole in this strip, at s_0, with $\operatorname{Im} s_0\neq 0$. Let f have residue r_0 at this pole, and let $|f|\le F$ on the edges of the strip, i.e. for all s with $|\operatorname{Im} s|=\alpha$. Then the quadrature error* (2.4) *obeys the bound*

$$|E_N| \;\le\; \frac{2\pi|r_0|}{e^{|\operatorname{Im} s_0|N}-1}+\frac{4\pi F}{e^{\alpha N}-1}\;. \tag{2.6}$$

Note that the first term dominates as N grows, and that $r_0=0$ recovers the Davis theorem.

Proof. Let Γ_1 and Γ_2 be the upper and lower strip boundaries respectively, both traversed with increasing real part. For the sum in (2.4) we apply the residue theorem to $\cot\frac{Ns}{2}f(s)$ in the strip, noticing that the vertical sides cancel due to periodicity. For the integral in (2.4) we apply the residue theorem to $f(s)$ in each of the upper and lower semi-strips, take their average. Combining these, (2.4) can be rewritten

$$E_N=\int_{\Gamma_1}\left(\frac{i}{2}\cot\frac{Ns}{2}-\frac{1}{2}\right)f(s)ds-\int_{\Gamma_2}\left(\frac{i}{2}\cot\frac{Ns}{2}+\frac{1}{2}\right)f(s)ds+2\pi ir_0\left(\frac{i}{2}\cot\frac{Ns_0}{2}\mp\frac{1}{2}\right)\;,$$

[2]Note that this Cauchy integral is not an example of Cauchy's theorem, because τ is not the boundary value of v. Rather, τ is purely real-valued.

with the choice of sign in the last term corresponding to the cases where $\operatorname{Im} s_0$ has sign $\pm$. The first bracketed term is bounded in size by $(e^{\alpha N}-1)^{-1}$ since $\operatorname{Im} s = \alpha$. The same is true for the second bracketed term since $\operatorname{Im} s = -\alpha$. The third bracketed term is bounded by $(e^{|\operatorname{Im} s_0| N}-1)^{-1}$. Combining these estimates and the bound on f on the strip boundary completes the proof. □

Now, since τ is real analytic on $\partial\Omega$, it may be continued as a bounded holomorphic function in the closure of some annular neighborhood A_α. Let us choose $\alpha > 0$ so that Z is also analytic and invertible in the closure of A_α as discussed in section 2.1. This is sufficient for the pullback $\tilde{\tau}$ to be bounded and holomorphic in the closed s-plane strip of half-width α. Consider a target evaluation point $z \in A_\alpha$, which then has a unique preimage $s = Z^{-1}(z)$ with $|\operatorname{Im} s| < \alpha$. Recalling the native evaluation (2.3), define the error function

$$\epsilon_N := \operatorname{Re}(v^{(N)} - v) \ . \tag{2.7}$$

Applying Lemma 2 in the strip $|\operatorname{Im} s| < \alpha$, noticing that the residue of the integrand in (2.2) is just $\tilde{\tau}(s)$, and bounding the dominant first term in (2.6) by a simple exponential, we have shown:

THEOREM 3. *Let Z be the conformal map and τ the density function defined at the beginning of this section. Let A_α, $\alpha > 0$, be an annular neighborhood in the closure of which τ is holomorphic and bounded, and Z^{-1} is holomorphic. Then at each target point $z \in A_\alpha \backslash \partial\Omega$ we have exponential convergence of the error of the Laplace double-layer potential evaluated with the N-point trapezoid rule in the s variable. That is, there exist constants C and N_0 such that*

$$|\epsilon_N(z)| \ \le \ C e^{-|\operatorname{Im} s| N} \qquad \text{for all } N \ge N_0 \tag{2.8}$$

where $Z(s) = z$. The constant C may be chosen to be any number greater than $|\tilde{\tau}(s)| = |\tau(z)|$.

To summarize: convergence is exponential with rate given by the imaginary part of the preimage of the target point under the complexification of the boundary parametrization.

REMARK 4. *It is possible to choose constants C and M for which* (2.8) *holds uniformly in any* compact *subset of $A_\alpha \backslash \partial\Omega$, but this is impossible over the entire set $A_\alpha \backslash \partial\Omega$ because the value of N at which exponential convergence sets in diverges as $1/|\operatorname{Im} s|$ as one approaches $\partial\Omega$. Intuitively, this failure occurs because, no matter how large N is, individual quadrature points are always "visible" from close enough to the boundary. However, the constant C may be chosen uniformly on the entire set to be any number greater than* $\sup_{z \in A_\alpha} |\tau(z)|$.

In Fig. 1.1 (a) we plot contours of constant $|\epsilon_N(z)|$ as the target point z is varied over a nonsymmetric interior domain Ω with analytic boundary, for fixed N, and the simplest case $\tau \equiv 1$ which generates the potential $u \equiv -1$ in Ω. We overlay (as darker curves) predicted contours using Theorem 3. In the annular neighborhood A_α (shown in grey in Fig. 2.1) the predicted contour for an error level ϵ is the curve $\Gamma_{-\log(\epsilon/C)/N}$, where the lower bound $C = 1$ was used. The match between the light and dark contours is almost perfect (apart from periodic "scalloping" due to oscillation in the error at the node frequency).

REMARK 5. *Note that it would be possible to use the properties of the cotangent function to improve Theorem 3 to include a* lower *bound on $|v^{(N)}(z) - v(z)|$ asymptotically approaching the upper bound. We have not pursued this, since after taking the real part the error has no lower bound; rather, it oscillates in sign at the node frequency as shown by the "fingers" in Fig. 1.1 (a).*

What happens further into the domain Ω, i.e. for α values larger than that for which Z is invertible? Here, since Z^{-1} starts to become multi-valued, there are multiple preimages which lie within a given strip $|\operatorname{Im} s| \le \alpha$ (e.g. see large dots in Fig. 2.1). To analyze this would require a variant of Lemma 2 with multiple poles. We prefer an intuitive explanation. Let us assume that $\tilde{\tau}$ remains holomorphic throughout such a wider strip. Then it is clear that the s-plane pole *closest to the real axis* will dominate the error for sufficiently large N because it creates the slowest exponential decay rate. Hence, to generate each predicted contour in Fig. 1.1 we use the

set of z-plane points which have their closest preimage a distance $\alpha = -\log(\epsilon/C)/N$ from the real axis. For each α, this curve is simply the boundary of A_α, i.e. the self-intersecting curve Γ_α with all its "loops trimmed off." We see in Fig. 1.1(a) that, throughout the interior of Ω, this leads to excellent prediction of the error down to at least 14 digits of accuracy. Even features such as the cusps which occur beyond the two closest interior Schwarz singularities (at roughly 4 o'clock and 7 o'clock) are as predicted.

Eventually, for a target point deep inside the domain, all of its preimages may be further from the real axis than the widest strip in which $\tilde{\tau}$ is holomorphic. In this case, one is able to apply only the Davis theorem, and the width of the strip in which $\tilde{\tau}$ is holomorphic will now control the error (this case is never reached in Fig. 1.1(a), although it will be in (d)).

We now perform some instructive variants on this numerical experiment. In Fig. 1.1(b) we use the same domain and N as in (a), but instead of using a given τ, we solve for τ via (1.4) with the N-point Nyström method, given (entire) Dirichlet data $u(x, y) = xy$. This is a typical BVP setting, albeit a simple one. We see that the errors are similar to (a) with the major difference that the errors bottom out at around 10^{-9}: this is because τ itself only has this accuracy for the $N = 60$ nodes used ($N \geq 130$ recovers full machine precision in τ).

In (c) we repeat (a) except using a boundary shape $\partial\Omega$ distorted by a localized Gaussian "bump" at around 11 o'clock. The errors are now never smaller than 10^{-8}: note that since $\tau \equiv 1$ this cannot be due to inaccuracy, nor to lack of sufficient analyticity, in τ. Rather, the mechanism is the rapid growth in distance from $\partial\Omega$ of the contours Γ_α in this region, as α increases. This is verified by the quite good agreement with predicted contours. We observe that such growth is typical in a region with rapidly-changing curvature, which explains the well-known empirical rule that, for high accuracy, N must be chosen large enough to resolve such spatial features.

To remind the reader that Theorem 3 predicts errors just as well in the exterior as in the interior, we suggest a glance at Fig. 1.1(d), to be discussed more later. We conclude with a remark about the universality of the "safe" distance from the boundary for accurate evaluation.

REMARK 6 ("$5h$ rule"). *In practical settings, if the evaluation point is a distance $5h$ or more from the boundary, where h is the local spacing between the nodes $Z(2\pi j/N)$, then around 14 digits of accuracy in $u^{(N)}$ is typical. This is because, when the local distortion induced by the conformal map Z is small, the preimage is then a distance roughly $5 \cdot 2\pi/N$ from the real axis, giving the term $e^{-5\cdot 2\pi} \approx 2 \times 10^{-14}$ in (2.8). This relies on two assumptions: i) τ is analytic and bounded in an annular neighborhood of sufficient width, and ii) the local distortion is small on a spatial scale of a few times h. Why should these hold in practice? The answer is that they are preconditions for the Nyström method to produce a highly accurate solution density τ in the first place.*

Finally, we note that if the periodic trapezoid rule were to be replaced by a panel-based quadrature formula with Chebyshev node density, such as Gauss–Legendre, a similar analysis to the above would show that the contours of error level are the images under Z of the Bernstein ellipses [14, 35] for the panel intervals.

2.3. Single-layer case. We now present a similar analysis for single-layer potential evaluation, including a numerical verification. Recalling the fundamental solution (1.3), the single-layer potential is

$$u(x) = \int_{\partial\Omega} \Phi(x, y)\sigma(y)ds_y \ , \qquad x \in \mathbb{C}\backslash\partial\Omega \ . \tag{2.9}$$

For the proof in the Laplace case ($\omega = 0$), we need the analytic function of which this is the real part. A real-valued analytic density $\sigma \in C(\partial\Omega)$ generates a potential $u = \operatorname{Re} v$, where,

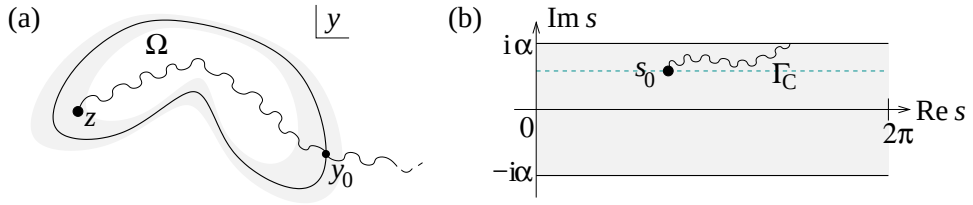

FIG. 2.2. *(a) A branch cut in the complex y plane for the function* $\log 1/(y-z)$ *in* (2.10) *for* $z \in \Omega$*, with exit point* $y_0 \in \partial\Omega$ *fixed; see Remark 7. (b) The analytic strip (excluding branch cut* Γ_C*) for the integrand in Lemma 8.*

analogously to (2.1) and (2.2),

$$v(z) \;=\; \frac{1}{2\pi}\int_{\partial\Omega}\left(\log\frac{1}{y-z}\right)\sigma(y)\,|dy| \;=\; \frac{1}{2\pi}\int_0^{2\pi}\left(\log\frac{1}{Z(s)-z}\right)\tilde{\sigma}(s)\,|Z'(s)|\,ds\;, \quad z\in\mathbb{C}\backslash\partial\Omega\;, \tag{2.10}$$

and $\tilde{\sigma}(s) := \sigma(Z(s))$ is the pullback. Note that now the magnitude of dy rather than its complex value is taken. The multiple sheets of the imaginary part of the logarithm cause the following complication.

REMARK 7. *Without a careful choice of the branch cuts of the kernel* $L(y,z) := \log 1/(y-z)$ *in* (2.10)*,* $\operatorname{Im} v$ *would fail to be a harmonic conjugate of* u*, and* v *would not be holomorphic. However, it is easy to check that sufficient conditions are as follows, which we will from now assume apply in the definition* (2.10)*. For the case* $z \in \Omega$*, the branch cut of* $L(\cdot,z)$ *must exit* Ω *only at a single point* $y_0 \in \partial\Omega$ *independent of* z *(see Fig. 2.2(a)), and, for each fixed* $y \in \partial\Omega$*,* $y \neq y_0$*,* $L(y,\cdot)$ *is continuous in* Ω*. For the case* $z \in \mathbb{C}\backslash\overline{\Omega}$*, the branch cut in* $L(\cdot,z)$ *must avoid* $\partial\Omega$*, and, for each* $y \in \partial\Omega$*, the branch cut in* $L(y,\cdot)$ *passing to infinity must avoid* Ω*.* [3]

As in the double-layer case, we will assume that $\tilde{\sigma}$ continues to a function analytic in some strip $|\operatorname{Im} s| \leq \alpha$ in which Z is analytic and invertible. The following is a variation on Lemma 2.

LEMMA 8. *Let* f *be* 2π*-periodic and analytic everywhere in the strip* $|\operatorname{Im} s| \leq \alpha$ *apart from on a branch cut* Γ_C *which starts from the point* s_0*, with* $\operatorname{Im} s_0 \neq 0$*, then proceeds to the nearer edge of the strip while avoiding the region* $|\operatorname{Im} s| \leq |\operatorname{Im} s_0|$ *(see Figure 2.2(b)). Let* $|f| \leq F$ *on the edges of the strip. Then the quadrature error* (2.4) *obeys the bound*

$$|E_N| \;\leq\; \frac{1}{e^{|\operatorname{Im} s_0|N}-1}\int_{\Gamma_C}|f^+(s)-f^-(s)|\,|ds| \;+\; \frac{4\pi F}{e^{\alpha N}-1}\;, \tag{2.11}$$

where f^+ *and* f^- *are the limiting values on either side of the branch cut.*

Proof. As in the proof of Lemma 2, we use the residue and Cauchy theorems to write E_N as

$$\int_{\Gamma_1}\left(\frac{i}{2}\cot\frac{Ns}{2}-\frac{1}{2}\right)f(s)ds - \int_{\Gamma_2}\left(\frac{i}{2}\cot\frac{Ns}{2}+\frac{1}{2}\right)f(s)ds + \int_{\Gamma_C}\left(\frac{i}{2}\cot\frac{Ns}{2}\mp\frac{1}{2}\right)[f^+(s)-f^-(s)]ds\,,$$

taking care to include the traversal of both sides of Γ_C in the relevant contours. As before, the first two integrals may be bounded to give the second term in (2.11). The last integral may be bounded by the first term in (2.11). □

We now apply this estimate to the native evaluation error for the single-layer potential (2.9) and find exponential convergence with the same rate as for the double-layer.

THEOREM 9. *Let* A_α*,* $\alpha > 0$*, be an annular neighborhood in the closure of which* σ *is holomorphic and bounded, and* Z^{-1} *is holomorphic. Let* ϵ_N *be the evaluation error function of*

[3] Note that, unless $\int_{\partial\Omega}\sigma(y)|dy| = 0$, i.e. total charge vanishes, then $\operatorname{Im} v$ itself cannot be single-valued outside Ω.

the Laplace single-layer potential (the real part of (2.10)*) with the N-point periodic trapezoid rule in the s variable. Then at each target point $z \in A_\alpha \backslash \partial\Omega$, there exist constants C and N_0 such that*

$$|\epsilon_N(z)| \;\le\; Ce^{-|\operatorname{Im} s|N} \qquad \text{for all } N \ge N_0$$

where $Z(s) = z$.

Proof. We use the notation $\overline{s}$ to indicate the complex conjugate of s. We note that the analytic continuation of $|Z'(s)|$ off the real axis is $\big(Z'(s)\overline{Z'(\overline{s})}\big)^{1/2}$, which is analytic and bounded in the strip. We also note that the imaginary part of (2.10) fails to be 2π-periodic, due to the jump in the logarithm by $2\pi i$ at $y = y_0$. We choose $y_0 = Z(0)$, and further restrict the branch cut of $\log 1/(y-z)$ in (2.10) to have a preimage which crosses the strip only along the imaginary s axis. Then we may then apply the function

$$f(s) = \frac{1}{2\pi}\left(\log\frac{1}{Z(s)-z} + is\right)\tilde{\sigma}(s)\big(Z'(s)\overline{Z'(\overline{s})}\big)^{1/2}$$

in Lemma 8. Here the new term is cancels the imaginary jump in the logarithm, making f periodic, yet does not affect the real part of the integral. Since the jump in the logarithm everywhere on Γ_C is $2\pi i$, then $\int_{\Gamma_C} |f^+(s) - f^-(s)|\,|ds|$ is bounded by a constant involving the lengths of $Z(\Gamma_C)$ and $Z(\overline{\Gamma_C})$ and an upper bound on $|\tilde{\sigma}|$ in the strip. Analogously to the proof of Theorem 3, this bounds the first exponential term in (2.11); the weaker second term can be absorbed into the constant. □

We now test the convergence for a combination of single- and double-layer densities, by checking the accuracy of the following Green's representation formula (GRF). For u a solution to the PDE (1.1) in a bounded domain $\Omega \subset \mathbb{R}^2$,

$$\int_{\partial\Omega} \Phi(x,y)\frac{\partial u}{\partial n}(y)\,ds_y - \int_{\partial\Omega}\frac{\partial\Phi(x,y)}{\partial n(y)}u(y)\,ds_y \;=\; \begin{cases} u(x), & x \in \Omega, \\ 0, & x \in \mathbb{R}^2\backslash\overline{\Omega}. \end{cases} \tag{2.12}$$

Errors for the native evaluation of this in both interior and exterior cases are shown together in Fig. 1.1(d): it is clear that the errors grow exponentially as x approaches $\partial\Omega$, from either inside or outside. In fact, this figure shows the case of the low-frequency Helmholtz equation with $\omega = 2$; the Laplace plot is almost identical. This confirms the intuition that the convergence for the low-frequency Helmholtz equation is very similar to that for the Laplace equation. (See section 4 for the formulae for the Helmholtz case.)

Why do the the errors stop decreasing with distance once an error of 10^{-12} is reached in Fig. 1.1(d)? This cannot be due to inaccurate boundary data, since the data is exact to machine precision. Rather, the answer lies in the fact that the boundary data derives from a point source (shown by a large dot) outside, but not too far from, Ω. This means that the densities cannot be analytic in a larger annular neighborhood, limiting the maximum convergence rate in N to the value at this singularity location. This illustrates the effect of densities which are not entire functions.

3. Close evaluation of Laplace layer potentials by surrogate local expansions. In this section we present and analyze the main new scheme. We have seen that the native evaluation error is exponentially small far from $\partial\Omega$, but unacceptably large close to $\partial\Omega$. Thus we define the "bad annular neighborhood" $A_{\text{bad}} := A_{\alpha_{\text{bad}}}$ where, following Remark 6, we will choose $\alpha_{\text{bad}} = 10\pi/N$, that is, a distance around $5h$ either side of $\partial\Omega$, giving around 14 digit expected accuracy throughout $\mathbb{R}^2\backslash A_{\text{bad}}$. (Obviously α_{bad} could be decreased if less accuracy is desired.) We now describe the new method for layer potential evaluation in this bad annular neighborhood,

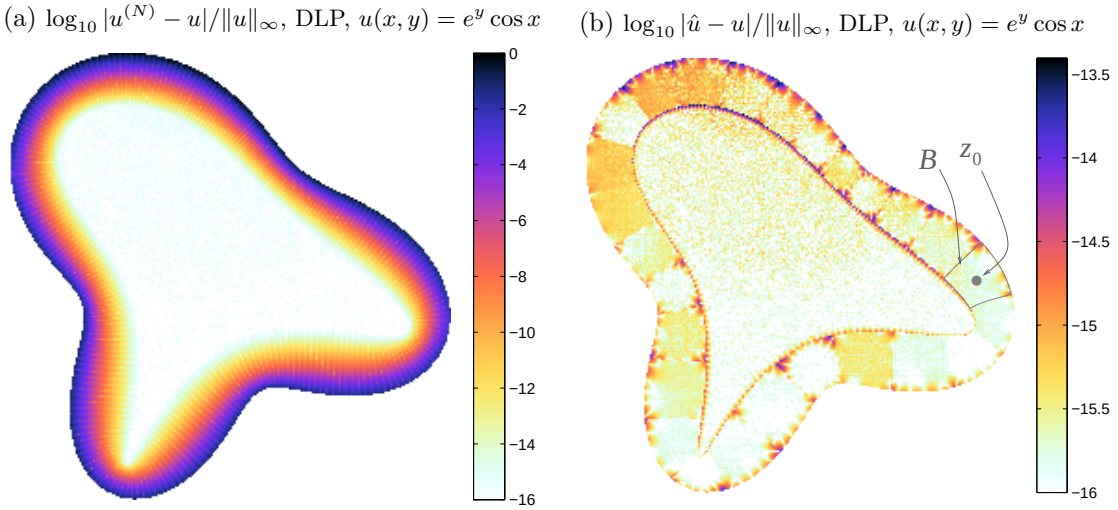

FIG. 2.3. *Error (relative to largest value) in evaluation of the solution of the Laplace BVP with Dirichlet data* $u(x,y)=e^y\cos x$*, with* $N=130$*, using the native quadrature (a), and the proposed scheme (b) with* $N_B=26$*,* $p=10$ *and* $M=4N$*. Note the change in color scale: the relative* L^∞ *error is* 4×10^{-14}*, and the relative* L^2 *error* 4×10^{-15}*. One box* B *and its center* z_0 *are labeled; many other boxes and the the annular half-neighborhood* $\Omega_{bad}=A_{bad}\cap\Omega$ *are visible due to jumps in the error.*

focusing on the case of $\Omega_{\rm bad} := A_{\rm bad}\cap\Omega$, for Ω an interior domain with $\partial\Omega$ traversed in the counter-clockwise sense. Thus, target preimages will have positive imaginary part. (The exterior case is analogous.)

Let us fix N and assume that a double-layer potential τ has been approximately computed by the Nyström method, so is represented by $\{\tau_j\}_{j=1}^N$, its values at the nodes. A precondition for accuracy of the Nyström method is that τ is well-approximated by its Nyström interpolant through these nodes. Recalling that $u=\operatorname{Re} v$, we wish to evaluate v via (2.1) in $\Omega_{\rm bad}$. We cover $\Omega_{\rm bad}$ by non-intersecting "boxes" B_b, $b=1,\ldots,N_B$. The bth box B_b is the image under Z of the s-plane rectangle $2\pi(b-1/2)/N_B\le\operatorname{Re}s<2\pi(b+1/2)/N_B$, $0\le\operatorname{Im}s\le\alpha_{\rm bad}$. By choosing $N_B=\lceil N/5\rceil$ the boxes become very close to being square. Fig. 2.3(b) shows the annular half-neighborhood and one box.

Let $B=B_b$ be the bth box. We choose an expansion center

$$z_0 := Z(2\pi b/N_B + i\alpha_0) \tag{3.1}$$

where the choice of imaginary distance $\alpha_0=\alpha_{\rm bad}/2$ places z_0 roughly central to B, and roughly $2.5\,h$ from the boundary. We represent v by a Taylor series

$$v(z)=\sum_{m=0}^{\infty} c_m(z-z_0)^m\ ,$$

which converges uniformly in B if v is analytic in some disc centered at z_0 with radius greater than R, where R is the maximum radius of the box,

$$R := \sup_{z\in B}|z-z_0|\ .$$

Each coefficient c_m can be computed, using the Cauchy formula for derivatives, as

$$c_m \;=\; \frac{-1}{2\pi i}\int_{\partial\Omega}\frac{\tau(y)}{(y-z_0)^{m+1}}dy = \frac{-1}{2\pi i}\int_0^{2\pi}\frac{\tilde\tau(s)}{(Z(s)-z_0)^{m+1}}Z'(s)ds\ , \qquad m=0,1,\ldots \tag{3.2}$$

To approximate the latter integral we use the periodic trapezoid rule with M new nodes, $s_j = 2\pi j/M$, $j = 1, \ldots, M$, which we call "fine" nodes, that is,

$$c_m \;\approx\; \hat{c}_m \;:=\; \frac{-1}{iM} \sum_{j=1}^{M} \frac{\tilde{\tau}(s_j)}{(Z(s_j) - z_0)^{m+1}} Z'(s_j) \;. \tag{3.3}$$

Since z_0 is in the bad annular neighborhood, clearly we need $M > N$ to evaluate even the coefficient c_0 accurately: one would expect, since $\alpha_0 = \alpha_{\rm bad}/2$, that $M = 2N$ would be sufficient. However, for large m the term $1/(Z(s) - z_0)^{m+1}$ is now *oscillatory*, so an even larger M will be needed. Note that $\tilde{\tau}(s_j)$ must be found by interpolation from its values at the N Nyström nodes. Our surrogate potential in the box B is then simply the truncated series with the above approximated coefficients, and p terms, i.e.

$$v(z) \approx \hat{v}(z) = \sum_{m=0}^{p-1} \hat{c}_m (z - z_0)^m \;, \qquad z \text{ in box } B \;. \tag{3.4}$$

Finally the real part must be taken, thus in each of the boxes we set $\hat{u} = \mathrm{Re}\,\hat{v}$. We claim that, if the parameters p and M are well chosen, then $\hat{u} \approx u$ uniformly to high accuracy throughout $\Omega_{\rm bad}$. (In $\Omega \backslash \Omega_{\rm bad}$ we revert to $\hat{u} = u^{(N)}$, the native evaluation scheme.)

This is illustrated in Fig. 2.3: we first solve (1.4) for τ, given entire boundary data, using $N = 130$, which is around the minimum N required to achieve full accuracy in τ. We then try the new evaluation scheme, using a standard trigonometric polynomial interpolant (e.g. see [26, Sec. 11.3], implemented easily via the FFT) to approximate τ at the fine nodes, which gives a $\hat{u}$ with uniform relative accuracy of 13.5 digits (14 digits in the $L^2(\Omega)$ norm). Note that, since the boundary data is entire, this is a well-behaved example; we consider more challenging cases after the following analysis.

3.1. A convergence theorem for the Laplace double-layer potential. Because the boxes touch $\partial\Omega$, uniform convergence of the Taylor series of v in any box generally demands that v have an analytic continuation some distance *outside* Ω, which might seem far-fetched. But it turns out that the region of analyticity of v is at least as large as that of τ.

PROPOSITION 10. *Let τ be analytic in some closed annular neighborhood A of $\partial\Omega$. Then v given by* (2.2) *is analytic in Ω and moreover continues as an analytic function throughout $A \backslash \Omega$.*

Proof. Let Γ be the exterior boundary of A. Consider the identity arising from (2.1),

$$\frac{-1}{2\pi i} \int_\Gamma \frac{\tau(y)}{y - z} dy \;=\; v(z) + \frac{-1}{2\pi i} \int_{\Gamma - \partial\Omega} \frac{\tau(y)}{y - z} dy \;,$$

where $-\partial\Omega$ means $\partial\Omega$ traversed in the opposite direction. Both sides equal $v(z)$ for $z \in \Omega$ since the second term on the right is zero by Cauchy's theorem. However, as a function of z, the left side is analytic throughout $A \cup \Omega$, and thus provides the desired analytic continuation of v. ∎

To state a convergence result, some control is needed of the distortion induced by the conformal map Z. Let A_α be an annular neighborhood in which Z^{-1} exists. Recall that box B has a center z_0 (which we assume is in A_α) with α_0 the imaginary part of its preimage, and radius R. Given this, we define a *geometric distortion quantity*,

$$\gamma = \gamma_{z_0,R} := \sup_{0<a<\alpha_0} \frac{R}{d(\Gamma_a, z_0)} \frac{\alpha_0 - a}{\alpha_0} \;, \tag{3.5}$$

where $d(X, y)$ means the minimum Euclidean distance from the point y to points in the set X. It is clear that, as $a > 0$ approaches α_0 from below, the curve Γ_a first touches z_0 at $a = \alpha_0$; an

interpretation of γ^{-1} is a scaled lower bound on the ratio between its distance and $\alpha_0 - a$ (we note that in peculiar geometries it may be that the nearest part of Γ_a is not within the box B). For an undistorted square box with z_0 at its center, $\gamma = \sqrt{2}$; in practical settings γ is around 1.5 to 2. For the geometry in Fig. 2.3(b), the median γ is 1.7 and the maximum 2.4 (occurring at around 7 o'clock, where the closest interior Schwarz singularity lies).[4]

Our result concerns convergence of the surrogate scheme simultaneously in the expansion order p and the number of fine nodes M.

THEOREM 11. *Let B be a box with radius R. Let the center z_0 have $\operatorname{Im} z_0 = \alpha_0$, with A_{α_0} an annular neighborhood in which Z^{-1} is holomorphic, and in which the density τ is holomorphic and bounded. Let the double-layer potential v given by (2.1) be analytic in an open neighborhood of some closed disc of radius $\rho > R$ about z_0. Let $\hat{v}$ be given by the surrogate scheme (3.4) with $\hat{c}_m$ given by (3.3), using the exact density at the fine nodes $\tilde{\tau}(s_j) = \tau(Z(s_j))$, $j = 1, \dots, M$. Then the error function*

$$\hat{\epsilon} = \operatorname{Re}(\hat{v} - v)$$

has the uniform bound

$$\hat{\epsilon}_B \;:=\; \sup_{z\in B} |\hat{\epsilon}(z)| \;\le\; C\left(\frac{R}{\rho}\right)^p + C(e\gamma\alpha_0 M)^p p^{1-p} e^{-\alpha_0 M} \qquad \text{for all } M \ge 1/\alpha_0,\ 1 \le p \le M/2\ , \tag{3.6}$$

where C indicates constants that depend on $\partial\Omega$, Z, τ, B and z_0, but not on p nor M.

The restriction to $M \ge 1/\alpha_0$ always holds in practice because N is already many times larger than $1/\alpha_0$, whilst $M > N$. The restriction on p also holds in practice, since usually $M > 10^2$. The first term in (3.6) is simply the truncation error of the Taylor series for v in the box, so is independent of M. Note that the required analyticity of v is already given by Prop. 10 and the analyticity of τ, with $\rho \approx \sqrt{2}R$, unless the local distortion is very large.

The M dependence of the second term may be written $e^{-\alpha_0 M + p\log M}$, showing that (at fixed p) this term is asymptotically exponentially convergent in M with rate α_0. Thus the whole scheme is also asymptotically exponentially convergent: given arbitrary $\epsilon > 0$, by fixing p such that the first term is smaller than $\epsilon/2$, one may find a value of M such that the second term is also smaller than $\epsilon/2$, and have $\hat{\epsilon}_B \le \epsilon$. There is a subtlety. Fixing M while increasing p is a *bad idea*: in practice it leads to exponential *divergence*, because the factor $e\gamma\alpha_0 M$ is usually large (around 300, independent of the problem size), and hence the second term grows exponentially in p for typical p (less than 30).

Proof. We will use C to indicate (different) constants that have only the dependence stated in the theorem. Comparing $\hat{v}$ to the exact Taylor series for v, and using $|z - z_0| \le R$,

$$\hat{\epsilon}_B \;\le\; \sup_{z\in B}\left|\sum_{m=0}^{p-1} \hat{c}_m (z-z_0)^m - v(z)\right| \;\le\; \sum_{m\ge p} |c_m| R^m + \sum_{m=0}^{p-1} |\hat{c}_m - c_m| R^m\ . \tag{3.7}$$

For the first term we use $|c_m| \le C/\rho_m$, which follows from the Cauchy integral formula in the disc of radius ρ (e.g. [33, Cor. 4.3]), and bound the geometric sum. For the second term we apply Theorem 1 to the s-integral (3.2) in a strip of width α, so that, for each $\alpha \in (0, \alpha_0)$,

$$|\hat{c}_m - c_m| \;\le\; \left(\sup_{\operatorname{Im} s = \pm\alpha} |\tilde{\tau}(s) Z'(s)|\right) \frac{2}{d(\Gamma_\alpha, z_0)^{m+1}} \frac{1}{e^{\alpha M} - 1} \;\le\; \frac{C}{d(\Gamma_\alpha, z_0)^{m+1}} \frac{1}{e^{\alpha M} - 1}\ , \tag{3.8}$$

[4] These γ values are accurately estimated using 20 values of α with 10^3 points on each curve Γ_α.

where the second estimate follows from the maximum modulus principle and boundedness of Z' and $\tilde{\tau}$ in the α_0-strip. Inserting the above two estimates into (3.7), then bounding each term in the sum by the last term since $R > d(\Gamma_\alpha, z_0)$, gives

$$\hat{\epsilon}_B \leq \frac{C}{1-R/\rho}\left(\frac{R}{\rho}\right)^p + \sum_{m=0}^{p-1} R^m \frac{C}{d(\Gamma_\alpha,z_0)^{m+1}} \frac{1}{e^{\alpha M}-1} \;\leq\; C\left(\frac{R}{\rho}\right)^p + Cp\left(\frac{R}{d(\Gamma_\alpha,z_0)}\right)^p \frac{1}{e^{\alpha M}-1}$$
$$\leq C\left(\frac{R}{\rho}\right)^p + Cp\left(\frac{\gamma\alpha_0}{\alpha_0-\alpha}\right)^p e^{-\alpha M} \qquad \text{for all } M > 1/2\alpha \;, \tag{3.9}$$

where in the last step we used (3.5) and simplified the exponential bound. This estimate holds for each $\alpha \in (0, \alpha_0)$; however, if α is chosen too small, the last exponential will decay very slowly. Conversely, if α approaches α_0 then the distance $d(\Gamma_\alpha, z_0)$ vanishes and its pth negative power blows up rapidly. For each p and M, the optimal value $\hat{\alpha}$ which minimizes this second term is found by setting $\partial/\partial\alpha$ of the logarithm of this term to zero, and solving, giving

$$\hat{\alpha} = \alpha_0 - \frac{p}{M}$$

which by the condition on p is never smaller than $\alpha_0/2$. Substituting $\alpha = \hat{\alpha}$ in (3.9) gives (3.6). □

To interpret the theorem we need to relate it to the original N Nyström nodes (notice that N does not appear in the theorem). We introduce two dimensionless parameters. Let

$$\delta \;:=\; \frac{\alpha_0 N}{2\pi}$$

be the distance of the center z_0 from the boundary in units of local node spacing h; recall that we set $\delta = 2.5$ in the above. Let

$$\beta \;:=\; \frac{M}{N}$$

be the "upsampling ratio", the ratio of the number of fine nodes to the number of original Nyström nodes.

It is clear that the effort to compute (3.3) as presented scales as $O(\beta)$, once N, N_B and p are fixed. Thus we wish to know the minimum β needed, and rewrite the second term in (3.6) as

$$Cp\exp\left(-2\pi\delta\beta + p\log\big(2\pi\delta\beta\frac{\gamma e}{p}\big)\right) . \tag{3.10}$$

For instance, fixing $\delta = 2.5$ and taking $\gamma \approx 1.7$, we may solve (via rootfinding) for the approximate β required for convergence by assuming that C is $O(1)$ then equating the exponential term in (3.10) to the desired error level, e.g. 10^{-14}. The predicted results are: for $p = 10$, $\beta = 4.2$ is sufficient. This matches well the finding that $\beta = 4.0$ was sufficient to give around 14-digit accuracy in the example of Fig. 2.3(b). For $p = 20$, $\beta = 5.9$ is sufficient, indicating that β need grow only weakly with p.

REMARK 12 (box centers). *Equation* (3.10) *suggests that increasing* δ *(moving the centers towards the far edge of their boxes) might increase accuracy. In practice, however, we find that this does not help because both* R/ρ *and* γ *increase for the worst-case boxes.*

3.2. Numerical performance and the effect of a nearby singularity. We now study in more detail the performance of the surrogate method for the Laplace double-layer potential, in the context of solving a Dirichlet BVP using the Nyström method.

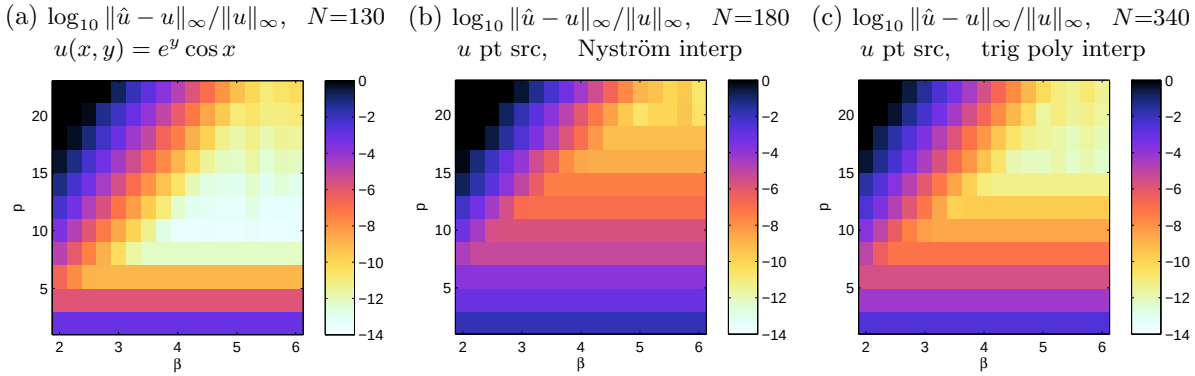

FIG. 3.1. *Convergence of the maximum relative $L^\infty(\Omega)$ error with respect to both parameters p and β, at fixed N and δ, for the surrogate evaluation scheme of the double-layer potential solving a Laplace BVP in the domain of Fig. 2.3. (a) Dirichlet data $u(x,y)=e^y\cos x$, with $N=130$, trigonometric interpolant. (b) $u(x,y)=\log\|(x,y)-(1,0.5)\|$, with $N=180$, Nyström interpolant. (c) $u(x,y)=\log\|(x,y)-(1,0.5)\|$, with $N=340$, trigonometric interpolant.*

We first return to the case presented in Fig. 2.3: boundary data $u(x,y)=e^y\cos x$, an entire function. We fix $N=130$ nodes (around the value at which the Nyström method has completely converged), and use a trigonometric polynomial interpolant to get τ at the fine nodes. With the N_B boxes fixed, we consider convergence with respect to p and β. Fig. 3.1(a) shows the resulting L^∞ errors (estimated on a spatial grid of spacing 0.02) relative to $\|u\|_\infty$. The two terms in Theorem 11 are clearly visible: the errors are always large at small p (due to the first term in (3.6)), but even at large p the errors are large when β is too small (due to the second term). The β needed for convergence grows roughly linearly in p, as one would expect if the value of the log in (3.10) is treated as roughly constant. In both directions convergence appears exponential, exceeding 13 digits once $p=10$ and $\beta\geq 4$.

REMARK 13. *For $p\geq 16$ the best achievable error (once β has converged) worsens slightly as p grows. We believe that this is due to catastrophic cancellation in the oscillatory integrand* (3.2) *for large m, combined with the usual double-precision round-off error. This seems to be a fundamental limit of the local expansion surrogate method implemented in floating-point arithmetic. However, the loss is quite mild: even at $p=22$ it is only 2.5 digits.*

We now change to boundary data $u(x,y)=\log\|(x,y)-(1,0.5)\|$, which is still a Laplace solution in Ω and is still real analytic on $\partial\Omega$, but whose analytic continuation outside Ω has a *singularity* at the exterior point $(1,0.5)$, a distance of only 0.24 from $\partial\Omega$. Its preimage has distance from the real axis $\alpha_* := -\operatorname{Im} Z^{-1}(1+0.5i)\approx 0.176$. The Nyström method fully converged by $N=180$, as assessed by the error at a distant interior point and by the density values at the nodes τ_j: see the first two curves in Fig. 3.2(a). In fact the convergence rate matches $e^{-\alpha_* N}$, as is to be expected, since the Nyström convergence rate is known to be the same as that of the underlying quadrature scheme (see discussion after [26, Cor. 12.9]), which here is controlled by the singularity via the Davis theorem.

However, turning to surrogate evaluation, Fig. 3.1(b) shows that the error due to Taylor truncation converges quite slowly with p, as is inevitable for a nearby singularity, pushing the optimal p up to around 22. In the best case only 10 digits are achieved. To evaluate τ at the fine nodes, the Nyström interpolant [26, Sec. 12.2] was used here rather than the trigonometric polynomial interpolant, for the following reason.

REMARK 14 (interpolants). *In the regime where a nearby singularity in the right-hand side data controls the convergence rate (rather than the kernel function), the Nyström interpolant*

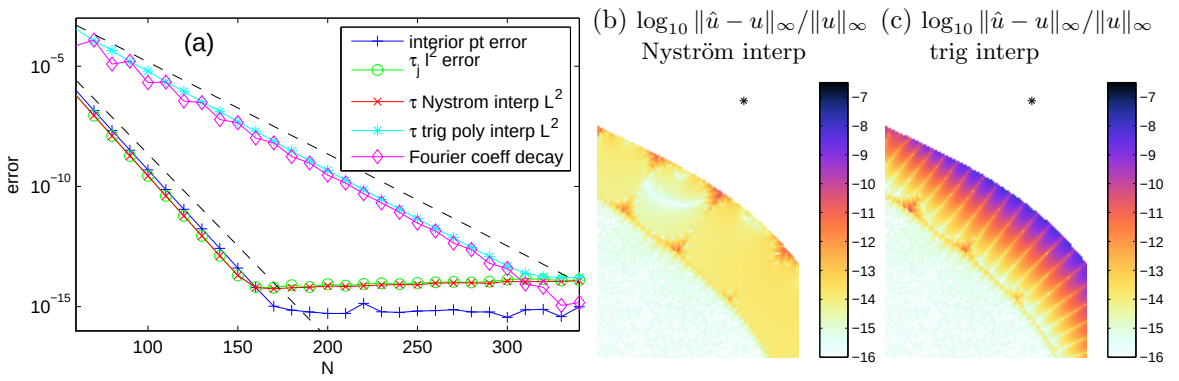

FIG. 3.2. *Comparing two interpolants in the Nyström solution of τ, for solution of the Dirichlet BVP for Laplace's equation, with data $u(x,y) = \log\|(x,y)-(1,0.5)\|$, in the domain of Fig. 2.3. (a) shows convergence of various errors: the solution at a distant interior point, the l_2 error at the nodes, the $L^2([0,2\pi))$ errors for the Nyström and trigonometric interpolants, and the size of the $N/2$ Fourier coefficient of τ. Dotted lines show exponential decay at the rates $e^{-\alpha_* N}$ and $e^{-\alpha_* N/2}$. Fixing $N = 180$, a zoomed plot of the relative error in the surrogate scheme is shown using (b) Nyström interpolant, and (c) trigonometric interpolant, near the singularity (shown by a $*$).*

converges twice as fast *as the trigonometric polynomial interpolant, as shown by Fig. 3.2(a). This reflects the fact that periodic trapezoid quadrature is "twice as good" as trigonometric interpolation [26, p.201], because the former is exact for Fourier components with index magnitudes up to $N-1$, but the latter is exact only up to $N/2$. Informally speaking, the trapezoid rule (and hence the Nyström interpolant) "beats the Nyquist sampling theorem by a factor of two!"*

Thus the Nyström interpolant is preferred in this context when it is desired that N be its smallest converged value. Fig. 3.2(b) and (c) show the loss in accuracy (and spurious evanescent waves which appear near $\partial\Omega$) that result from attempting to use the inferior trigonometric interpolant. Note that this loss would occur for *any* accurate close-evaluation scheme, since it is a loss of accuracy in the function τ itself. Nor is the Nyström interpolant perfect: it requires the application of a M-by-N dense matrix, and it appears to cause up to 1 digit more roundoff error than in the trigonometic case.

However, by nearly doubling N to 340 (which is somewhat wasteful), the Fourier coefficents of $\tilde{\tau}$ decay to machine precision by index $N/2$ (see Fig. 3.2(a)) making the trigonometric interpolant as accurate as the Nyström one, whilst box radii decrease so that the p-convergence is faster. We show convergence for this trigonometric case in Fig. 3.1(c): 12 digits accuracy result at $p = 16$ and $\beta = 5$.

3.3. The Laplace single-layer case and a Neumann problem. Recall that the single-layer potential is the real part of v given by (2.10). By writing $\log 1/(y-z) = \log 1/(y-z_0) - \log(1-\frac{z-z_0}{y-y_0})$, and using the Taylor series for the second logarithm, we find that the single-layer version of (3.2) is

$$c_0 \;=\; \frac{1}{2\pi}\int_{\partial\Omega}\left(\log\frac{1}{y-z_0}\right)\sigma(y)|dy|, \qquad c_m \;=\; \frac{1}{2\pi m}\int_{\partial\Omega}\frac{\sigma(y)}{(y-z_0)^m}|dy|, \quad m=1,2,\ldots \tag{3.11}$$

Here for c_0 the branch cuts of Remark 7 apply. The convergence of the surrogate scheme is then as least as good as for the double-layer case.

THEOREM 15. *The version of Theorem 11 corresponding to the single-layer potential holds. That is, with the same conditions, the uniform error bound* (3.6) *holds, but with τ changed to*

σ, the potential v given by (2.10)*, and $\hat{c}_m$ given by the M-node periodic trapezoid rule applied to the parametrized version of* (3.11)*.*

Proof. Using (3.11) and the Davis theorem, the coefficient errors analogous to (3.8) are

$$|\hat{c}_0 - c_0| \le C\big|\log d(\Gamma_\alpha, z_0)\big| \frac{1}{e^{\alpha M} - 1} \ , \qquad |\hat{c}_m - c_m| \le \frac{C}{m\, d(\Gamma_\alpha, z_0)^m} \frac{1}{e^{\alpha M} - 1} \ , \quad m = 1, 2, \dots$$

But for each $m = 0, 1, \dots$, this error is smaller than that in (3.8), after possibly a change in the constant C. The rest of the proof follows through. □

REMARK 16. *Note that the change in the constant C referred to is generally in the favorable direction: since $d(\Gamma_\alpha, z_0) < R$, the constant may be multiplied by a factor given by the larger of R and $R|\log R|$, which are usually small. Also note that, since factors of $1/m$ arise in the single-layer case, but were not taken advantage of, the p-convergence could probably be improved slightly.*

As an application, we now report numerical results for the interior Laplace–Neumann BVP

$$\Delta u = 0 \qquad \text{in } \Omega \tag{3.12}$$

$$\partial u/\partial n = f \qquad \text{on } \partial\Omega \tag{3.13}$$

which has a solution only if f has zero mean on $\partial\Omega$, and in that case the solution is unique only up to an additive constant. Following [1, Sec. 7.2] we use the integral equation

$$(D^* + K + \tfrac{1}{2}I)\sigma = f \ , \tag{3.14}$$

where D^* has kernel $k(x, y) = \partial\Phi(x, y)/\partial n(x)$, and K is the boundary operator which returns the value of its operand at some (fixed but arbitrary) point on $\partial\Omega$. The solution is then recovered up to an unknown constant by (2.9); we compare against the exact interior solution after subtracting the value of the constant measured at a single point.

As with the Dirichlet case, we test with boundary data coming from the entire function $u(x, y) = e^y \cos x$, or from the function with a nearby singularity $u(x, y) = \log \|(x, y) - (1, 0.5)\|$. In the entire case, we find, as for the Dirichlet case, that the convergence of the Nyström method saturates at around $N = 130$. Surrogate expansion (using the Nyström interpolant) then gives a maximum relative error of 6×10^{-15} at $p = 10$ and $\beta = 4$.[5] For the singularity case, convergence of the Nyström method is complete at around a value $N = 200$ and the maximum relative surrogate evaluation error is then found to be 6×10^{-13} at $p = 24$ and $\beta = 5.5$. Notice that, for both data types, these single-layer errors are improved by at least one digit over the double-layer errors reported in Sec. 3.2; this may be explained by Remark 16. Indeed, the convergence plots analogous to Fig. 3.1(a) and (b) are very similar but show a gain of around 1 extra digit of accuracy.

4. The Helmholtz equation and an $O(N)$ close evaluation scheme. We now move to a PDE for which we no longer have theorems, but which has important applications.

4.1. Implementation and convergence test for the Helmholtz equation. The above close evaluation scheme for $\omega = 0$ with real-valued potentials is very easily adapted to the Helmholtz equation ($\omega > 0$) with complex potentials, by replacing a couple of formulae. The

[5] Curiously, the trigonometric interpolant in this case does not become fully accurate until $N = 340$, which is *much* more than the value $N = 130$ for the Dirichlet case. We cannot explain this—neither result is as predicted by Remark 14—and it tells us that there is still more to understand about the Nyström method convergence rates in analytic BVP settings.

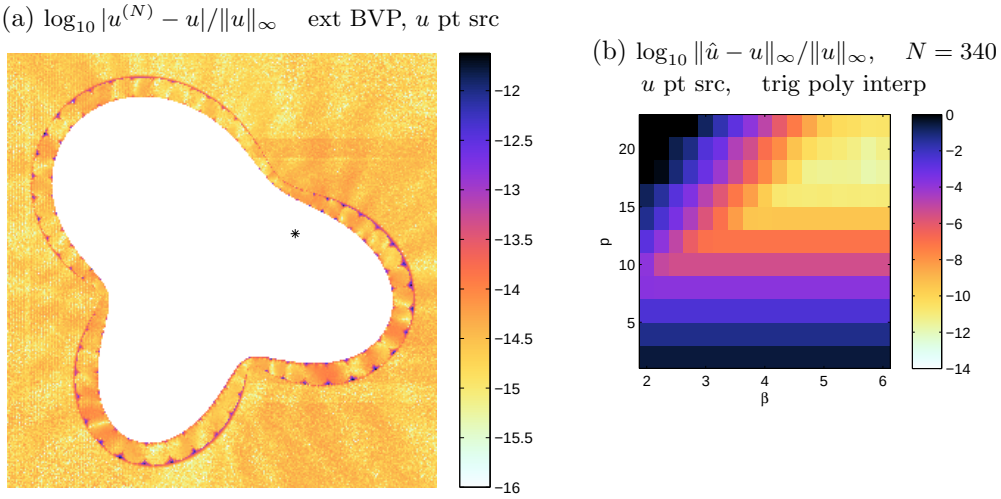

FIG. 4.1. *Relative error for the surrogate scheme for the exterior Helmholtz Dirichlet problem at $\omega = 30$ (the diameter is 12 wavelengths), with known solution $\Phi(\cdot, x_0)$ with x_0 as shown by the $*$ symbol. (a) error plot for $N = 340$, $p = 18$ and $\beta = 6$ (many boxes are apparent). (b) $L^\infty(\Omega)$ error convergence with respect to parameters p and β, at fixed $N = 340$.*

p-term Taylor expansion (3.4) for v (whose real part was taken to get u), is replaced by a local (Fourier-Bessel) expansion with $2p-1$ terms,

$$\hat{u}(z) = \sum_{|m|<p} c_m e^{im\theta} J_m(\omega r) \;, \qquad \text{where } \; z - z_0 = re^{i\theta}, \qquad z \text{ in box } B \;,$$

i.e. (r,θ) is the polar coordinate system with origin z_0. For the single-layer potential, recalling (1.3) and Graf's addition formula [31, (10.3.7)], the formula (3.11) is replaced by

$$c_m \;=\; \frac{i}{4}\int_{\partial\Omega} e^{-im\theta_y} H_m^{(1)}(\omega r_y)\sigma(y)ds_y \;, \qquad |m| < p \;, \tag{4.1}$$

where (r_y, θ_y) are the polar coordinates of the point y relative to the origin z_0. Using the addition formula, the reflection formulae, after some simplification, the Cauchy formula (3.2) is replaced by

$$c_m \;=\; \frac{i\omega}{8}\int_{\partial\Omega} \left[e^{-i(m-1)\theta_y - i\nu_y} H_{m-1}^{(1)}(\omega r_y) - e^{-i(m+1)\theta_y + i\nu_y} H_{m+1}^{(1)}(\omega r_y)\right] \tau(y)ds_y \;, \qquad |m| < p \;, \tag{4.2}$$

where ν_y is the angle of the outward normal at $y \in \partial\Omega$. The use of the above three formulae is all the change needed to make the Helmholtz version of the scheme.

We have already seen in Fig. 1.1(d) that the Laplace equation predictions for the native evaluation error (theorems 3 and 9) also hold well for the low-frequency Helmholtz equation. This is to be expected, since for $\omega > 0$ the fundamental solution (1.3) remains analytic away from the origin, so the Davis theorem applies to the parametrized (1.5) and (2.9), allowing the Laplace convergence rate to be approached. (We leave the Helmholtz equivalents of the tight theorems 3 and 9 for future work.) Therefore we apply the same criterion for the bad annular neighborhood Ω_{bad} as in section 3.

We test the scheme in the context of BVP applications, namely the Dirichlet problem in an exterior domain Ω, with the usual Sommerfeld radiation condition [11, (3.62)], for which the

integral equation (1.4) is replaced by the so-called combined-field formulation [11, p. 48]

$$(D - i\omega S + \tfrac{1}{2}I)\tau = f \tag{4.3}$$

which is well conditioned for all $\omega > 0$. Here the potential is represented as

$$u(x) = \int_{\partial\Omega} \left[\frac{\partial\Phi(x,y)}{\partial n(y)} - i\omega\Phi(x,y) \right] \tau(y) ds_y \ , \qquad x \in \Omega \ , \tag{4.4}$$

and we choose the right-hand side data as coming from an interior point-source shown by the $*$ in Fig. 4.1(a). To achieve spectral accuracy in the Nyström method, special quadratures for the weak logarithmic singularity are needed; we use the scheme of Kress [25] (also see [18]).

We fix $\omega = 30$ (thus $\partial\Omega$ is 12 wavelengths across), and find that the Nyström method converges at around $N = 340$. A maximum relative surrogate evaluation error around 3×10^{-12} is then achieved at $p = 18$, $\beta = 6$, as shown in Fig. 4.1(a). This is dominated by errors in the box corners furthest from $\partial\Omega$, especially those with distortion due to a convex part of $\partial\Omega$. Since this holds largely independently of the singularity location x_0, we believe it is instead controlled by the wavenumber ω. The $L^\infty(\Omega)$ convergence with p and β is very similar to the Laplace case, as shown in Fig. 4.1(b). $p = 18$ is close to optimal, since for $p > 20$ the error worsens—the explanation is believed to be as in Remark 13, but demands further study. However, simply by making the boxes slightly narrower and more numerous by setting $N_B = \lceil N/4 \rceil$, with 25% more effort we cut the maximum relative error to 3×10^{-13}.

REMARK 17 (interpolants revisited). *A natural question is: does Remark 14 hold for the Helmholtz equation? The answer is no, at least for the Kress scheme (which is one of the best known [18]). Although a Nyström interpolant does exist for the weakly-singular kernels [25, (3.3)], we believe it has little advantage over the trigonometric interpolant. This is because the product-quadrature scheme of Kress is only accurate for Fourier components of indices up to $N/2$. In data with a nearby singularity with preimage at a distance α_* from the real axis, the convergence rate is thus only $e^{-\alpha_s N/2}$, or half that of the Laplace case.*

4.2. Evaluation via a correction to the fast multipole method. At higher frequencies ω and/or complex geometries $\partial\Omega$, the value of N needed for convergence of the Nyström method is pushed higher. For instance, with fixed geometry the high frequency asymptotics is empirically $N = O(\omega)$, i.e. a constant number of nodes per wavelength [3, 19]. The surrogate scheme as presented requires $O(pN^2)$ effort, by evaluating p coefficients at $N_B = O(N)$ centers using $M = O(N)$ fine node kernel evaluations for each. Clearly, a better scaling with N would be preferred. We now show that at fixed frequency, linear scaling is easy to achieve by locally correcting the FMM.

Fixing a box B with center z_0, and taking e.g. the Helmholtz single-layer potential (2.9), we split the integral into "near" and "far" parts using a cut-off radius G. We apply the surrogate local expansion only for the near part to get

$$\hat{u}(x) \ = \ \sum_{|m|<p} c_m^{(\mathrm{near})} e^{im\theta} J_m(\omega r) \ + \ \int_{y\in\partial\Omega, |y-z_0|>G} \Phi(x,y)\sigma(y) ds_y \ , \qquad x \text{ in box } B \ , \tag{4.5}$$

where, adapting (4.1), the coefficients of the potential due to the near part of the integral only are,

$$c_m^{(\mathrm{near})} = \frac{i}{4} \int_{y\in\partial\Omega, |y-z_0|\le G} e^{-im\theta_y} H_m^{(1)}(\omega r_y)\sigma(y) ds_y \approx \frac{i\pi}{2M} \sum_{j\in J_{\mathrm{near}}} e^{-im\theta_{Z(s_j)}} H_m^{(1)}(\omega r_{Z(s_j)}) \tilde{\sigma}(s_j) Z'(s_j) \ , \tag{4.6}$$

and where $J_{\rm near} := \{j : |Z(s_j) - z_0| \le G\}$ is the index set of the "near" subset of the fine nodes $\{s_j\}_{j=1}^M$. The cut-off G must be large enough that the fictitious singularities induced at the ends of the near interval are distant enough not to slow down the convergence of the local expansion, thus we choose G several times the box radius R, and so evaluating these sums takes $O(p)$ effort. Applying the fine quadrature to the far part of the integral in (4.5) gives

$$\int_{y\in\partial\Omega, |y-z_0|>G} \Phi(x,y)\sigma(y)ds_y \;\approx\; \frac{2\pi}{M}\sum_{j=1}^{M} \Phi(x,Z(s_j))\tilde\sigma(s_j)Z'(s_j) - \frac{2\pi}{M}\sum_{j\in J_{\rm near}} \Phi(x,Z(s_j))\tilde\sigma(s_j)Z'(s_j)\,. \tag{4.7}$$

The first sum can be evaluated *for all target points x in all boxes* in a single FMM call. Making the reasonable assumption that the user demands $O(1)$ targets per box, this FMM call requires $O(N)$ effort. Finally, the second sum is a local correction that takes $O(1)$ effort. The total effort to evaluate $\hat u(x)$ for all x in all N_B boxes is thus $O(pN)$. A similar scheme applies for the double-layer potential.

Since the first (FMM) term in (4.7) accurately approximates u everywhere except in a narrower neighborhood A_α with $\alpha = 10\pi/\beta N$, we in fact may, and will, shrink the boxes in the normal direction, as long as they cover A_α, whilst using this first term in the remaining part of $\Omega_{\rm bad}$. This has the advantage of avoiding larger errors that tend to occur in the distant corners of boxes.

REMARK 18 (no end corrections). *At this point the reader might very well suspect that, since both of the above integrals are on non-periodic intervals, the trapezoid rule would give at best low-order $O(1/M)$ convergence unless higher-order end correction rules were used. In fact, spectral accuracy equal to that of the trapezoid rule on the original periodic integrand is observed. The reason is slightly subtle: the endpoint errors in the two integrals cancel, since their sum ultimately represents to high order the result of the trapezoid rule applied to the periodic analytic integrand in* (2.9)*. Thus our simple splitting achieves spectral accuracy, needing neither a partition of unity nor end-point quadrature corrections.*

4.3. High frequency scattering example. We now detail the application of the $O(N)$ method just described to a high frequency scattering problem with Dirichlet boundary condition; this corresponds to an acoustically sound-soft obstacle. We choose a complicated (but analytic) boundary $\partial\Omega$ given by the polar Fourier series $f(\theta) = 1 + \sum_{n=1}^{40} a_n\cos(n\theta) + b_n\sin(n\theta)$, with a_n, b_n uniform random in $[-0.04, 0.04]$, and parametrized by θ; see Fig. 4.2(a). We choose $\omega = 250$ such that the obstacle is 100 wavelengths across. For the scattering problem with plane wave $u_{\rm inc}(x) = e^{i\omega\hat d\cdot x}$ incident at angle $\hat d = (\cos -\pi/5, \sin -\pi/5)$, the total potential (physical field) is $u_{\rm tot} = u_{\rm inc} + u$ where the (radiative) scattered potential u solves the exterior Dirichlet problem with boundary data $f = -u_{\rm inc}|_{\partial\Omega}$. For this BVP, as before, we solve the integral equation (4.3) then evaluate u via (4.4).

We find that $N = 9000$ is needed to get 13-digit convergence of the Nyström method for u, as assessed at a variety of exterior points lying outside the bad annular neighborhood $\Omega_{\rm bad}$. This corresponds to an average of 13 nodes per wavelength, although it is as little as 4.6 nodes per wavelength where $|Z'(s)|$ is largest. It takes 73 s to fill the dense $N\times N$ matrix, and a further 14 s to solve the system via GMRES using dense matrix-vector products (needing 95 iterations to exceed a relative residual of 10^{-12}).

REMARK 19. *Here all timings are reported for a System76 laptop with 2.6 GHz Intel i7-3720QM CPU, 16GB of RAM, running MATLAB 2012b [34],* `MPSpack` *version 1.32 [4], and* `FMMLIB2D` *version 1.2 [17]. We use MATLAB's native Bessel functions. For Hankel functions in* (4.6) *and the second sum in* (4.7) *(and their double-layer analogs) we use a MEX interface to* `hank103.f` *[17] for $m = 0, 1$, and upwards recurrence [31, (10.6.1)]. Most operations use*

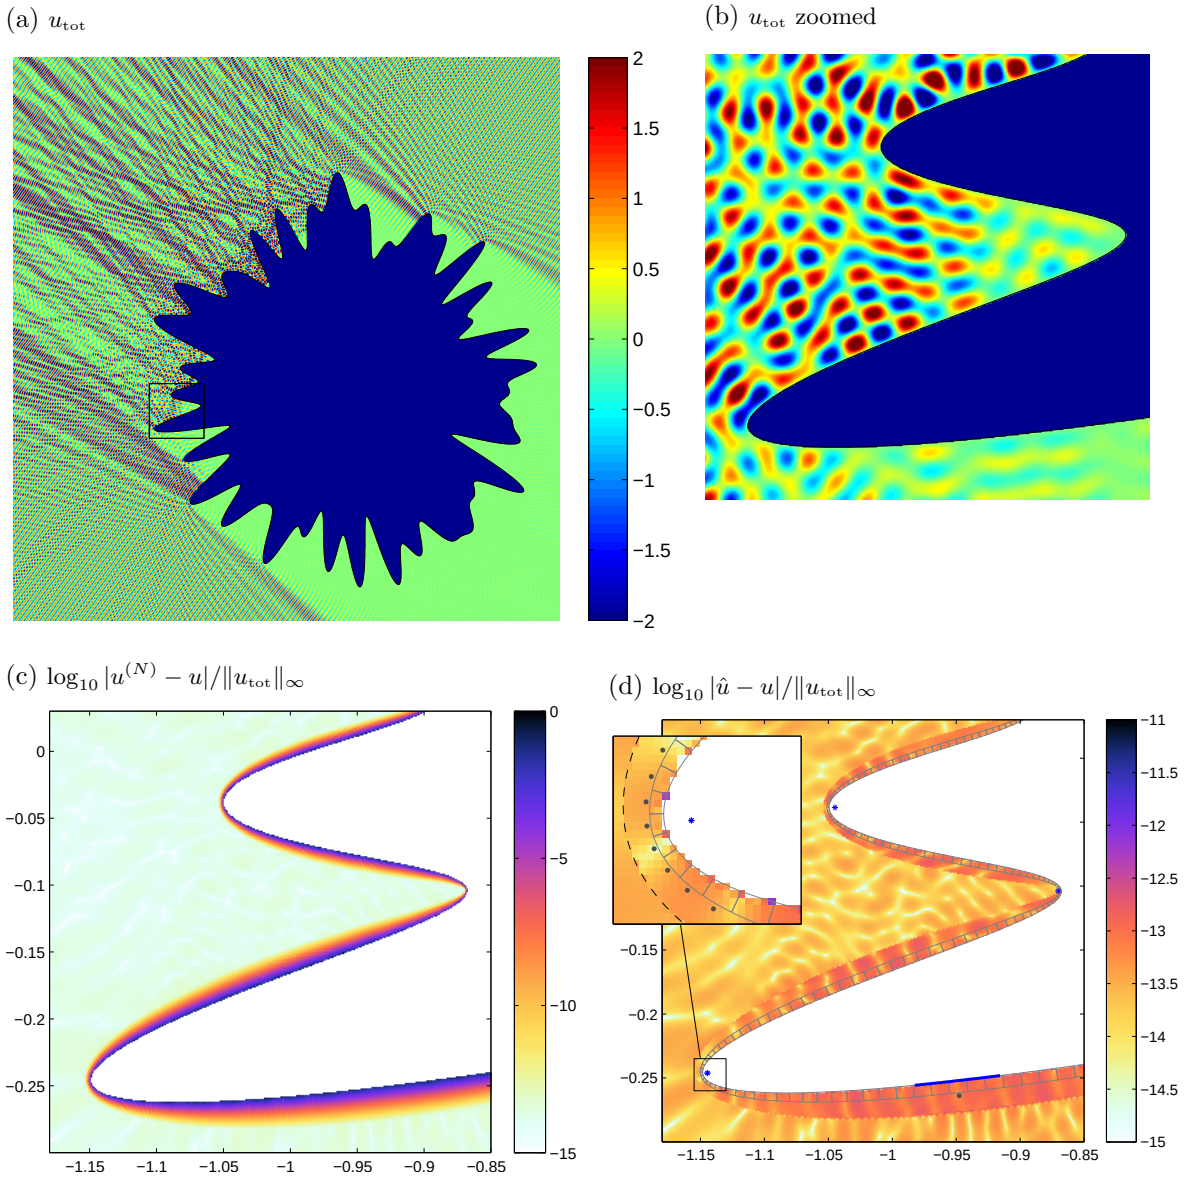

FIG. 4.2. *High frequency sound-soft scattering example, $\omega = 250$ (diameter is 100 wavelengths), needing $N = 9000$ nodes. (a) Total field $u_{tot} = u_{inc} + u$. (b) Zoom of total field for black box shown in (a). (c) Relative error (relative to $\|u_{tot}\|_\infty = 6.2$) in native evaluation $u^{(N)}$ in same region as (b). (d) Relative error in surrogate scheme $\hat{u}$ in same region as (b), with $p = 26$, $\beta = 6$. Boxes are shown in grey, and a single center z_0 (grey dot) with its corresponding "near" set of boundary points (thick blue line), and domain Schwarz singularities ($*$). Note the change in color scale between (c) and (d). The inset in (d) is a zoom of the highly convex region, and also shows all centers (grey dots) and $\Gamma_{\alpha_{bad}}$, the boundary of Ω_{bad} (dotted line).*

only a single core; the only ones which exploit all four cores are `FMMLIB2D` *and the matrix-vector products in GMRES.*

We fix a set of around 8×10^6 evaluation points, namely those on a 3300×3400 grid of spacing 10^{-3} which lie in the exterior of $\partial\Omega$; see Fig. 4.2. Evaluation of u at these targets using

the N native nodes and the FMM takes 24 s, and gives the relative errors in Fig. 4.2(c): large errors are apparent near $\partial\Omega$. We now apply the surrogate close evaluation to the points lying in $\Omega_{\rm bad}$, which number around 2.2×10^5. We set $N_B=\lceil N/3\rceil=3000$, rather more than the value $\lceil N/5\rceil$ recommended before: this helps reduce errors by shrinking the box radii R. We find convergence at around $p=26$, $\beta=6$ (this high p is needed because boxes are up to 0.7 wavelengths in size). A cut-off radius G that does not induce additional error was found to be 1.5 times the maximum width of the bad annular neighborhood in the surrounding 7 boxes; this gives between 84 and 326 "near" fine points, but with a mean of only 107 (note that this is less than the 288 that would be required for three 16-node panels and the same β; see remark 22). With the above parameters we achieve a maximum error in u of 6×10^{-12}, in 28 s computation time. Much of this is spent on direct sums in (4.6) and (4.7) (the fine FMM in (4.7) takes only 1.5 s), as well as geometry and inevitable MATLAB overheads. We found this to be 50 times faster than directly applying the $O(N^2)$ formulae from section 4.1, justifying the utility of our proposed $O(N)$ scheme. Fig. 4.2(d) plots a zoom of the resulting relative error.

REMARK 20 (Reference solution). *In previous examples u was analytically known. In order to assess the error in $\hat{u}$ for this example we compute a reference u in the following expensive fashion. Outside Ω_{bad} we used the FMM from the density $\{\tau_j\}_{j=1}^N$ refined by a factor 10 by trigonometric interpolation. Inside Ω_{bad} we did the same but with a factor 10^3, apart from the few hundred points in the narrow ribbon inside $\Gamma_{10^{-3}\alpha_{bad}}$. For these last points we used 9th-order polynomial extrapolation from points distances $\{2^j(10^{-2}\pi/N)|Z'(s)|\}_{j=0}^9$ along the normal direction. This appears to give around 12 digits; we are not able reliably to get more.*

REMARK 21 (Convex vs concave). *The largest errors occur at the corners of boxes near highly-convex parts of the boundary; we believe that this is due to the nearby interior Schwarz singularity of $\partial\Omega$, which generically induces a singularity in the Helmholtz continuation of the solution u [29] (briefly reviewed in [3, Sec. 3.1]). The latter in turn slows the p-convergence of the expansion. By contrast, concave parts have Schwarz singularities on the physical side of $\partial\Omega$ which thus do not affect u. We believe this explains why lower accuracy is reported at convex (but not concave) locations in QBX [23].*

5. Conclusion and discussion. Firstly, we analyzed the spatial distribution of the error in evaluating Laplace layer potentials using the popular global periodic trapezoid rule. The key tools were generalizations of the Davis theorem, and the annular conformal map between the complex parameter plane and the physical plane. We found (Theorems 3 and 9) that the exponential convergence rate at a point is simply the *imaginary part of its preimage under this map*, a result believed to be new. Error contours are thus given by "imaginary parameter translations" of the boundary; these sweep out an annular neighborhood $\Omega_{\rm bad}$ where errors are unacceptable. Empirically, errors are similar in the Helmholtz case.

Secondly, we devised a surrogate local expansion method for accurate evaluation in $\Omega_{\rm bad}$. Our main analytical result is its exponential convergence (and hence that of QBX [23]) in the analytic Laplace case (Theorems 11 and 15). The scheme can be implemented via the FMM in $O(N)$ time, and generalizes easily to the Helmholtz equation, as we showed in a challenging high-frequency scattering application. Our scheme gives errors close to machine precision, given density values at only the number N of nodes sufficient for the Nyström method to converge. Our experiments also highlighted the need for more understanding of the Nyström exponential convergence rate and of the situations in which the trigonometric interpolant is inferior to the Nyström one.

REMARK 22 (Adaptivity). *We used global quadrature in this study; an adaptive panel-based (composite) underlying quadrature, however, would be preferred in a production code, and would allow refinement at corners. The surrogate scheme is easy to implement with Gaussian panels, needing only one Lagrange interpolation to the fine nodes. Furthermore, the fast scheme*

of section 4.2 becomes simpler: the only FMM needed is the native evaluation $u^{(N)}$, corrected in $O(N)$ effort by refining only the 3 nearest panels. Comparison against the recent panel-based scheme of Helsing [19] would be desired.

We expect the generalization to surfaces in $\mathbb{R}^3$ to be fruitful, both to evaluate close to a surface, and to build singular Nyström quadratures on the surface (generalizing the QBX scheme in $\mathbb{R}^2$ [23]). Since only smooth interpolation, the local expansion, and the addition theorem are needed, this is simpler to implement than most existing high-order schemes [11, 9, 37, 8]. Indeed, since initial submission of the present work, numerical results in $\mathbb{R}^3$ have been encouraging [6], and an analysis of QBX, including the Helmholtz single-layer potential in $\mathbb{R}^3$, has appeared [15]. In $\mathbb{R}^3$, since one can no longer exploit a simple link between Laplace solutions and holomorphic functions, the analysis relies on estimates for spherical harmonics. We suggest that an analysis for the present work in the case of the Helmholtz equation in $\mathbb{R}^2$ should be tractable via Vekua's map from holomorphic functions to Helmholtz solutions [36, 21, 7, 30].

Acknowledgments. The author is very grateful to Hanh Nguyen, via the support of the Women in Science Project (WISP) at Dartmouth College, for testing a Laplace double-layer implementation in the spring of 2011. The author also benefited from discussions with Stephen Langdon, Zydrunas Gimbutas, Leslie Greengard, Andreas Klöckner, Mike O'Neil, and Nick Trefethen. This work is supported through the National Science Foundation via grants DMS-0811005 and DMS-1216656.

Appendix A. Code. While we have not yet released a formal package implementing the methods of this paper, we make available codes that generated the figures here:

`http://math.dartmouth.edu/~ahb/software/closeeval`

These require MATLAB [34], `MPSpack` version at least 1.32 [4], and `FMMLIB2D` version 1.2 [17].

REFERENCES

[1] K. Atkinson, *The numerical solution of integral equations of the second kind*, Cambridge University Press, 1997.

[2] K. Atkinson and Y.-M. Jeon, *Algorithm 788: Automatic boundary integral equation programs for the planar Laplace equation*, ACM Trans. Math. Software, 24 (1998), pp. 395–417.

[3] A. H. Barnett and T. Betcke, *Stability and convergence of the Method of Fundamental Solutions for Helmholtz problems on analytic domains*, J. Comput. Phys., 227 (2008), pp. 7003–7026.

[4] ———, `MPSpack`*: A MATLAB toolbox to solve Helmholtz PDE, wave scattering, and eigenvalue problems*, 2008–2012. `http://code.google.com/p/mpspack/`.

[5] A. H. Barnett and L. Greengard, *A new integral representation for quasi-periodic fields and its application to two-dimensional band structure calculations*, J. Comput. Phys., 229 (2010), pp. 6898–6914.

[6] A. H. Barnett, L. Greengard, and Z. Gimbutas, *Efficient and robust integral equation methods for acoustic scattering from doubly-periodic media in three dimensions*, 2013. in preparation.

[7] T. Betcke, *Computations of Eigenfunctions of Planar Regions*, PhD thesis, Oxford University, UK, 2005.

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

[9] O. P. Bruno and L. A. Kunyansky, *A fast, high-order algorithm for the solution of surface scattering problems: basic implementation, tests, and applications*, J. Comput. Phys., 169 (2001), pp. 80–110.

[10] S. N. Chandler-Wilde, I. G. Graham, S. Langdon, and E. A. Spence, *Numerical-asymptotic boundary integral methods in high-frequency acoustic scattering*, Acta Numer., (2012), pp. 89–305.

[11] D. Colton and R. Kress, *Inverse acoustic and electromagnetic scattering theory*, vol. 93 of Applied Mathematical Sciences, Springer-Verlag, Berlin, second ed., 1998.

[12] P. J. Davis, *On the numerical integration of periodic analytic functions*, in Proceedings of a Symposium on Numerical Approximations, R. E. Langer, ed., University of Wisconsin Press, 1959.

[13] P. J. Davis, *The Schwarz function and its applications*, The Mathematical Association of America, Buffalo, N. Y., 1974. The Carus Mathematical Monographs, No. 17.

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

[15] C. L. Epstein, L. Greengard, and A. Klöckner, *On the convergence of local expansions of layer potentials*, 2012. `arXiv:1212.3868`.

[16] S. D. Gedney, *On deriving a locally corrected Nyström scheme from a quadrature sampled moment method*, IEEE Trans. Antennas Propag., 51 (2003), pp. 2402–2412.
[17] Z. Gimbutas and L. Greengard, *FMMLIB2D, Fortran libraries for fast multipole method in two dimensions*, 2011. `http://www.cims.nyu.edu/cmcl/fmm2dlib/fmm2dlib.html`.
[18] S. Hao, A. H. Barnett, P. G. Martinsson, and P. Young, *High-order accurate Nyström discretization of integral equations with weakly singular kernels on smooth curves in the plane*, 2013. accepted, Adv. Comput. Math., `arxiv:1112.6262v2`.
[19] J. Helsing, *Solving integral equations on piecewise smooth boundaries using the RCIP method: a tutorial*, 2012. preprint, 34 pages, `arXiv:1207.6737v3`.
[20] J. Helsing and R. Ojala, *On the evaluation of layer potentials close to their sources*, J. Comput. Phys., 227 (2008), pp. 2899–2921.
[21] P. Henrici, *A survey of I. N. Vekua's theory of elliptic partial differential equations with analytic coefficients*, Z. Angew. Math. Phys., 8 (1957), pp. 169–203.
[22] D. B. Hunter, *The evaluation of integrals of periodic analytic functions*, BIT Numer. Math., 11 (1971), pp. 175–180.
[23] A. Klöckner, A. H. Barnett, L. Greengard, and M. O'Neil, *Quadrature by expansion: a new method for the evaluation of layer potentials*, 2013.
[24] J. Korevaar, *Book review of "The Schwarz function and its generalization to higher dimensions" by H. S. Shapiro*, Bull. Amer. Math. Soc., 31 (1994), pp. 112–116.
[25] R. Kress, *Boundary integral equations in time-harmonic acoustic scattering*, Mathl. Comput. Modelling, 15 (1991), pp. 229–243.
[26] R. Kress, *Linear Integral Equations*, vol. 82 of Applied Mathematical Sciences, Springer, second ed., 1999.
[27] A. Mayo, *Fast high order accurate solution of Laplace's equation on irregular regions*, SIAM J. Sci. Stat. Comput., 6 (1985), pp. 144–157.
[28] A. McKenney, *An adaptation of the fast multipole method for evaluating layer potentials in two dimensions*, Computers Math. Applic., 31 (1996), pp. 33–57.
[29] R. F. Millar, *Singularities and the Rayleigh hypothesis for solutions to the Helmholtz equation*, IMA J. Appl. Math., 37 (1986), pp. 155–171.
[30] A. Moiola, R. Hiptmair, and I. Perugia, *Vekua theory for the Helmholtz operator*, Z. Angew. Math. Phys., 62 (2011), pp. 779–807.
[31] F. W. J. Olver, D. W. Lozier, R. F. Boisvert, and C. W. Clark, eds., *NIST Handbook of Mathematical Functions*, Cambridge University Press, 2010. `http://dlmf.nist.gov`.
[32] C. Schwab and W. L. Wendland, *On the extraction technique in boundary integral equations*, Math. Comp., 68 (1999), pp. 91–122.
[33] E. M. Stein and R. Shakarchi, *Complex Analysis (Princeton Lectures in Analysis, No. 2)*, Princeton University Press, 2003.
[34] The MathWorks, Inc., *MATLAB software*, Copyright (c) 1984–2012. `http://www.mathworks.com/matlab`.
[35] L. N. Trefethen, *Approximation Theory and Approximation Practice*, SIAM, 2012. `http://www.maths.ox.ac.uk/chebfun/ATAP`.
[36] I. N. Vekua, *Novye metody rezhenija ellipticheskikh uravnenij, (OGIZ, Moscow and Leningrad, 1948); English translation: New methods for solving elliptic equations*, North-Holland, 1967.
[37] L. Ying, G. Biros, and D. Zorin, *A high-order 3D boundary integral equation solver for elliptic PDEs in smooth domains*, J. Comput. Phys., 216 (2006), pp. 247–275.