

A fast algorithm with error bounds for Quadrature by Expansion

Matt Wala*, Andreas Klöckner

Department of Computer Science, University of Illinois at Urbana-Champaign, 201 North Goodwin Ave, Urbana, IL 61801, United States



ARTICLE INFO

Article history:

Received 23 January 2018
 Received in revised form 30 April 2018
 Accepted 4 May 2018
 Available online xxxx

Keywords:

Fast algorithms
 Fast multipole method
 Integral equations
 Quadrature
 Singular integrals

ABSTRACT

Quadrature by Expansion (QBX) is a quadrature method for approximating the value of the singular integrals encountered in the evaluation of layer potentials. It exploits the smoothness of the layer potential by forming locally-valid expansions which are then evaluated to compute the near or on-surface value of the potential. Recent work towards coupling of a Fast Multipole Method (FMM) to QBX yielded a first step towards the rapid evaluation of such integrals (and the solution of related integral equations), albeit with only empirically understood error behavior. In this paper, we improve upon this approach with a modified algorithm for which we give a comprehensive analysis of error and cost in the case of the Laplace equation in two dimensions. For the same levels of (user-specified) accuracy, the new algorithm empirically has cost-per-accuracy comparable to prior approaches. We provide experimental results to demonstrate scalability and numerical accuracy.

© 2018 Elsevier Inc. All rights reserved.

1. Introduction

Integral equation methods for the solution of boundary value problems of partial differential equations offer a number of numerically attractive properties, including boundary-only discretizations for homogeneous problems, seamless treatment of exterior problems, and mesh-independent conditioning. However, their effective numerical realization presents a number of technical challenges. A key prerequisite for the use of these methods is the scalable and accurate evaluation of layer potentials on, near, and away from the source layer. This in turn involves *singular and near-singular quadrature* and so-called *fast algorithms* (like the Fast Multipole Method) to facilitate the evaluation of $O(N^2)$ interactions in linear or near-linear time. To maintain accuracy and efficiency, both aspects need to be well-integrated, and, as a unit, have well-understood error behavior.

The single layer and double layer potential integral operators \mathcal{S} , \mathcal{D} for the Laplace equation with boundary density function $\mu(y) : \Gamma \rightarrow \mathbb{R}$ are defined as

$$\mathcal{S}\mu(x) = -\frac{1}{2\pi} \int_{\Gamma} \mu(y) \log |y - x| ds(y), \quad (1)$$

$$\mathcal{D}\mu(x) = -\frac{1}{2\pi} \int_{\Gamma} \mu(y) \hat{n} \cdot \nabla_y \log |y - x| ds(y). \quad (2)$$

* Corresponding author.

E-mail addresses: wala1@illinois.edu (M. Wala), andreask@illinois.edu (A. Klöckner).

The capability of evaluating such layer potentials can be used for the solution of homogeneous PDE boundary value problems. We demonstrate this by the following example. Consider, for specificity, the exterior Neumann problem in two dimensions for the Laplace equation:

$$\begin{aligned} \Delta u(x) &= 0 \quad (x \in \mathbb{R}^2 \setminus \Omega), \\ (\hat{n}(x) \cdot \nabla u(y)) &\rightarrow g \quad (x \in \partial\Omega, y \rightarrow x_+), \\ u(x) &\rightarrow 0 \quad (x \rightarrow \infty), \end{aligned} \quad (3)$$

where Ω is a bounded domain with a smooth boundary, $\hat{n}(x)$ is the unit normal to $\partial\Omega$ at x , and $\lim_{y \rightarrow x_+}$ denotes a limit approaching the boundary from the exterior of Ω . Then, by choosing $\Gamma = \partial\Omega$ and representing the solution u in terms of a single layer potential $u(x) = S\mu(x)$ with an unknown density function μ , we obtain that the Laplace PDE and the far-field boundary condition are satisfied by u . The remaining Neumann boundary condition becomes, by way of the well-known jump relations for layer potentials (see [35, Theorem 6.28]) the integral equation of the second kind

$$-\frac{\mu}{2} + S'\mu = g. \quad (4)$$

The boundary Γ and density μ may then be discretized using piecewise polynomials and, using the action of the normal derivative S' of S as supplied by our method below, solved for the unknown density μ . Once μ is known, the representation of u in terms of the single-layer potential (1) may be evaluated anywhere in $\mathbb{R}^2 \setminus \Omega$, again using the method described below, to obtain the sought solution u of the boundary value problem.

Quadrature by Expansion ('QBX') is an approach to singular and near-singular quadrature in the setting of layer potential evaluation that is kernel- and dimension-independent. QBX makes use of the fact that the layer potential is analytic and accurately resolved via regular quadrature methods for smooth functions (such as Gaussian quadrature) when the target point is sufficiently far away from the surface. Accuracy in the near or on-surface regime is recovered through analytic continuation by ways of, e.g. a Taylor ('local') expansion of the potential about a center in the well-resolved regime.

Since the earliest days of the numerical use of integral equation methods [2], acceleration of the otherwise quadratic (in the number of degrees of freedom) runtime of the associated matrix-vector products has been a concern. If no acceleration is used, the integrals of (1) and (2) must be evaluated from scratch for each of $O(N)$ target points, where each such integration involves evaluation of the integrand at $O(N)$ quadrature nodes. Acceleration approaches range from custom methods based on the hierarchical decomposition of curves [48] to evaluation methods [13] based on Barnes–Hut-style [3] tree codes. In these methods, the coexistence of quadrature and acceleration is a pervasive concern. When used to solve PDE BVPs as described above, layer potential evaluation may be viewed as two distinct tasks: First, evaluation of the potential on the surface itself as required for the solution of the integral equation to obtain the density, and, second, evaluation of the potential in the volume to obtain the actual solution of the boundary value problem. Kussmaul–Martensen quadrature [36,39] as a representative of singularity subtraction techniques, the polar coordinate transform [35] as a representative of singularity cancellation techniques, or Generalized Gaussian quadrature [55] are examples of a quadrature scheme only suited to the evaluation of (weakly) singular on-surface layer potential integrals.

Meanwhile, the evaluation of layer potentials in the volume is in principle straightforward as no singular integrals are involved, for example by adaptive quadrature. Careful management of accuracy that avoids dramatic performance degradation however is less straightforward to achieve [6,27]. QBX, as a quadrature scheme, unifies on-surface and off-surface evaluation [4] with only minor accommodations, and we would like to retain this feature of the method in its accelerated version.

Beyond this overview, we will not attempt to review the vast literature on singular quadrature (e.g. [1,5,7,9,10,12,15,18–20,23–26,29–31,37,38,40,50,51,53,56]). We instead refer the reader to [34] for a rough overview. Here, we will continue by focusing instead on approaches to combining singular quadrature with a fast algorithm into a single scheme.

The use of hierarchically-based fast algorithms for the evaluation of layer potentials considerably predates the Fast Multipole Method itself, such as Rokhlin's early work aimed squarely at accelerated quadrature [48]. Within the framework of the Fast Multipole Method, quadrature methods that require special treatment of near interactions often proceed by replacing the direct interactions (of 'List 1' in the FMM) with their own procedure. Unfortunately, no guarantees of geometric separation between source and target can be derived from membership in List 1, and so methods requiring this may subtract out unwanted interactions already mediated by the FMM, at an additional computational cost.

Within the realm of the acceleration of QBX, early work [17,34] remarked on the apparent ease with which QBX might be integrated into Fast Multipole-type algorithms, by slightly modifying the algorithm to yield local expansions containing contributions from the entire source geometry in what has come to be called *global QBX*. First steps towards the realization of such an integration were soon made, first in unpublished work. These early attempts were plagued by uncontrolled and poorly-understood accuracy issues. An initial approach to recovering accuracy through an increase of the FMM order [44,46] succeeded, but provided only empirical evidence for its accuracy. We refer to this order-increase scheme as the 'conventional QBX FMM' throughout this text. QBKIX [47] (Quadrature by Kernel-Independent Expansion) emerged as a related, global QBX-based numerical method that is built upon the machinery of kernel-independent Fast Multipole Methods [57].

QBX may also be operated as a local correction applied in the near-field of an FMM, as described above. These schemes, broadly classified as *local QBX* [44,45], are algorithmically much simpler than global QBX since a fast algorithm for point

potential evaluation may be used largely unmodified. However, to allow the transition between QBX-regularized near-field and point-potential far-field to occur without loss of accuracy, schemes based on local QBX generically require very high QBX expansion orders, which in turn requires a large amount of oversampling. Recent work [52] has been seeking to mitigate the computational cost of this effect.

This contribution is concerned with presenting a version of a global QBX-based FMM coupling that provides *rigorous error bounds*, thus providing one approach for a compatible coupling of a singular quadrature rule with acceleration. To accomplish this, we make substantial modifications to the Fast Multipole Algorithm itself. We list these comprehensively in Section 4. Some versions of some of these modifications have been discussed in the literature, though in contexts unrelated to layer potential evaluation. For example, we restrict the set of allowable multipole-to-local translations to be between boxes separated by a distance of at least twice their own size. Greengard already discusses the possibility of such a modification, in the context of a three-dimensional generalization of the FMM, in his thesis work [21]. We also introduce the notion of sizing for targets, to accommodate the unique requirements of global QBX centers. A related need emerged in the chemical physics community, where Coulomb interactions between (extent-bearing) ‘clouds’ of charge need to be evaluated [54], though the algorithm ultimately constructed is substantially different from ours.

We refer to our algorithm as *GIGAQBX* (for ‘Geometric Global Accelerated QBX’), to contrast with prior versions of the scheme. In this paper, we take the point of view that the cumulative error in an accelerated QBX scheme effectively splits into three additive components:

$$|\text{accelerated QBX error}| \leq |\text{truncation error}| + |\text{quadrature error}| + |\text{FMM error}|. \quad (5)$$

Here, *truncation error* refers to the analytical truncation error, and *quadrature error* is error in evaluating the QBX-regularized integral using quadrature. The *FMM error* refers to the FMM’s achieved accuracy in approximating the output of unaccelerated QBX (see Section 2 for details). We choose to split error contributions in this way, rather than relying on direct approximation of the layer potential by the FMM. This interpretation allows us to rely on the existing body of work establishing bounds on the truncation and quadrature components of the error (e.g. [17,32,33,46]). A complicating factor for the FMM error analysis is that traditional FMM error estimates apply only to the approximation of potentials at point-shape targets, whereas our version of these estimates must account for the approximation of a *QBX local expansion* and its accuracy when evaluated as an approximation the potential. The FMM error in this setting is not well-studied.

2. Overview and motivation

In this section, after reviewing the basic operating principle and convergence theory of QBX in Section 2.1, we summarize recent progress made towards making global, unaccelerated QBX geometrically robust [46] which we continue to leverage (see Section 2.2.1). We summarize the approach taken toward acceleration in the same article for comparison and to contrast with our version in Section 2.2. We assume that the reader has some familiarity with Fast Multipole methods but not necessarily the details of [46]. The major difference between our new algorithm and the algorithm of [46] occurs in the handling of QBX centers within the FMM. The motivation for how we handle QBX centers is detailed in Section 2.4.

FMMs conventionally evaluate potentials from point sources. In the remainder of this paper, we will often refer to that family of algorithms as *point FMMs*, especially to distinguish from our intended task, which is the evaluation of potentials originating from a layer Γ , i.e. layer potentials.

2.1. Quadrature by expansion

We consider the problem of evaluating the Laplace single layer potential over a smooth simple closed curve Γ in a region that is close to the boundary itself (which may include points on the boundary). This is done solely to simplify and focus the discussion; as with many other parts of the QBX literature, our approach generalizes straightforwardly to many more layer potentials and kernels. We discretize the integrals of (1) and (2) by subdividing Γ into disjoint pieces Γ_k , each parametrized by $\gamma_k : [-1, 1] \rightarrow \Gamma_k$, and by using Gaussian quadrature of a fixed node count p_{quad} on each piece. Because of the nearly or weakly singular behavior of the integrand, Gaussian quadrature will not yield accurate results when the target point x is near the boundary Γ . The precise width of this neighborhood depends on both the quadrature node count p_{quad} and the panel length h but it is usually on the order of a panel length. For a thorough experimental account of this phenomenon, see the first section of [34]. By contrast, when x is far enough away from the boundary, the integrand is smooth and the integral may be easily evaluated to high accuracy with a smooth quadrature rule.

Roughly speaking, QBX extends the ‘high accuracy’ quadrature region by making use of the fact that the potential is *analytic* in the set $\mathbb{R}^2 \setminus \Gamma$. QBX proceeds in two stages:

First stage. First, Taylor expansion centers are placed away from the boundary Γ in the high accuracy region. Let c be the expansion center associated with a target/evaluation point x . For simplicity, in the remainder of the paper we will identify

\mathbb{R}^2 with \mathbb{C} and make use of the complex logarithm, which satisfies $\text{Re} \log(z) = \log|z|$ for all $z \in \mathbb{C} \setminus \{0\}$. We truncate the Taylor expansion and write the p_{QBX} -th order expansion as follows, interchanging the order of summation and integration:

$$S\mu(x) \approx -\frac{1}{2\pi} \text{Re} \sum_{n=0}^{p_{\text{QBX}}} \left[\frac{1}{n!} \int_{\Gamma} \mu(y) \left(\frac{d^n}{dc^n} \log(y-c) \right) ds(y) \right] (x-c)^n.$$

Second stage. The occurring (nonsingular) integrals are discretized using (in our case Gaussian) quadrature. Let $\{y_i\}_{i=1}^N$ be the set of quadrature nodes with weights and arc length elements $\{w_i\}_{i=1}^N$. Then the formula for QBX is

$$S\mu(x) \approx S_{\text{QBX}(p_{\text{QBX}}, N)}\mu(x) := -\frac{1}{2\pi} \text{Re} \sum_{n=0}^{p_{\text{QBX}}} \left[\frac{1}{n!} Q_N \left\{ \int_{\Gamma} \mu(y) \left(\frac{d^n}{dc^n} \log(y-c) \right) ds(y) \right\} \right] (x-c)^n, \tag{6}$$

where Q_N denotes the approximate computation of an integral of a smooth function through the application of a quadrature rule:

$$Q_N \left\{ \int_{\Gamma} f(y) ds(y) \right\} = \sum_{i=1}^N w_i f(y_i).$$

For purposes of solving integral equations, one is mostly interested in the case where $x \in \Gamma$. QBX handles this case with high order accuracy, provided that μ and Γ are smooth [17]. More precisely, we state the error estimates in the case where the quadrature is composite Gaussian quadrature over panels of equal length h . Let r be the expansion radius for the Taylor expansion. Let $\bar{B}(c, r)$ denote the closed disk of radius r centered at c . Assume that $x \in (\bar{B}(c, r) \cap \Gamma)$ and that $\Gamma \setminus \{x\}$ does not intersect $\bar{B}(c, r)$. Let p_{quad} denote the node count of the Gaussian quadrature. Then it can be shown that the error components scale as

$$\begin{aligned} |\text{truncation error}| &\leq \frac{1}{2\pi} \left| \sum_{n=p_{\text{QBX}}+1}^{\infty} \left[\frac{1}{n!} \int_{\Gamma} \mu(y) \left(\frac{d^n}{dc^n} \log(y-c) \right) ds(y) \right] (x-c)^n \right| \\ &\leq C_1(p_{\text{QBX}}, \Gamma) \|\mu\|_{C^p(\Gamma)} r^{p_{\text{QBX}}+1} \log \frac{1}{r}, \end{aligned} \tag{7}$$

$$\begin{aligned} |\text{quadrature error}| &\leq \frac{1}{2\pi} \left| \sum_{n=0}^{p_{\text{QBX}}} \left[\frac{1}{n!} \left(\int_{\Gamma} - Q_N \left\{ \int_{\Gamma} \right\} \right) \mu(y) \left(\frac{d^n}{dc^n} \log(y-c) \right) ds(y) \right] (x-c)^n \right| \\ &\leq C_2(p_{\text{quad}}, p_{\text{QBX}}, \Gamma) \left(\frac{h}{4r} \right)^{2p_{\text{quad}}} \|\mu\|_{C^{2p_{\text{quad}}}(\Gamma)}. \end{aligned} \tag{8}$$

If we choose $r = \alpha h$ with a proportionality factor α so that $h/(4r) < 1$, we obtain a scheme that is accurate of order $p_{\text{QBX}} + 1$ in the mesh spacing h up to controlled precision $1/(4\alpha)^{2p_{\text{quad}}}$. See [17] for more details about the error estimates for QBX. This choice reveals p_{quad} as a free parameter that governs the controlled precision term $1/(4\alpha)^{2p_{\text{quad}}}$, where it is mostly the requirement of resolving the high derivatives of the kernel occurring in (6) that governs the magnitude of p_{quad} .

2.2. Ensuring accuracy

The assumptions required by QBX convergence theory will not necessarily be met by input geometries supplied by a user. In addition, quadrature resolution and placement of centers need to be carefully controlled so as to retain bounds on quadrature and truncation error.

An efficient algorithm that accomplishes this is the main contribution of [46]. We briefly and informally review this procedure.

2.2.1. Preprocessing the geometry

To ensure the accuracy of QBX (independently of any fast algorithm), we check for the following situations which may result in inaccuracy:

- When a QBX disk intersects the source curve at more points than the target point. This comes from the analytical requirement that the QBX expansion disk must not be obstructed by any piece of the source curve.

- *When resolution and separation of source geometry from a QBX expansion center do not guarantee accurate coefficient computation.* This issue occurs in the presence of varying panel sizes. Depending on the quadrature rule and the panel size, the quadrature contribution from the singular integrand may not be resolved adequately if the source geometry is too close to the QBX center. This happens when a large panel is close to a small panel: quadrature from the large panel may not adequately resolve the integrand when evaluated at a QBX center near the small panel.

Because of the assumption of smooth, non-self-intersecting source geometry Γ , both of these sources of error can be controlled by *iterative refinement*, such as by repeated bisection of panels. In the first case, bisecting the source panel of the ‘disturbed’ center, by the proportionality $h = \alpha r$, draws the expansion center closer to the panel and hence, if applied often enough, the resulting expansion disk will eventually avoid the conflicting geometry. In the second case, the offending source panels can be iteratively refined to grow the region in which accurate coefficient quadrature is achieved to include the target center. Both of these checks can be efficiently implemented using a mechanism termed *area queries* in [46].

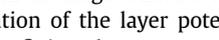
Once the geometry has been processed to ensure that there are no obstructions to accurate quadrature and control of truncation error, it remains to choose a quadrature node count. The resulting p_{quad} typically exceeds what might be required to resolve the density and the geometry by some factor. Especially when solving integral equations, it is thus natural to maintain density and geometry at a suitable resolution and ‘upsample’ them to p_{quad} nodes for QBX computation. p_{quad} can be empirically estimated, as in Table 1 of [46], or adaptively determined based on analytic knowledge [33].

2.2.2. Placing centers and identifying QBX targets

Evaluation of the potential at a target point requires special treatment (e.g. by QBX) when it is so close to a source panel that the underlying (here, Gaussian) quadrature cannot resolve the integrand for the target. Any such targets need to be identified and associated with a QBX center whose expansion disk contains it. Like the geometry processing tasks of Section 2.2.1, both *identifying* a target that is too close to the source and *finding* a QBX center for the target can be accomplished efficiently using area queries as described in [46].

On the question of center placement, the most straightforward approach (and the one used here) is to place QBX centers at $t \pm \frac{h_k}{2} \hat{n}(t)$, where $t \in \Gamma$ is a target point, h_k is the length of the panel Γ_k

$$h_k := \frac{1}{2} \int_{-1}^1 |\gamma'_k(t)| dt$$

containing t , and \hat{n}_t is the unit normal at t . This ensures that the QBX disks cover most of the area near the (smooth) source curve. There may be gaps in coverage where a target point needing close evaluation does not fall inside a QBX disk; this occurs during volume evaluation of the layer potential for target points  very close to the surface. We currently treat such gaps by refining the source geometry until all required target points are covered, at considerable computational expense. Empirically, we have observed that simply associating targets with QBX centers even if they fall outside their closest QBX center’s expansion disk by a given factor, possibly by up to 20%, leads to little or no observable loss in accuracy, though such use is not covered by theoretical guarantees. Improvements on either strategy are the subject of future investigation.

2.3. Evaluating the potential with an FMM

If QBX is implemented following (6) directly, then a quadratically-scaling computational cost $O(NM)$ is incurred by evaluating the contribution of the N source points at each of the M targets. Making use of the point-discrete form of the quadrature-discretized sources and interpreting the summation in (6) as the evaluation of a local expansion of a potential due to N source charges provides an avenue by which QBX can be accelerated to an $O(N + M)$ scheme by ways of a variant of the Fast Multipole Method (FMM).

To achieve this, our work follows the strategy used in [46] which treats a QBX center as a special kind of FMM target at which the FMM *forms a local expansion* instead of *evaluating* the potential. These expansion center targets participate in the FMM algorithm in much the same way point-shape targets do; in particular, they are each ‘owned’ by a box in the FMM tree. In principle, the only additional capability required of a Taylor global QBX FMM is the accurate evaluation of local expansion coefficients, i.e. higher-order partial derivatives of the potential. The ability to compute one or two derivatives of the potential is a common feature in production FMM codes. Such numerical differentiation is commonly associated with some loss of accuracy. Since the coefficient order (and hence the number of derivatives) in QBX can be substantial, there is the possibility of substantial loss of accuracy. This notion is empirically confirmed in [46], and an empirically-determined increase in the FMM order is suggested as a remedy. The key contribution of this article is to (a) introduce a modified algorithm that does not require an artificial order increase and (b) provide convergence theory that gives concrete error bounds for layer potential approximated in this manner.

Recall that, in an adaptive point-evaluation FMM [11], the potential at a target point is computed from three parts (each of which is often the sum of further contributions), summarized in the first two columns of Table 1. The conventional QBX FMM [46] replaces each of these with a translation to a local expansion as appropriate, summarized in the last column of Table 1.

Table 1

Contributions to the potential in a point-evaluation FMM and the conventional QBX FMM of [46]. ‘mpole’ and ‘local’ refer to multipole and local expansions respectively, and ‘eval.’ and ‘transl.’ refers to expansion evaluation and translation respectively.

Interaction [List]	Point FMM [11]	Conv. QBX FMM [46]
Near neighbor boxes [1/U]	point → point eval.	point → local transl.
Sep. smaller mpoles [3/W]	mpole → point eval.	mpole → local transl.
Far field	local → point eval.	local → local transl.

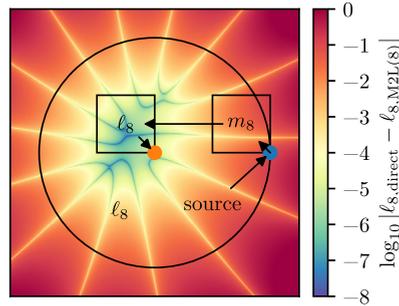


Fig. 1. Accuracy of obtaining a QBX expansion by multipole-to-local translation (vs. direct computation) for an interaction that may be encountered in ‘List 2’ of an FMM. (For interpretation of the colors in the figure(s), the reader is referred to the web version of this article.)

2.4. Accuracy of using translated local expansions for QBX

By *accuracy*, here and in Section 3, we mean the fast algorithm’s accuracy in approximating the terms of (6). This notion is distinct from (though closely related to) the accuracy of the underlying point FMM. The main difference is that the point FMM approximates a potential but the QBX FMM approximates the local expansion of a potential.

A key detail not explicitly considered in the modifications of Table 1 is that QBX expansion disks, unlike target points, have an extent. Since no geometric constraints are imposed, some expansion disks will almost inevitably cross box boundaries. Using the above notion of accuracy, it is easy to imagine that this might have an adverse influence on the accuracy of the computed QBX expansion, owing to either reduced separation from source boxes or larger separation of evaluation points from expansion centers than allowed by FMM separation criteria.

To frame the discussion and give the reader an intuitive sense of this issue, this section presents numerical examples of interactions that *may plausibly occur* in the conventional QBX FMM which lead to large losses in accuracy. We also give an intuitive idea of how our method avoids these errors. We defer a precise statement of the algorithm and a proof of its accuracy to Sections 3 and 4. We consider a number of different types of interactions occurring in an FMM, and we demonstrate the possibility of inaccuracy in each. We consider the potential originating from a single source of unit strength and a single expansion center separated by a reference distance. All expansions have order 8 unless stated otherwise.

Fig. 1 portrays an interaction from a point source (blue, right) to a QBX expansion center (orange, middle). A black circle indicates the size of the QBX expansion disk about the center. In a typical usage scenario of QBX, the source may contribute to an expansion of the layer potential about the center, which is then evaluated back at the source, for, say, the computation of the one-sided limit of the layer potential at the source. We now consider an evaluation scenario in which FMM acceleration mediates this interaction through a multipole (shown as ‘ m_s ’) and a local expansion (shown as ‘ l_s ’). While we have chosen this placement to be ‘adversarial’ (i.e. to lead to a large loss in accuracy), the scenario is permissible under the rules of the conventional QBX FMM, since all required conditions are met: The source point lies inside the box for which the multipole expansion is formed, the source and target box are ‘well-separated’, and the target QBX expansion center lies inside of the target box.

A sufficient criterion to assure that the accelerated and unaccelerated version of the scheme yield the same potential is that the QBX local expansions computed directly and by ways of intermediate expansions evaluate to the same potential, up to FMM accuracy. The false-color plot of Fig. 1 shows the magnitude of the difference between those two expansions in the scenario. In the point FMM, a coarse estimate of multipole-to-local (‘M2L’ for short) accuracy for eighth-order expansions evaluated within the target box gives

$$\left(\frac{\text{dist}(\text{box center, furthest target})}{\text{dist}(\text{box center, closest source})} \right)^{p+1} \leq \left(\frac{\sqrt{2}}{4 - \sqrt{2}} \right)^9 \approx 4.4 \cdot 10^{-3}. \quad (9)$$

We have not yet demonstrated the applicability of such an estimate to the QBX case (cf. Lemma 5), but the data show that, for evaluation when the source point is also the target, the expansion computed through intervening multipole and local expansions misses this accuracy goal by a noticeable margin.

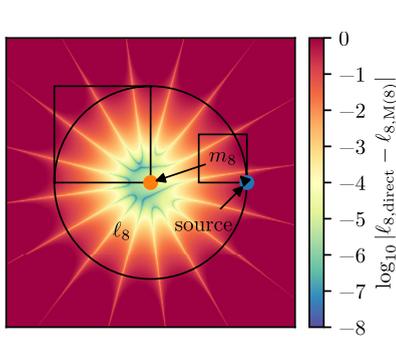


Fig. 2. Accuracy of obtaining a QBX expansion by multipole-to-target translation (vs. direct computation) for an interaction that may be encountered in ‘List 3’ of an FMM.

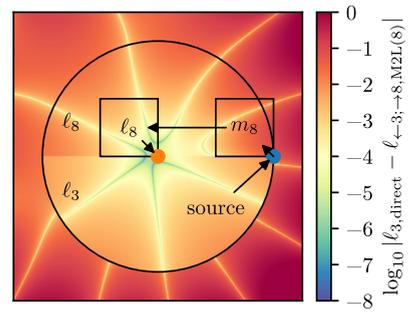


Fig. 3. Accuracy of obtaining a QBX expansion by multipole-to-local translation (vs. direct computation) for an interaction that may be encountered in ‘List 2’ of an FMM. In this experiment, the QBX order is lower than the order of the intermediate multipole and local (FMM) expansions.

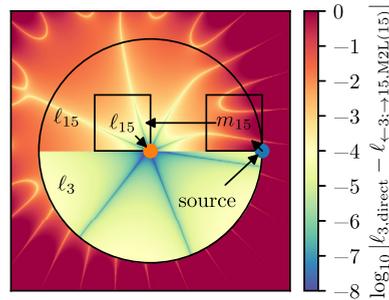


Fig. 4. Accuracy of obtaining a QBX expansion by multipole-to-local translation (vs. direct computation) for an interaction that may be encountered in ‘List 2’ of an FMM. In this experiment, the QBX order is lower than the order of the intermediate (FMM) multipole and local expansions. Compared with Fig. 3, this experiment explores the effect of increasing the order of the intermediate (FMM) multipole and local expansions.

Analogously inaccurate approximation of QBX interactions for evaluation back at the source may occur not just in multipole-to-local, but also in other types of interactions in the FMM. In Fig. 2, the false-color plot again shows the magnitude of the difference between the potential obtained from evaluating the QBX expansion computed directly from the source and the QBX expansion obtained indirectly by ways of intermediate expansions, this time from a multipole expansion associated with a small box containing the source point directly to QBX expansion center ‘target’ within a larger target box. Such an interaction may occur through *List 3* in the conventional QBX FMM. Fig. 5 similarly shows a source-to-local interaction of the type one might encounter in a *List 4* of the conventional QBX FMM.

The experiments described so far still paint an incomplete picture of the translation process involved in accelerating QBX. For a more complete understanding, consider that the FMM order affects the accuracy of the potential by ways of the multipole-to-local error, in a form like (9), whereas the QBX order controls h -convergence as in (7) up to controlled precision, as in (8). As a result, the QBX order is typically lower than the FMM order, a fact that is not reflected in our experiments thus far. Fig. 3 shows the result of a first experiment that takes this into consideration. Denote the lower-order QBX expansion obtained directly from the source, here of order q , by $\ell_{q,\text{direct}}$. Further, denote the local expansion of order q centered at the same location, obtained through a multipole-to-local chain of order p as pictured by $\ell_{q,M2L(p)}$. Then the top and outer parts of Fig. 3 show $|\ell_{3,\text{direct}} - \ell_{8,M2L(8)}|$, while the bottom part shows $|\ell_{3,\text{direct}} - \ell_{3,M2L(8)}|$. A first observation from this experiment is that $\ell_{3,M2L(8)}$ better approximates $\ell_{3,\text{direct}}$ than $\ell_{8,M2L(8)}$. While atypical from the point of view of conventional M2L error estimation theory (where high order entails higher accuracy), this is also not entirely surprising, as the translation chain is bound to approximate lower-order coefficients more accurately than higher-order ones. In other words, simply truncating $\ell_{8,M2L(8)}$ leads to higher accuracy. While this argument is intuitively immediately appealing, we are not aware of any estimates that would aid in quantifying the effect. Next, we observed in our earlier experiments that M2L-mediated expansions did not achieve ‘conventional’ M2L accuracy for QBX evaluation at the source point. Based on the results of our latest experiment, we still cannot confidently assert that these tolerances are being met here. We can however predict that, as long as the order of the final M2L-mediated (‘QBX’) local expansion is being kept constant, increasing the intervening M2L orders should improve the approximation of the individual coefficients of $\ell_{3,\text{direct}}$. This prediction is borne out by the experiment of Fig. 4 which analogously to Fig. 3 compares $|\ell_{3,\text{direct}} - \ell_{15,M2L(15)}|$ with $|\ell_{3,\text{direct}} - \ell_{3,M2L(15)}|$.

This is the basic mechanism by which the conventional QBX FMM of [46] achieves accuracy. Looking ahead to the results obtained, Table 4 summarizes the accuracy achieved in a verification of Green’s formula $\mathcal{S}(\partial_n u) - \mathcal{D}(u) = u/2$ for a harmonic u across a range of FMM orders p_{FMM} and p_{QBX} for a reasonably simple test geometry by that scheme. We find these results unsatisfactory for two reasons: First, the evidence supporting the attained accuracy, while surprisingly

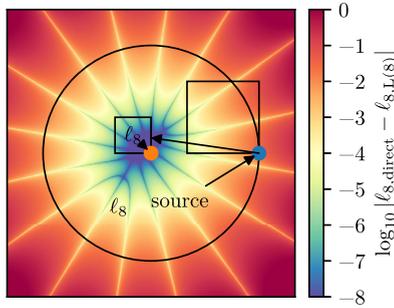


Fig. 5. Accuracy experiments for QBX-FMM coupling with interactions as found in an FMM. QBX FMM error for a point-to-local ('List 4') interaction.

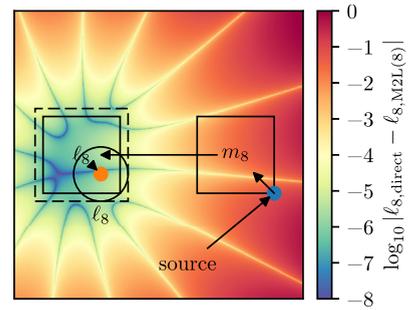


Fig. 6. Accuracy experiments for QBX-FMM coupling with interactions as found in an FMM. QBX FMM error for a List 2 interaction, with the expansion confined to a small region extending beyond the box containing the center.

robust across geometries in practice, is empirical. Second, considering the results shown in Table 4 for high QBX orders p_{QBX} (and thus high relative accuracies), the required FMM order quickly becomes very large if high accuracy is desired. In the remainder of this contribution, we pursue a different strategy that addresses both of these issues.

2.5. Improved accuracy through a geometric criterion

Perhaps the foremost problem with the above translation schemes in the context of the conventional error estimates for multipole and local expansions is that they permit—for QBX purposes—inaccurate near-field contributions mediated through multipole and local expansions to enter the QBX local expansion. As shown, increasing the order of those expansions can (empirically) mitigate this circumstance. We prefer to rework the fast algorithm so as to prevent those contributions in the first place. Roughly speaking, this requires separating contributions in the 'far field' of the QBX local expansion disk (not just the box) from the inaccurate 'near-field' ones. As we will show, the far-field ones may be computed with intervening translations without endangering accuracy. We have shown above that intervening translations can considerably damage accuracy for the near-field evaluation. It is useful to realize that the Fast Multipole algorithm already contains a mechanism for handling this type of issue; its chief purpose, after all, is to separate a far-field that is easily approximated from a near-field that does not tolerate approximation. It is thus natural to seek to broaden the FMM's notion of a near-field so as to respect the needs of QBX. Wishing to avoid the scenarios that led to loss of accuracy above, we begin with the coarse notion that we wish to avoid expansion-'accelerated' contributions to QBX local expansions which would not meet the same accuracy target as the FMM itself.

A first algorithmic variant that provides sufficiently strong accuracy guarantees is nearly immediate: One may require that the entirety of a QBX expansion disk be contained inside some FMM (potentially non-leaf-level) target box. From there, it is at least intuitively plausible that the conventional FMM interaction patterns and their associated error estimates might generalize to guarantee accurate multipole-/local-mediated far-field contributions to QBX local expansions inside each box. This relatively simple generalization of the FMM already represents a somewhat large algorithmic change: While in the original FMM, target particles can only occur in leaf boxes, confining a center to its box entails that *non-leaf* boxes may also contain QBX expansion disk targets. We call this restriction a *target confinement rule*, named this way because the target QBX disks are confined to the inside of a box. This modification achieves the desired accuracy (rigorously, as we will show in Section 3). Unfortunately, it is unsuitable in practice because it no longer has linear complexity—neither in theory nor in practice. In fact, the restriction may lead QBX expansion disk targets that overlap the boundaries of boxes near the root of the tree to exist at near-root levels. Such QBX expansion disks, of which there could be a large number, interact with nearly the entire geometry without the benefit of multipole acceleration.

A second algorithmic variant that remedies this is again virtually immediate: Let QBX expansion disks with a center inside a target box not be confined to the strict extent of their containing box, but instead allow them to extend beyond it by a constant factor of the box size, called the *target confinement factor* ('TCF'). Intuitively, this ensures that each expansion disk may propagate down the tree (away from the root) until it reaches a box whose size is commensurate with the disk's own diameter. It is perhaps plausible that such a scheme might no longer be subject to superlinear complexity. However, the price for the lower cost is that obtaining guaranteed accuracy requires a more complicated algorithm than in the previous case, which we may describe as having a target confinement factor of zero.

Fig. 6 provides a graphical representation. The larger target confinement region is shown with a dashed line. It extends beyond the boundaries of the FMM box, which are drawn using a solid line. The figure also shows a computational experiment analogous to the one of Fig. 1 demonstrating that, at least in this situation, mediated expansions accurately approximate the directly-obtained QBX expansion if they are contained in the target confinement region.

As we will see, modifying the FMM algorithm to retain its benign characteristics in terms of accuracy and cost under this modification presents a considerable set of challenges. At the heart of this modification process is the choice of the target confinement factor, which represents the main control point for the cost-accuracy trade-off inherent in our algorithm. To

illustrate: a larger TCF may result in worse convergence factors for nearly all FMM interactions, while yielding smaller cost by allowing QBX expansion disks to settle closer to the leaves of the tree. To obtain good convergence factors in two and three dimensions, we have chosen to modify the basic notion of ‘well-separated-ness’ inherent in the FMM, from, roughly speaking, ‘1-away’ to ‘2-away’, similar to the three-dimensional FMM of [21]. Similarly, we had to considerably rework the criteria for interaction lists of well-separated smaller and bigger boxes (‘List 3’ and ‘List 4’).

The purpose of the remainder of this paper is to make rigorous the heuristic arguments of the previous paragraphs. In Section 3, we present a novel, more versatile version of the expansion translation error estimates of [22] that allow us to estimate the accuracy achieved by a chain of translation operators in the presence of varying expansion orders. In Section 4, we precisely state our algorithm and provide a complexity analysis that provides a set of benign conditions under which linear complexity is retained. We further point out how the analysis of Section 3 can be used to understand the accuracy of the full algorithm. We close with a comprehensive set of accuracy and complexity experiments in Section 5.

3. Analysis on translation operators with varying orders

To arrive at a better, more quantitative understanding of the accuracy of the translation chains examined in the previous section, we prove new results that explore the accuracy behavior of FMM translation operators when the expansion order varies throughout the chain. We are particularly interested in the effect of truncation of an ‘upstream’ expansion on the accuracy of ‘downstream’ expansion coefficients. The lemmas proven below, while useful in our analysis of the GI-GAQBX FMM, are entirely independent of our particular usage scenario and may prove useful in other settings. For simplicity and conciseness, we prove these results in the setting of the Laplace equation in two dimensions.

3.1. Analytical preliminaries

First, we recall standard facts regarding multipole and local expansions. For proofs of these facts, we refer the reader to [22].

The *multipole expansion* centered at the origin due to a unit strength source s_1 at z_0 takes the form

$$\phi(z) = \log(z - z_0) = a_0 \log(z) + \sum_{k=1}^{\infty} \frac{a_k}{z^k} \tag{10}$$

with $a_0 = 1$ and $a_k = -z_0^k/k$ for $k > 0$. This series converges for $|z| > |z_0|$, where $R = |z_0|$ is called the radius of the multipole expansion. The multipole expansion (10) can also be truncated to $(p + 1)$ terms, which we term a *p-th order expansion*.

Two important operations on the (truncated or non-truncated) expansion (10) are (1) shifting the center of the expansion and (2) conversion to a local expansion. The center of the multipole expansion (10) may be shifted to a new center y , obtaining another multipole expansion, with coefficients $(\alpha_k)_{k=0}^{\infty}$, given by

$$\alpha_m = \begin{cases} a_0 & m = 0, \\ -\frac{a_0(-y)^m}{m} + \sum_{k=1}^m \binom{m-1}{k-1} a_k (-y)^{m-k} & m > 0. \end{cases} \tag{11}$$

The resulting expansion converges for $|z| > R + |y|$. The multipole expansion (10) can also be converted to a local (Taylor) expansion, centered at y for $|y| > R$, with coefficients $(\beta_k)_{k=0}^{\infty}$ given by

$$\beta_m = \begin{cases} a_0 \log(y) + \sum_{k=1}^{\infty} \frac{a_k}{y^k} & m = 0, \\ (-1)^m \left(\frac{a_0}{m y^m} + \sum_{k=1}^{\infty} \binom{m+k-1}{k-1} \frac{a_k}{y^{m+k}} \right) & m > 0. \end{cases} \tag{12}$$

The series $\sum_{k=0}^{\infty} \beta_k (z - y)^k$ converges when $|z - y| < R - |y|$.

The local expansion centered at the origin of the potential due to a unit strength source at z_0 is the Taylor expansion

$$\phi(z) = \log(z - z_0) = \sum_{k=0}^{\infty} b_k z^k. \tag{13}$$

This converges for $|z| < |z_0|$. The coefficients $(b_k)_{k=0}^{\infty}$ are given by $b_0 = \log(-z_0)$ and $b_m = -1/(m z_0^m)$ for $m > 0$. The main operation on expansions of the type (13) of importance to this discussion is shifting the center of the local expansion. The center of a p -th order local expansion of the form (13) may be shifted to another center y , with new coefficients $(\beta_m)_{m=1}^p$, given by

$$\beta_m = \sum_{k=m}^p \binom{k}{m} b_k y^{k-m}. \tag{14}$$

3.2. Error estimates for chained translations

Given our intended usage pattern in accelerated QBX, we are interested in the following types of translation chains:

1. Source \rightarrow Multipole(p) \rightarrow Local(q) (Lemma 3)
2. Source \rightarrow Local(p) \rightarrow Local(q) (Lemma 4)
3. Source \rightarrow Multipole(p) \rightarrow Local(p) \rightarrow Local(q) (Lemma 5)

The main distinction among these we encounter is whether the interaction is mediated through an intermediate multipole or local expansion, or both. The list above shows, abstractly, the order of each expansion through the values p and q . In our envisioned usage scenario, q represents the order of the final QBX local expansion and will generically be lower than p . The reader familiar with conventional adaptive FMMs (e.g. [11]) may discover a direct correspondence of these types of translation chains and the various interaction lists used in those algorithms.

Without loss of generality, we may assume that an interaction goes through at most a single intermediate multipole expansion and intermediate local expansion, occupying a single level of the FMM’s hierarchy. This is due to the fact that, absent additional truncation, the FMM ‘forgets’ intermediate translations in the following way: the value of a local expansion shifted downward through a sequence of local-to-local (14) translations only depends on the source and the *initial* local expansion center. Similarly, the value of a multipole expansion shifted upward through a sequence of multipole-to-multipole (11) translations only depends on the source and the *final* multipole expansion center. (See [22, Lemma 2.3 and Lemma 2.5].)

We recall a technique from complex analysis for bounding the n -th derivative of a complex analytic function. The proof can be found in [14, IV.2.14 on page 73].

Proposition 1. *Let $U \subseteq \mathbb{C}$ be open and let $\phi : U \rightarrow \mathbb{C}$ be a complex analytic function. Let $z \in U, r > 0$ and suppose that $\overline{B}(z, r) \subseteq U$. Then for all $n \geq 0$*

$$|\phi^{(n)}(z)| \leq \frac{n!}{r^n} \left(\max_{w \in \overline{B}(z,r)} |\phi(w)| \right).$$

Remark 2. Although Lemmas 3, 4, and 5 are stated for a single source charge of unit strength, the statements can be straightforwardly generalized for an ensemble of m charges of strengths q_1, \dots, q_m , with the error bound scaled by $\sum_{k=1}^m |q_k|$.

See Fig. 7 for context on the following lemma.

Lemma 3 (Truncating a mediating multipole to p -th order on a q -th order local). *Let $\lambda, c, r > 0$. Suppose that a single unit strength charge is placed in the closed disk $\overline{B}(z_0, \lambda r)$ with radius λr and center z_0 , such that $|z_0| \geq (c + 1 + \lambda)r$. The corresponding multipole expansion with coefficients $(a_k)_{k=0}^\infty$ converges in the closed disk $\overline{B}(0, r)$ of radius r centered at the origin.*

Suppose that $y, z \in \overline{B}(0, r)$. Then if $|z| < r$ and $|y - z| \leq r - |z|$, the potential due to the charge is described by a power series

$$\phi(y) = \sum_{k=0}^\infty \beta_k (y - z)^k.$$

Fix the intermediate multipole order $p \geq 0$. For $n \geq 0$, let $\tilde{\beta}_n$ be the n -th coefficient of the local expansion centered at z obtained by translating the p -th order multipole expansion of ϕ :

$$\tilde{\beta}_n = \frac{1}{n!} \frac{d^n}{dz^n} \left(a_0 \log(z - z_0) + \sum_{k=1}^p \frac{a_k}{(z - z_0)^k} \right).$$

Define $\omega = 1/(1 + \frac{c}{\lambda})$. Fix the local expansion order $q \geq 0$. Then

$$\left| \sum_{k=0}^q \beta_k (y - z)^k - \sum_{k=0}^q \tilde{\beta}_k (y - z)^k \right| \leq \left(\frac{q + 1}{p + 1} \right) \left(\frac{\omega^{p+1}}{1 - \omega} \right).$$

Proof. We may write $\beta_n - \tilde{\beta}_n$ as

$$\beta_n - \tilde{\beta}_n = \frac{1}{n!} \frac{d^n}{dz^n} R_p(z)$$

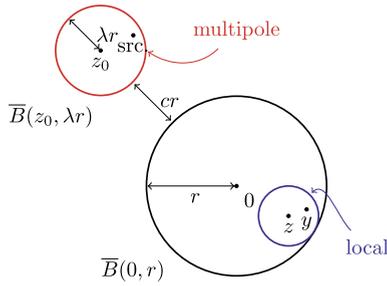


Fig. 7. Obtaining the local expansion of a point potential using an intermediate multipole expansion. The local expansion of the potential due to the source charge is formed by first forming a multipole expansion inside $\overline{B}(z_0, \lambda r)$ and then shifting to z . This provides the setting for Lemma 3.

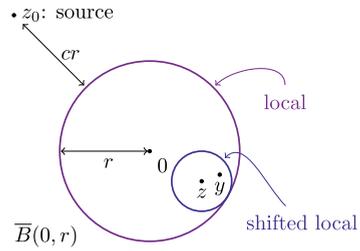


Fig. 8. Obtaining the local expansion of a point potential using an intermediate local expansion. The local expansion of the potential due to the source charge is formed inside the disk $\overline{B}(0, r)$ and then shifted to the center z . This provides the setting for Lemma 4.

where the function $R_p : (\mathbb{C} \setminus \overline{B}(z_0, \lambda r)) \rightarrow \mathbb{C}$, defined as

$$R_p(z) = \sum_{k=p+1}^{\infty} \frac{a_k}{(z - z_0)^k}$$

is what remains after truncating the multipole expansion of ϕ to p -th order.

We bound the n -th derivative of R_p at z . R_p is complex analytic and its domain contains the closed disk $\{w : |w - z| \leq r - |z|\}$, so by Proposition 1

$$|R_p^{(n)}(z)| \leq \frac{n!}{(r - |z|)^n} \left(\max_{y \in \overline{B}(z, r - |z|)} |R_p(y)| \right). \tag{15}$$

Recall (e.g., from (10)) that the multipole coefficients $(a_k)_{k=1}^{\infty}$ satisfy

$$|a_k| \leq \frac{(\lambda r)^k}{k}, \quad k > 0.$$

Using Fig. 7 and noting $z \neq z_0$, when $|y - z| \leq r - |z|$ and $k > 0$, we have

$$\left| \frac{a_k}{(y - z_0)^k} \right| \leq \frac{(\lambda r)^k}{k(cr + \lambda r)^k} = \frac{\omega^k}{k}.$$

Noting $\omega < 1$, we find

$$\max_{y \in \overline{B}(z, r - |z|)} |R_p(y)| \leq \sum_{k=p+1}^{\infty} \frac{\omega^k}{k} \leq \frac{1}{p+1} \left(\frac{\omega^{p+1}}{1 - \omega} \right). \tag{16}$$

Combining (15) and (16) yields

$$|\beta_n - \tilde{\beta}_n| = \left| \frac{R_p(z)}{n!} \right| \leq \frac{1}{p+1} \left(\frac{1}{(r - |z|)^n} \right) \left(\frac{\omega^{p+1}}{1 - \omega} \right). \tag{17}$$

From the triangle inequality and (17), we obtain the claim:

$$\left| \sum_{k=0}^q \beta_k(y - z)^k - \sum_{k=0}^q \tilde{\beta}_k(y - z)^k \right| \leq \sum_{k=0}^q \left(\frac{1}{p+1} \cdot \frac{\omega^{p+1}}{1 - \omega} \cdot \frac{1}{(r - |z|)^k} \right) (r - |z|)^k = \left(\frac{q+1}{p+1} \right) \left(\frac{\omega^{p+1}}{1 - \omega} \right). \quad \square$$

See Fig. 8 for context on the following lemma.

Lemma 4 (Truncating a mediating local to p -th order on a q -th order local). Let $c, r > 0$. Suppose that a single unit strength charge is placed at z_0 , with $|z_0| \geq (c + 1)r$. Consider the closed disk $\overline{B}(0, r)$ of radius r centered at the origin. Suppose that $y, z \in \overline{B}(0, r)$. If $|z| < r$ and $|y - z| \leq r - |z|$, the potential ϕ due to the charge is described by a power series

$$\phi(y) = \sum_{l=0}^{\infty} \beta_l(y - z)^l.$$

Fix the intermediate local order $p \geq 0$. For $n \geq 0$, let $\tilde{\beta}_n$ be the n -th coefficient of a local expansion centered at z obtained by translating a p -th order local expansion of ϕ centered at the origin:

$$\tilde{\beta}_n = \frac{1}{n!} \frac{d^n}{dz^n} \left(\sum_{k=0}^p \frac{\phi^{(k)}(0)}{k!} z^k \right).$$

Fix the local expansion order $q \geq 0$. Define $\alpha = 1/(1 + c)$. Then

$$\left| \sum_{k=0}^q \beta_k(y - z)^k - \sum_{k=0}^q \tilde{\beta}_k(y - z)^k \right| \leq \left(\frac{q + 1}{p + 1} \right) \left(\frac{\alpha^{p+1}}{1 - \alpha} \right).$$

Proof. This lemma may be proved with an argument almost identical to the proof of Lemma 3, so we only sketch the proof. We have $\beta_n - \tilde{\beta}_n = (1/n!)(d^n/dz^n)R_p(z)$ where the complex analytic function $R_p : \overline{B}(0, r) \rightarrow \mathbb{C}$ given by $R_p(z) = \sum_{k=p+1}^\infty -z^k/(kz_0^k)$ is the Taylor remainder of the Taylor series for ϕ (cf. (13)) centered at 0 and evaluated at z .

Noting $\alpha < 1$, applying Proposition 1 to R_p yields $|\beta_n - \tilde{\beta}_n| \leq 1/(p + 1) (1/(r - |z|)^n) (\alpha^{p+1}/(1 - \alpha))$. From this and the triangle inequality, the claim follows. \square

Once again, see Fig. 7 for context on the following lemma.

Lemma 5 (Truncating mediating multipole and local to p -th order on a q -th order local). Let $c, \lambda, r, (a_k)_{k=0}^\infty, \overline{B}(0, r), \overline{B}(z_0, \lambda r)$, and ϕ be as in Lemma 3. Let $y, z \in \overline{B}(0, r)$. Then if $|z| < r$ and $|y - z| \leq r - |z|$, the potential due to the charge is described by a power series

$$\phi(y) = \sum_{k=0}^\infty \beta_k(y - z)^k.$$

Fix the intermediate multipole and local order $p \geq 0$. For $n \geq 0$, let $\tilde{\zeta}_n$ be the n -th coefficient of the local expansion at the origin of the potential arising from the p -th order multipole expansion at z_0 of ϕ :

$$\tilde{\zeta}_n = \frac{1}{n!} \frac{d^n}{dz^n} \Big|_{z=0} \left(a_0 \log(z - z_0) + \sum_{k=1}^p \frac{a_k}{(z - z_0)^k} \right).$$

Also, let $\tilde{\beta}_n$ be the n -th coefficient of the local expansion at z of the potential arising from the p -th order local expansion with coefficients $(\tilde{\zeta}_k)_{k=0}^p$:

$$\tilde{\beta}_n = \frac{1}{n!} \frac{d^n}{dz^n} \left(\sum_{k=0}^p \tilde{\zeta}_k z^k \right).$$

Fix the final local expansion order $q \geq 0$. Define $\alpha = 1/(1 + c)$ and $\omega = 1/(1 + \frac{c}{\lambda})$. Then

$$\left| \sum_{k=0}^q \tilde{\beta}_k(y - z)^k - \sum_{k=0}^q \beta_k(y - z)^k \right| \leq (q + 1) \left(\frac{\omega^{p+1}}{1 - \omega} \right) + \left(\frac{q + 1}{p + 1} \right) \left(\frac{\alpha^{p+1}}{1 - \alpha} \right).$$

Proof. For $n \geq 0$, define τ_n as the n -th coefficient of the q -th order local expansion at z of the potential arising from the p -th order local expansion of the source potential ϕ at the origin:

$$\tau_n = \frac{1}{n!} \frac{d^n}{dz^n} \left(\sum_{k=0}^p \frac{\phi^{(k)}(0)}{k!} z^k \right).$$

From the triangle inequality,

$$\left| \sum_{k=0}^q \tilde{\beta}_k(y - z)^k - \sum_{k=0}^q \beta_k(y - z)^k \right| \leq \left| \sum_{k=0}^q \tilde{\beta}_k(y - z)^k - \sum_{k=0}^q \tau_k(y - z)^k \right| + \left| \sum_{k=0}^q \tau_k(y - z)^k - \sum_{k=0}^q \beta_k(y - z)^k \right|.$$

Realizing that the expansion with coefficients $(\tau_k)_{k=0}^q$ is the result of p -th order truncation of an intermediate local expansion, we can apply Lemma 4 to obtain

$$\left| \sum_{k=0}^q \tau_k(y - z)^k - \sum_{k=0}^q \beta_k(y - z)^k \right| \leq \left(\frac{q + 1}{p + 1} \right) \left(\frac{\alpha^{p+1}}{1 - \alpha} \right). \tag{18}$$

To estimate

$$\left| \sum_{k=0}^q \tilde{\beta}_k (y-z)^k - \sum_{k=0}^q \tau_k (y-z)^k \right|,$$

write

$$\tau_n - \tilde{\beta}_n = \frac{1}{n!} \frac{d^n}{dz^n} R_p(z)$$

for the complex analytic function

$$R_p(z) = \sum_{k=0}^p \left(\frac{\phi^{(k)}(0)}{k!} - \tilde{\zeta}_k \right) z^k.$$

Then we have from Proposition 1 that

$$|R_p^{(n)}(z)| \leq \frac{n!}{(r-|z|)^n} \left(\max_{y \in \overline{B}(z, r-|z|)} |R_p(y)| \right). \tag{19}$$

Realizing that R_p embodies the difference between a direct local expansion of the source and one mediated by the p -truncated multipole expansion with coefficients $(\tilde{\zeta}_k)_{k=0}^p$, both centered at the origin, we may apply Lemma 3 to find that for $y \in \overline{B}(z, r-|z|) \subseteq \overline{B}(0, r)$

$$|R_p(y)| \leq \frac{\omega^{p+1}}{1-\omega}. \tag{20}$$

Combining (19) and (20) we obtain

$$|\tau_n - \tilde{\beta}_n| = \left| \frac{R_p^{(n)}(z)}{n!} \right| \leq \frac{1}{(r-|z|)^n} \left(\frac{\omega^{p+1}}{1-\omega} \right).$$

This implies

$$\left| \sum_{k=0}^q \tilde{\beta}_k (y-z)^k - \sum_{k=0}^q \tau_k (y-z)^k \right| = \left| \sum_{k=0}^q (\tilde{\beta}_k - \tau_k) (y-z)^k \right| \leq (q+1) \left(\frac{\omega^{p+1}}{1-\omega} \right). \tag{21}$$

The claim follows from combining (18) and (21). \square

4. The GIGAQBX algorithm

The algorithm in this section is a hierarchical fast algorithm modeled on the adaptive Fast Multipole Method [11]. Section 2.5 provided a glimpse of the differences between the conventional point FMM and our modified version. For the benefit of readers familiar with point FMMs, these modifications in brief amount to:

- **Targets** (in the form of QBX centers) **may have a non-zero size or ‘extent’**. This extent is considered during tree construction. Interactions to the target from sources not well-separated from its radius are evaluated without expansion-based acceleration.
- To retain efficiency in the presence of this constraint, **each box has an associated ‘target confinement region’** (‘TCR’) which protrudes beyond the box by a fixed multiple of the box size. Targets (with their full extent) must fit within that region of a box to be eligible for inclusion in that box. If they do not fit, they will remain in a (larger) ancestor box.
- As a result of the prior point, **targets may occur in non-leaf boxes**. Interaction list generation must be modified to permit this.
- To avoid degradation of the expansion convergence factors in the presence of the target confinement region larger than the box, we **modify the basic recursion structure to consider a neighborhood two boxes wide** when measured from the target box instead of the classical point FMM’s one-wide region.
- To further retain convergence in the presence of the TCR, we **divide** some of the conventional **interaction lists (particularly, List 3 and 4) into ‘close’ and ‘far’ parts**, based on whether expansion-accelerated evaluation provides sufficient accuracy. The ‘close’ sub-lists are then evaluated directly, without acceleration.

The remainder of this section is devoted to a precise statement of the modified algorithm as well as an analysis of its complexity.

The input to the algorithm consists of: (a) a curve Γ discretized into panels equipped with a piecewise Gaussian quadrature rule; (b) a chosen accuracy $\varepsilon > 0$; (c) a density μ with values at the points of the discretized geometry; and (d) a set of (potentially on-surface) target points at which the potential is to be evaluated. The global accuracy parameter ε is used to determine the order and truncation parameters p_{QBX} , p_{FMM} , and p_{quad} , which we describe in Section 4.1. The geometry and targets should be preprocessed according to Section 4.2.

4.1. Choice of algorithm parameters

The splitting (5) of the error into FMM error, truncation error, and quadrature error, allows us to control the error components separately, so that in total they do not exceed the allowed precision ε .

The error estimates in Section 3 can be used to guarantee that the FMM error component is of order $\approx (\frac{1}{2})^{p_{\text{FMM}}+1}$ (cf. Theorem 7). Thus we may set $p_{\text{FMM}} \approx \lceil \log_2 \varepsilon \rceil$.

The QBX order p_{QBX} controls the truncation error component and can be set independently of the FMM order. Unlike the algorithm of [46], our algorithm does not require an artificial order increase to the FMM order to maintain accuracy depending on the value of p_{QBX} . Using (7) as the truncation error estimate, we see that the truncation error should be $O(r^{p_{\text{QBX}}+1})$. This means the choice of p_{QBX} depends on the expansion radius, and hence the length of the associated panel.

Finally, the quadrature error depends chiefly on the node count p_{quad} of the (upsampled) Gaussian quadrature. This error typically decays quickly in comparison to the other error components. For instance, assuming the QBX centers are placed a distance of $h/2$ from the panels, where h is the panel width, then the estimate (8) and the surrounding discussion imply the convergence factor for the quadrature error is approximately $(1/2)^{2p_{\text{quad}}}$. For calculations on curves in the plane, the rapid increase in accuracy with p_{quad} makes it expedient (if not necessarily efficient) to choose a generic, high value (e.g. $p_{\text{quad}} = 64$), ensuring the smallness of the quadrature error term. The contributions [32,33] provide precise means of estimating this error contribution.

4.2. Preparing geometry and targets

As with the algorithm in [46], a number of preprocessing steps are required on the inputs, which we include in this section by reference. The motivation and postconditions of these steps are described in Section 2.2. Detailed algorithms can be found in [46].

Specifically, we require that Γ has been refined according to the algorithm in Section 5 of [46]. The refinement procedure described there controls for quadrature and truncation error. Additionally, preprocessing needs to ensure that geometry and density are upsampled to the chosen quadrature node count p_{quad} . Lastly, our algorithm expects that targets needing QBX-based evaluation have been associated to expansion centers according to the algorithm in Section 6 of [46].

Unless otherwise noted, we make the same parameter choices as that contribution, concerning, e.g., center placement (cf. also Section 2.2.2) and oversampling.

4.3. Tree and interaction lists

The inputs to our algorithm give rise to a variety of entities in the plane, specifically source quadrature nodes, QBX centers, and target points not needing QBX-based evaluation. We will refer to these generically as ‘particles’. We disregard target points requiring QBX-based potential evaluation at this stage because their potential evaluation needs can be met simply by evaluating the local expansion that was obtained at the end of the algorithm at the target’s associated QBX center.

Our algorithm is based on a quadtree whose axis-aligned root box includes all these particles as well as all the entirety of each placed expansion disk. Each box (even a non-leaf box) may ‘own’ a subset of particles. The quadtree is formed by repeatedly subdividing boxes, starting with the root box. A box is subdivided if it owns more than n_{max} particles eligible to be owned by its child boxes. If a QBX disk does not fully fit within the target confinement region (see below) of the subdivided box, it is not eligible to be owned by the child box, and its center remains owned by the parent box.

4.3.1. Notation

In the context of the quadtree described above, we introduce the following notation: We will use $\overline{B_\infty}(r, c)$ to denote the closed ℓ^∞ ball (i.e., square) of radius r centered at c .

Let b be a box in the quadtree. We will use $|b|$ to denote the radius of b (i.e., half its width). The *target confinement region* (‘TCR’, also $\text{TCR}(b)$) of a box b with center c is $\overline{B_\infty}(|b|(1+t_f), c)$, where t_f is the *target confinement factor* (‘TCF’). We assume $0 \leq t_f < 1$. A typical value is $t_f = 0.9$ (cf. Theorem 7).

The *k-near neighborhood* of a box b with center c is the region $\overline{B_\infty}(|b|(1+2k), c)$. The *k-colleagues* of a box b are boxes of the same level as b that are contained inside the *k-near neighborhood* of b . T_b denotes the set of 2-colleagues of a box b . Two boxes at the same level are *k-well-separated* if they are not *k-colleagues*. The parent of b is denoted $\text{Parent}(b)$. The set of ancestors is $\text{Ancestors}(b)$. The set of descendants is $\text{Descendants}(b)$. Ancestors and Descendants are also defined in the natural way for sets of boxes. A box owning a point or QBX center target is called a *target box*. A box owning a source quadrature node is called a *source box*. Ancestors of target boxes are called *target-ancestor boxes*.

Definition 1 (*Adequate separation relation, \prec*). We define a relation \prec on the set of boxes and target confinement regions, with $a \prec b$ read as ‘ a is adequately separated from b , relative to the size of a ’.

We write $a \prec b$ for boxes a and b if the ℓ^∞ distance between a and b is at least $2|a|$, i.e. the ℓ^∞ distance between the centers of a and b is at least $3|a| + |b|$.

We write $a \prec \text{TCR}(b)$ for boxes a and b if the ℓ^∞ distance between a and $\text{TCR}(b)$ is at least $2|a|$, i.e. the ℓ^∞ distance between the centers of a and b is at least $3|a| + |b|(1 + t_f)$.

We write $\text{TCR}(a) \prec b$ for boxes a and b if the ℓ^∞ distance between $\text{TCR}(a)$ and b is at least $2|a|(1 + t_f)$, i.e. the ℓ^∞ distance between the centers of a and b is at least $3|a|(1 + t_f) + |b|$.

Because the size of the TCR is proportional to the box size, $\text{Parent}(a) \prec b$ implies $a \prec b$. We refer to this property as the ‘monotonicity’ of ‘ \prec ’.

4.3.2. Interaction lists

The core function of the FMM is to convey interactions between boxes by ways of multipole and local expansions. It is common for implementations to store lists of source boxes, one per expansion/interaction type and target or target-ancestor box. These lists are called *interaction lists*.

Roughly, the FMM proceeds by obtaining multipole expansions of the sources in each box, propagating them upwards in the tree (towards larger boxes), then using multipole-to-local translation to convert those to local expansions where allowable. These local expansions are then propagated down the tree and evaluated, yielding an approximation of the far field of the box. The near field is evaluated directly, completing the evaluation of the potential from all source boxes at each target box. In adaptive trees (like ours), it cannot be assumed that all subtrees have the same number of levels; additional interaction lists were introduced in [11] to deal with the arising special cases.

We motivate and define the interaction lists used in our implementation in this section. Building on these, a precise, step-by-step statement of our version of the FMM can be found in Section 4.4. For a target or target-ancestor box b , the interaction lists $U_b, V_b, W_b, W_b^{\text{close}}, W_b^{\text{far}}, X_b, X_b^{\text{close}}, X_b^{\text{far}}$ are sets of boxes defined as follows:

List 1 (U_b) enumerates interactions from boxes adjacent to b for which no acceleration scheme is used. This includes the interaction of the box with itself and, since target boxes can be non-leaf boxes, also interactions with b ’s descendants.

Definition 2 (*List 1, U_b*). For a target box b , U_b consists of all leaf boxes from among $\text{Descendants}(b) \cup \{b\}$ and the set of boxes adjacent to b .

List 2 (V_b) enumerates interactions from boxes of the same size/level as b with separation to b sufficient to satisfy the assumptions for required error bounds on multipole-to-local translation.

Definition 3 (*List 2, V_b*). For a target or target-ancestor box b , V_b consists of the children of the 2-colleagues of b ’s parent that are 2-well-separated from b .

List 3 (W_b) enumerates interactions between non-adjacent, not 2-well-separated sources/target box pairs in which the target box b is too large (considering its separation) to receive the contribution of the source box through multipole-to-local translation. These interactions are typically conveyed through evaluation of the source box’s multipole expansion and are implied to cover any children of the source box. Of the descendants of the 2-colleagues of b , the boxes of List 3 are the first ones to become non-adjacent to b as one descends the tree towards the leaves.

Definition 4 (*List 3, W_b*). For a target box b , a box $d \in \text{Descendants}(T_b)$ is in W_b if d is not adjacent to b and, for all $w \in \text{Ancestors}(d) \cap \text{Descendants}(T_b)$, w is adjacent to b .

The following observations are immediate: (a) List 3 of b contains the immediate children of any 2-colleagues of b not adjacent to b . (b) Any box in W_b is strictly smaller than b . (c) Any box $d \in W_b$ is separated from b by at least the width of d .

W_b specifies no geometric relationship of its constituent boxes to b ’s TCR. As a result, W_b never occurs explicitly in our algorithm. We merely use W_b as a stepping stone to define two sub-lists, W_b^{far} and W_b^{close} (‘List 3 far’ and ‘List 3 close’) whose definitions take into account the existence of the TCR. Considering Fig. 2, some elements of W_b may be too close to b for evaluation of the source multipole to deliver the required accuracy. Interactions between such boxes and b may be handled via direct evaluation. Since direct evaluation, unlike multipole evaluation, does not include information from children, child boxes of too-close boxes must also be considered. Boxes sufficiently far from b make up W_b^{far} (‘List 3 far’), while close leaf (source) boxes comprise W_b^{close} . Because of monotonicity, children of boxes in W_b^{far} also satisfy the TCR separation requirement. As an easy consequence, observe that while $W_b^{\text{close}} \cup W_b^{\text{far}} \subseteq W_b$ does not hold in general, $W_b^{\text{close}} \cup W_b^{\text{far}} \subseteq \text{Descendants}(W_b) \cup W_b$ is generally true.

Definition 5 (List 3 close, W_b^{close}). For a target box b , a leaf box d is said to be in W_b^{close} if $d \in \text{Descendants}(W_b) \cup W_b$ such that $d \not\prec \text{TCR}(b)$.

Definition 6 (List 3 far, W_b^{far}). For a target box b , a box d is in W_b^{far} if $d \in \text{Descendants}(W_b) \cup W_b$ such that $d \prec \text{TCR}(b)$ and, for all $w \in \text{Ancestors}(d) \cap (\text{Descendants}(W_b) \cup W_b)$, $w \not\prec \text{TCR}(b)$.

List 4 (X_b) enumerates interactions between non-adjacent, not 2-well-separated source/target box pairs in which the source/leaf box d is too large (considering its separation) to transmit its contribution to the target box b through multipole-to-local translation. These interactions are typically conveyed through formation of a local expansions from the source box's sources. Since this local expansion can then participate in the downward propagation, the interaction from d does not also need to be conveyed to b 's children by way of List 4. List 4 consists of non-adjacent 2-colleagues of b or 2-colleagues of its ancestors.

Definition 7 (List 4, X_b). For a target or target-ancestor box b , a source/leaf box d is in X_b if d is a 2-colleague of b and d is not adjacent to b . Additionally, a leaf box d is in List 4 of b if d is a 2-colleague of some ancestor of b and d is adjacent to $\text{Parent}(b)$ but not b itself.

The following observations are immediate: (a) Any box in X_b is at least the width of b . (b) Any box in X_b is separated from b by at least the width of b . (c) For any $d \in X_b$, either $b \in W_d$ or d is a 2-colleague of b .

Again, X_b specifies no geometric relationship of its constituent boxes to b 's TCR. As a result, X_b never occurs explicitly in our algorithm. We merely use X_b as a stepping stone to define two sub-lists, X_b^{far} and X_b^{close} ('List 4 far' and 'List 4 close') whose definitions take into account the existence of the TCR. Considering Fig. 5, some elements of X_b may be too close to b to allow the resulting local expansion to deliver the required accuracy. Interactions between such boxes and b may be handled via direct evaluation. If a source box d meets the separation requirement of $\text{Parent}(b)$, by monotonicity it will meet the separation requirements of b and its descendants. Hence it will enter the downward propagation at $\text{Parent}(b)$ and thus need not be part of either X_b^{far} or X_b^{close} . Conversely, if $d \in X_b$ while *not* meeting the separation requirement, it will need to be added to List 4 close of b and its descendants down to the level at which it meets the requirement, at which point its contribution enters the downward propagation via X_b^{far} .

Definition 8 (List 4 close, X_b^{close}). A box d is in X_b^{close} if for some $w \in \text{Ancestors}(b) \cup \{b\}$ we have $d \in X_w$ and furthermore $\text{TCR}(b) \not\prec d$.

Definition 9 (List 4 far, X_b^{far}). A box $d \in X_b$ is in List 4 far if $\text{TCR}(b) \prec d$. Furthermore, if b has a parent, a box $d \in X_{\text{Parent}(b)}^{\text{close}}$ is in List 4 far if $\text{TCR}(b) \prec d$.

Remark 6. In some cases, it is computationally cheaper to choose a direct interaction instead of an indirect one (see [16, Section 2.4] for an example of a treecode optimization based on this observation). The accuracy of the algorithm is not impacted negatively by this choice. For example, a multipole-to-QBX local interaction can be more expensive than the corresponding direct interaction if only a small number of sources contribute to the multipole expansion. Such a situation can occur for a box $b' \in W_b^{\text{far}}$. In this case we remove b' from W_b^{far} and place its leaf descendants in W_b^{close} . We make use of this possibility in Section 4.6.

4.4. Formal statement of the algorithm

We use the following notation: $P_b^{\text{near}}(t)$ denotes the potential at a target point t due to all sources in $U_b \cup W_b^{\text{close}} \cup X_b^{\text{close}}$, $P_b^W(t)$ denotes the potential at a target t due to all sources in W_b^{far} , $L_b^{q,\text{near}}(t)$ denotes the (QBX) local expansion of the potential at target/center t due to all sources in $U_b \cup W_b^{\text{close}} \cup X_b^{\text{close}}$, $L_b^{q,W}(t)$ denotes the (QBX) local expansion at target/center t due to all sources in W_b^{far} , and $L_b^{q,\text{far}}(t)$ denotes the local expansion at target/center t due to all sources not in $U_b \cup W_b \cup X_b^{\text{close}}$.

Algorithm: GIGAQBX Fast Multipole Method.

Require: The maximum number of FMM targets/sources n_{max} per box for quadtree refinement and a target confinement factor t_f are chosen.

Require: Based on the precision ε to be achieved, a QBX order p_{QBX} , an FMM order p_{FMM} , and an oversampled quadrature node count p_{quad} are chosen in accordance with Section 4.1.

Require: The input geometry and targets are preprocessed according to Section 4.2.

Ensure: An accurate approximation to the potential $S(\mu)$ at all target points is computed.

Stage 1: Build tree

Create a quadtree on the computational domain containing all sources, targets, and expansion centers.

repeat

Subdivide each box containing more than n_{\max} particles into four children, pruning any empty child boxes. If an expansion center cannot be placed in a child box with target confinement factor t_f due to its radius, it remains in the parent box.

until each box no longer needs to be subdivided or an iteration produced only empty child boxes

*Stage 2: Form multipoles***for all** boxes b **do**

Form a p_{FMM} -th order multipole expansion M_b centered at b due to sources owned by b .

end for**for all** boxes b in postorder **do**

For each child of b , shift the center of the multipole expansion at the child to b . Add the resulting expansions to M_b .

end for*Stage 3: Evaluate direct interactions***for all** boxes b **do**

For each conventional target t owned by b , add to $P_b^{\text{near}}(t)$ the contribution due to the interactions from sources owned by boxes in U_b to t .

end for**for all** boxes b **do**

For each expansion center target t owned by b , add to the expansion coefficients $L_b^{q,\text{near}}(t)$, the contribution due to the interactions from U_b to t .

end for*Stage 4: Translate multipoles to local expansions***for all** boxes b **do**

For each box $d \in V_b$, translate the multipole expansion M_d to a local expansion centered at b . Add the resulting expansions to obtain L_b^{far} .

end for*Stage 5(a): Evaluate direct interactions due to W_b^{close}*

Repeat Stage 3 with W_b^{close} instead of U_b .

Stage 5(b): Evaluate multipoles due to W_b^{far} **for all** boxes b **do**

For each conventional target t owned by b , evaluate the multipole expansion M_d of each box $d \in W_b^{\text{far}}$ to obtain $P_b^W(t)$.

end for**for all** boxes b **do**

For each expansion center target t owned by b , compute the expansion coefficients $L_b^{q,W}(t)$, due to the multipole expansion M_d of each box $d \in W_b^{\text{far}}$.

end for*Stage 6(a): Evaluate direct interactions due to X_b^{close}*

Repeat Stage 3 with X_b^{close} instead of U_b .

Stage 6(b): Form locals due to X_b^{far} **for all** boxes b **do**

Convert the field of every particle owned by boxes in X_b^{far} to a local expansion about b . Add to L_b^{far} .

end for*Stage 7: Propagate local expansions downward***for all** boxes b in preorder **do**

For each child d of b , shift the center of the local expansions L_b^{far} to the child. Add the resulting expansions to L_d^{far} respectively.

end for*Stage 8: Evaluate final potential at targets***for all** boxes b **do**

For each conventional target t owned by b , evaluate $L_b^{\text{far}}(t)$.

Add $P_b^{\text{near}}(t)$, $P_b^W(t)$, $L_b^{\text{far}}(t)$ to obtain the potential at t .

end for**for all** boxes b **do**

For each expansion center target t owned by b , translate L_b^{far} to t , obtaining $L_b^{q,\text{far}}(t)$.

Add $L_b^{q,\text{near}}(t)$, $L_b^{q,W}(t)$, $L_b^{q,\text{far}}(t)$ to obtain the QBX local expansion at t .

end for

4.5. Accuracy of the computed potential

Section 3 contains the information necessary to derive an accuracy estimate for the GIGAQBX FMM. The interaction lists are designed so that each interaction mediated through them has a provable convergence factor. As we shall see in the next theorem, only the List 2 convergence factor depends on t_f , and the convergence factor for Lists 3 and 4 far is fixed at $\sqrt{2}/3$. Thus, the overall accuracy of the GIGAQBX FMM is primarily determined by the choice of t_f .

Theorem 7. Fix a target confinement factor $0 \leq t_f < 6 - 2\sqrt{2}$ and define $\alpha = (t_f + \sqrt{2})/(6 - \sqrt{2}) < 1$. There exists a constant C such that for every target point $x \in \mathbb{R}^2$

$$|S_{\text{QBX}(p_{\text{QBX}}, N)}\mu(x) - \mathcal{G}_{p_{\text{FMM}}} [S_{\text{QBX}(p_{\text{QBX}}, N)}\mu(x)]| \leq \frac{1}{1 - \alpha} CA(p_{\text{QBX}} + 1) \max\left(\frac{\sqrt{2}}{3}, \alpha\right)^{p_{\text{FMM}}+1},$$

where $\mathcal{G}_{p_{\text{FMM}}}[\cdot]$ denotes approximation by the GIGAQBX FMM of order p_{FMM} and

$$A = \sum_{i=0}^N |w_i \mu(y_i)|,$$

with the $\{w_i\}$ the quadrature weights and the $\{y_i\} \subset \Gamma$ the quadrature nodes. In particular, for $t_f \leq 3 - 3/\sqrt{2} \approx 0.87$, we obtain

$$|S_{\text{QBX}(p_{\text{QBX}}, N)}\mu(x) - \mathcal{G}_{p_{\text{FMM}}} [S_{\text{QBX}(p_{\text{QBX}}, N)}\mu(x)]| \leq CA(p_{\text{QBX}} + 1) \left(\frac{1}{2}\right)^{p_{\text{FMM}}}.$$

C is independent of μ , t_f , p_{QBX} , p_{FMM} , p_{quad} , and of the curve Γ and its discretization.

Proof. The proof of the statement results from applying the error estimates of Section 3 to the geometric situations resulting from the definitions of the interaction lists in Section 4.3. We fix a target point x . Without loss of generality, we assume that x is associated with a QBX center. Let c be the QBX center associated with x . Every source point y_i contributing to the summation (6) contributes via either $L_b^{q,\text{near}}(c)$, $L_b^{q,W}(c)$, or $L_b^{q,\text{far}}(c)$, where b is the box that owns c . For $L_b^{q,\text{near}}(c)$, the contribution must arrive via a direct interaction. This contribution incurs no error.

For $L_b^{q,W}(c)$, the contribution must arrive via a W_b^{far} interaction. The contribution from all W_b^{far} interactions incurs an error of at most

$$\frac{3A}{3 - \sqrt{2}} \left(\frac{p_{\text{QBX}} + 1}{p_{\text{FMM}} + 1}\right) \left(\frac{\sqrt{2}}{3}\right)^{p_{\text{FMM}}+1}.$$

See Fig. 11 and Lemma 3. For $L_b^{q,\text{far}}(c)$, the contribution must arrive via a $V_{b'}$ or $X_{b'}^{\text{far}}$ interaction, where b' is either b or an ancestor of b . The contribution from all $X_{b'}^{\text{far}}$ interactions incurs an error of at most

$$\frac{3A}{3 - \sqrt{2}} \left(\frac{p_{\text{QBX}} + 1}{p_{\text{FMM}} + 1}\right) \left(\frac{\sqrt{2}}{3}\right)^{p_{\text{FMM}}+1}.$$

See Fig. 12 and Lemma 4. The contribution from all $V_{b'}$ interactions incurs an error of at most

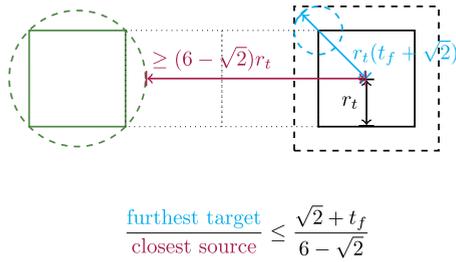
$$\frac{A}{1 - \alpha} (p_{\text{QBX}} + 1) \left(1 + \frac{1}{p_{\text{FMM}} + 1}\right) \alpha^{p_{\text{FMM}}+1}.$$

See Fig. 9 and Lemma 5. Figs. 9, 11, and 12 reinterpret the convergence factor geometrically in terms of ratios involving distances to sources and distances to targets. These are equivalent to the definitions of the convergence factors encountered in Lemmas 3, 4, and 5. Combining the estimates above yields the final error estimate. \square

The analysis in Section 3 that leads to the bound in Theorem 7 is not sharp due to various mathematical simplifications, as we will see in Section 5. Techniques similar to the ones of [42] may lead to sharper bounds.

4.6. Complexity

The purpose of this section is demonstrate that the GIGAQBX algorithm has a running time that is, roughly speaking, linear in the size of the input, assuming, roughly, that the number of sources within a neighborhood of each QBX disk is constant. Let S be a set of source points with $N_S = |S|$ (as obtained from the discretization of a curve Γ in accordance with Section 2), and let C be a set of expansion centers with $N_C = |C|$. Let $N = N_S + N_C$. Let the quadtree have N_B boxes and L levels.



$$\frac{\text{furthest target}}{\text{closest source}} \leq \frac{\sqrt{2} + t_f}{6 - \sqrt{2}}$$

Fig. 9. Separation criteria for List 2, with convergence factor calculation for a multipole-to-local-to-QBX local interaction. The target box with radius r_t is on the right. See Lemma 5.

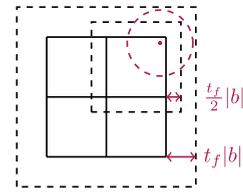
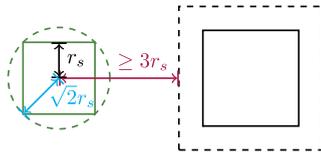
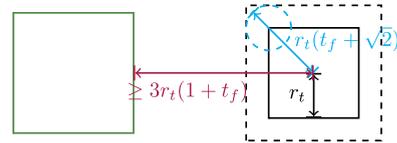


Fig. 10. The expansion disk of a suspended center must have radius at least $t_f/2$ times the radius of the box b that owns the center.



$$\frac{\text{furthest source}}{\text{closest target}} \leq \frac{\sqrt{2}}{3}$$

Fig. 11. Separation criteria for List 3 far, with convergence factor calculation for a multipole-to-QBX local interaction. The source box with radius r_s is on the left. See Lemma 3.



$$\frac{\text{furthest target}}{\text{closest source}} \leq \frac{\sqrt{2} + t_f}{3(1 + t_f)} \leq \frac{\sqrt{2}}{3}$$

Fig. 12. Separation criteria for List 4 far, with convergence factor calculation for a local-to-QBX local interaction. The target box with radius r_t is on the right. See Lemma 4.

Table 2
Complexity of each stage of the GIGAQBX Algorithm.

Stage	Modeled operation count	Note
Stage 1	NL	
Stage 2	$N_S p_{\text{FMM}} + N_B p_{\text{FMM}}^2$	There are N total particles with at most L levels of refinement.
Stage 3	$9N n_{\text{max}} p_{\text{QBX}} + N_C M_C p_{\text{QBX}}$	$N_S p_{\text{FMM}}$ for forming multipoles and the rest for shifting multipoles upward, with each shift costing p_{FMM}^2
Stage 4	$75N_B p_{\text{FMM}}^2$	Lemma 11
Stage 5	$N_C M_C p_{\text{QBX}} + 64N_C p_{\text{FMM}} p_{\text{QBX}} + 8N_S L n_{\text{max}} p_{\text{QBX}}$	Lemma 12
Stage 6	$63N_B n_{\text{max}} p_{\text{FMM}} + 42N_C n_{\text{max}} p_{\text{QBX}}$	Lemma 13
Stage 7	$4N_B p_{\text{FMM}}^2$	Lemma 15
Stage 8	$N_C p_{\text{FMM}} p_{\text{QBX}}$	The cost of shifting a local expansion downward is p_{FMM}^2 . There are at most 4 children per box. Cost of translating the box local expansions to N_C centers.

As a simplifying assumption, we restrict the complexity analysis to the case that the set of targets at which the potential is to be evaluated is covered by QBX expansion disks, or, in other words, we eliminate from consideration any targets whose potential can be evaluated through the conventional FMM algorithm without the use of QBX expansions.

We provide worst case running time bounds that apply to all particle distributions (although they do not imply linear complexity for all particle distributions). In particular, we do not attempt to obtain tight bounds on the leading constant terms from the complexity analysis. To complement the theoretical analysis with a more precise cost of each stage of the algorithm, we offer an empirical cost model in Section 5.2.2.

From the standpoint of complexity analysis, perhaps the most significant difference between a point FMM and the GIGAQBX FMM is that in the GIGAQBX algorithm it is no longer the case that there always are at most n_{max} particles per box. For the point FMM, this feature allows for bounding the number of near-neighborhood interactions between two boxes [11]. In the GIGAQBX FMM, it is possible for any number of QBX centers to cluster inside a box in the tree due to target confinement restrictions, and so another technique is needed to count the near-neighbor interactions.

Because we cannot rely on there being at most n_{max} centers per box, our analysis takes into account whether a QBX center is *suspended* in an upper level of the tree or not. A QBX center that is owned by a leaf box and *could* be owned by a hypothetical child of the leaf box is called *leaf-settled*. A QBX center that is not leaf-settled is called *suspended*. There are never more than n_{max} leaf-settled QBX centers in a box. For suspended centers, we make use of Proposition 9 below, which relates the size of the ‘near neighborhood’ of a suspended center to the size of the near neighborhood of its owner box.

A summary of the complexity results for each stage is given in Table 2. The rest of this section provides the details of this complexity analysis. We provide our complexity analysis in terms of ‘modeled floating point operations’. This means that while we include constants throughout, depending on interpretation, these constants may omit a flop-related constant

factor independent of problem parameters when we felt that no information was gained from including it for added realism. For instance, we model the cost of evaluating an expansion of order p as p , when more realistic operation counts might range from $2p + 1$ if multiplications and additions are counted, to $p + 1$ if a fused-multiply-add operation is assumed, to yet different counts if the computation of powers is taken into account.

4.6.1. Near neighborhoods of suspended QBX centers

We start with the following basic observation about suspended QBX centers.

Proposition 8. *Let c be a suspended QBX center of radius r_c owned by the box b_c . Then the closed square $\overline{B_\infty}(8r_c/t_f, c)$ is, geometrically, a superset of the 1-near neighborhood of b_c .*

Proof. Because c is suspended, c cannot be placed in any (hypothetical) child of b_c because it will not fit in the target confinement region. Since a child has radius $\frac{1}{2}|b_c|$, it follows that $r_c > \frac{t_f}{2}|b_c|$. This situation is illustrated in Fig. 10.

Regardless of where c is located in b_c , $\overline{B_\infty}(c, 4|b_c|)$ must contain the 1-near neighborhood of b_c . The claim follows since $\frac{8}{t_f}r_c > 4|b_c|$. \square

Based on this proposition, we define the following parameter. Consider the input to the GIGAQBX algorithm, with a set of centers C and a set of sources S . Then M_C is defined as

$$M_C = \frac{1}{N_C} \sum_{c \in C} \left| S \cap \overline{B_\infty} \left(c, \frac{8}{t_f} r_c \right) \right|$$

where r_c denotes the radius of the center c . In other words, M_C is the average number of sources that intersect with a square of radius $\frac{8}{t_f}r_c$ surrounding a QBX center c .

The running time of GIGAQBX algorithm depends non-trivially on the particle distribution in the tree. However, M_C is a geometry-dependent parameter that is independent of the tree structure, and can be used to provide a worst case bound on the 1-near neighborhood interaction cost for a suspended center. Moreover, if the geometry is smooth and refined in such a way that the panel sizes are locally uniform, M_C will not depend on the total number of particles. We give some values of M_C for actual geometries in Section 5.

The main utility of M_C is in the following proposition.

Proposition 9. *The number of source-center pairs $(s, c) \in S \times C$, such that c is a suspended center and s is in the 1-near neighborhood of the box that owns c , is at most $N_C M_C$.*

Proof. This follows immediately from Proposition 8 and the definition of M_C . \square

4.6.2. Detailed complexity analysis

In this section, we will use b_x to refer to the box that owns particle x and B to refer to the set of boxes in the quadtree.

Proposition 10. *Let b be an arbitrary box. Then there are at most 9 leaf boxes at least as large as b that intersect the 1-near neighborhood of b .*

Proof. Each such leaf box can be mapped injectively to b or one of the 8 colleagues of b it contains. \square

Lemma 11. *The amount of work done in Stage 3 (direct evaluation of the potential from List 1 source boxes) is at most*

$$9Nn_{\max}p_{\text{QBX}} + N_C M_C p_{\text{QBX}}.$$

Proof. We model the cost of all Stage 3 interactions as $p_{\text{QBX}}|U|$, where $U = \{(s, c) \in S \times C \mid b_s \in U_{b_c}\}$. U may be written as the disjoint union $U = U_{\text{big}} \cup U_{\text{small}}$ where U_{big} is the set of pairs (s, c) , such that $|b_s| \geq |b_c|$.

We bound $|U_{\text{big}}|$ as follows. Consider a center c . Then there are at most 9 leaf boxes larger than b_c that are in or adjacent to b_c , by Proposition 10. Therefore there are at most $9n_{\max}$ sources s such that $(s, c) \in U_{\text{big}}$. Thus $|U_{\text{big}}| \leq 9N_C n_{\max}$.

Now we bound $|U_{\text{small}}|$. We can group the particle-center interactions (s, c) in U_{small} according to whether c is suspended or leaf settled. We consider these two cases separately. Suppose that c is suspended. Then any s such that $(s, c) \in U_{\text{small}}$ must be in the 1-near neighborhood of b_c . By Proposition 9, the number of such (s, c) pairs is at most $M_C N_C$. Now, to consider settled centers, let s be a source. From Proposition 10, the number of leaf boxes larger than b_s and adjacent to it is at most 9. Therefore there are at most $9n_{\max}$ leaf settled centers c such that $(s, c) \in U_{\text{small}}$.

It follows that $|U_{\text{small}}| \leq M_C N_C + 9N_C n_{\max}$. Therefore the total cost of Stage 3 is at most $9Nn_{\max}p_{\text{QBX}} + M_C N_C p_{\text{QBX}}$. \square

Lemma 12. *The amount of work done in Stage 4 (translation of multipole to local expansions) is at most $75N_B p_{\text{FMM}}^2$.*

Proof. For each box b , $|V_b| \leq 10^2 - 5^2 = 75$ because there are at most 10^2 children of the parent of b and its 2-colleagues, and b and its 2-colleagues cannot be in V_b . There are N_B boxes, and each multipole to local translation has a modeled cost of p_{FMM}^2 operations. \square

Lemma 13. *Assume that $0 \leq t_f < 1$. The amount of work done in Stage 5 (handling Lists 3 close and far) is at most*

$$N_C M_C p_{\text{QBX}} + 64N_C p_{\text{FMM}} p_{\text{QBX}} + 8N_S L n_{\text{max}} p_{\text{QBX}}.$$

Proof. In order to simplify the analysis, we assume that, for every box b , W_b^{far} is disjoint from the 1-near neighborhood of b . In other words, for our cost analysis, we treat W_b -type interactions from within the 1-near neighborhood of b as being mediated through W_b^{close} . This assumption is conservative and will not lead to an underestimation of the cost because (1) direct interactions through W_b^{close} do not subsume interactions from their children and as such are more numerous and (2) any multipole evaluation that turns out to be more expensive than direct evaluation of the interaction from that box and its children may be replaced by the latter. (In fact, this latter strategy is a viable, if minor, cost optimization. See Remark 6 for details.)

We model the cost of Stage 5 as $p_{\text{FMM}} p_{\text{QBX}} |W_{\text{far}}| + p_{\text{QBX}} |W_{\text{close}}|$, where

$$W_{\text{far}} = \left\{ (b, c) \in B \times C \mid b \in W_{b_c}^{\text{far}} \right\}, \quad \text{and} \quad W_{\text{close}} = \left\{ (s, c) \in S \times C \mid b_s \in W_{b_c}^{\text{close}} \right\}.$$

We bound $|W_{\text{far}}|$ first. Let c be an expansion center. By our assumption on W_b^{far} and the 1-near-neighborhood, every box b with $(b, c) \in W_{\text{far}}$ is a descendant of one of the $5^2 - 3^2 = 16$ boxes that are 2-colleagues of b_c that are not adjacent to b_c . Since $0 \leq t_f < 1$, the direct descendants of b_c 's 2-colleagues satisfy Definition 6, and no children of the direct descendants can be in $W_{b_c}^{\text{far}}$. Thus there are at most 16×4 boxes b such that $(b, c) \in W_{\text{far}}$. In other words, $|W_{\text{far}}| \leq 64N_C$.

We handle $|W_{\text{close}}|$ next. Because $0 \leq t_f < 1$, for every center c , every box in $W_{b_c}^{\text{close}}$ is contained inside the 1-neighborhood of b_c . As in the proof of Lemma 11, we will consider separately the cases that the center is suspended or leaf settled. In the case that the center is suspended, Proposition 9 bounds the number of all such (s, c) pairs by $M_C N_C$. For the leaf settled case, let s be a source particle. Consider a pair $(s, c) \in W_{\text{close}}$. Then an ancestor of b_s is adjacent to b_c . Since there are L levels, there are at most $8L$ boxes adjacent to ancestors of b_s . So there are at most $8L n_{\text{max}}$ leaf settled centers c such that $(s, c) \in W_{\text{close}}$. It follows that $|W_{\text{close}}| \leq N_C M_C + 8N_S L n_{\text{max}}$.

Therefore the total cost of Stage 5 is at most $N_C M_C p_{\text{QBX}} + 64N_C p_{\text{FMM}} p_{\text{QBX}} + 8N_S L n_{\text{max}} p_{\text{QBX}}$. \square

Remark 14. The factor of $N_S L$ in the cost estimate of Lemma 13 implies that the cost of Stage 5 of the algorithm has a worst-case dependence on the number of particles times the number of levels in the tree. Since there are $\Omega(\log N)$ levels, this could lead to $\Omega(N \log N)$ algorithmic scaling. In practice, we have not observed this for the particle distributions we have tried and we expect $\Omega(N \log N)$ behavior to be uncommon in geometries used for layer potential evaluation. It is conceivable that a sharper analysis might be able to eliminate the factor.

In fact, if we assume the quadtree to be *level-restricted*, i.e. if we assume that adjacent leaf boxes' levels differ by at most one, then for every leaf box b , it can be shown that $|W_b^{\text{close}} \cup W_b^{\text{far}}|$ is at most a (dimension-dependent) constant. Because of this, every center in a leaf-settled box will interact via W_b^{close} with at most a constant number of source particles, and via W_b^{far} with at most a constant number of boxes. This allows us to remove the dependence on L in the complexity estimates.

Furthermore, using a level-restricted quadtree does not affect the asymptotic cost of any other stage of the algorithm. Starting with an arbitrary adaptive quadtree, the cost of converting it into a level-restricted quadtree is $O(N_B)$ and the resulting tree has $O(N_B)$ boxes. See [41, Theorem 1].

Lemma 15. *The amount of work done in Stage 6 (handling Lists 4 close and far) is at most*

$$63N_B n_{\text{max}} p_{\text{FMM}} + 42N_C n_{\text{max}} p_{\text{QBX}}.$$

Proof. To aid with the later analysis, we first bound $|X_b|$. Let $b \in B$. By definition of X_b , any box $b' \in X_b$ is not adjacent to b , and must be adjacent to the parent of b and at least as large as the parent of b , or a 2-colleague of b . Of the former, it is an easy result based on the same argument as Proposition 10 that there at most 5 such boxes. There are at most 16 2-colleagues of b not adjacent to b . Thus $|X_b| \leq 21$.

We begin by bounding the size of X_b^{close} . Recall from Definition 8 that X_b^{close} is a subset of the List 4's of b and its ancestors. In fact, as we now show, $X_b^{\text{close}} \subseteq X_b \cup X_{\text{Parent}(b)}$.

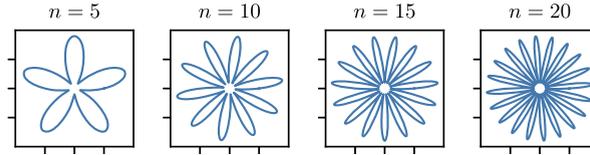


Fig. 13. A subset of the 'starfish' test geometries used to obtain many of the results of Section 5.

For any ancestor b' of b that is k levels above b , b is separated by at least a distance of $2^{k+1}|b|$ from any box in $X_{b'}$. Now suppose that b' is an ancestor of b that is $k \geq 2$ levels above b . Let $e \in X_{b'}$. Then $d(b, e) \geq 8|b|$. Furthermore, for any point $t \in \text{TCR}(b)$, $d(t, b) < |b|$ since $0 \leq t_f < 1$. From the reverse triangle inequality,

$$4|b| < 7|b| < |d(t, b) - d(b, e)| \leq d(t, e).$$

This means that $\text{TCR}(b) \prec e$, i.e. $e \notin X_b^{\text{close}}$. Since e was taken from the List 4 of a grandparent of b or higher, X_b^{close} must be a subset of $X_b \cup X_{\text{Parent}(b)}$. By the earlier argument, $|X_b^{\text{close}}| \leq 42$.

Recall from Definition 9 that X_b^{far} is a subset of $X_b \cup X_{\text{Parent}(b)}^{\text{close}}$. It follows that $|X_b^{\text{far}}| \leq 21 + 42 = 63$. Therefore the total cost of Stage 6 is at most is at most $63N_B n_{\max} p_{\text{FMM}} + 42N_C n_{\max} p_{\text{QBX}}$. \square

The next theorem summarizes the contents of this section. It is useful to consider the tree build phase (Stage 1) and the evaluation phase (the remaining stages) as conceptually distinct phases of the algorithm, with their own cost analysis. In the context of solving integral equations, the tree build typically only needs to be done once, while evaluation may need to be done many times, such as in the inner iteration of an iterative method like GMRES [49]. The tree build has a straightforward cost of $O(NL)$. The cost of the evaluation phase is more complicated to analyze.

As is the case for most adaptive point FMMs, linear running time is achievable under some set of additional assumptions on the particle distribution. For instance, in the paper [11], the algorithm is shown to be linear-time only for particle distributions with at most $|\log_2 \varepsilon|$ levels in the tree. Unlike [11], we choose not to make any assumptions on the number of levels of the tree. We instead establish that our algorithm (with a level-restricted quadtree) has a cost at most *linear in the number of boxes* N_B . Under the additional assumption of $N_B = O(N)$, cost linear in N_B implies that cost is also linear in the number of particles N . This assumption is weaker than limiting the number of levels of the tree. Recent work on adaptive FMMs [43] seeks to show that the adaptive FMM has time complexity linear in the number of particles irrespective of the particle distribution, by some modification to the basic FMM algorithm. Our variant of the FMM should be amenable to these modifications should the need arise.

Theorem 16. (a) The cost of the tree build phase of the GIGAQBXM is $O(NL)$.

(b) Assume that $p_{\text{FMM}} = O(|\log \varepsilon|)$, and that $p_{\text{QBX}} \leq p_{\text{FMM}}$. For a fixed value of n_{\max} , the modeled cost of the evaluation stage of the GIGAQBXM is $O((N + N_B)|\log \varepsilon|^2 + NL|\log \varepsilon| + N_C M_C |\log \varepsilon|)$. Using a level-restricted quadtree, the modeled cost is $O((N + N_B)|\log \varepsilon|^2 + N_C M_C |\log \varepsilon|)$. If the particle distribution satisfies $N_B = O(N)$ and $M_C = O(1)$, the worst-case modeled cost using a level-restricted quadtree is linear in N .

Proof. The proof of this theorem is evident from adding up the cost of the individual stages of the algorithm as given in Table 2. The cost of Stage 1 (the tree build phase) is immediate from the table. The cost of the remaining stages may be obtained by bounding p_{QBX} and p_{FMM} by $O(|\log \varepsilon|)$. The factor of L in Stage 5 may be eliminated by using a level-restricted quadtree, cf. Remark 14. \square

5. Results

In this section, we illustrate the numerical accuracy and cost scaling of the GIGAQBXM algorithm. We also perform a cost comparison of this algorithm with the prior algorithm of [46].

The experimental setup is as follows. We use a family of test geometries, parameterized by $n \in \mathbb{N}$, that form a 'starfish' curve $\gamma_n : [0, 1] \rightarrow \mathbb{R}^2$ whose parametrization is given by

$$\gamma_n(t) = (1 + 0.8 \sin(2\pi nt)) \begin{pmatrix} \cos(2\pi t) \\ \sin(2\pi t) \end{pmatrix}. \quad (22)$$

With increasing n , the starfish geometry has a larger number of more closely-spaced 'arms'. See Fig. 13 for graphical renditions of some of these geometries. We make use of these geometries because we have empirically found them to present a demanding scenario for layer potential evaluation, with varying feature and panel sizes, close spacing of unconnected parts of the geometry, and large (and scalable) overall size. We have found these characteristics to present an adequate challenge both the accuracy and the scalability of a layer potential evaluation code in a way that is representative of smooth source geometries 'in the wild'.

Table 3

ℓ^∞ error in Green's formula $\mathcal{S}(\partial_n u) - \mathcal{D}(u) = u/2$, scaled by $1/\|u\|_\infty$, for the 65-armed starfish γ_{65} , using the GIGAQBX algorithm. p_{FMM} denotes the FMM order and p_{QBx} the QBx order. The geometry was discretized with 3250 Gauss–Legendre panels, with 33 nodes per panel. Idealized point FMM error $(1/2)^{p_{\text{FMM}}+1}$ included for comparison. Entries in bold indicate that the FMM error is negligible.

$(1/2)^{p_{\text{FMM}}+1}$	p_{FMM}	$p_{\text{QBx}} = 3$	$p_{\text{QBx}} = 5$	$p_{\text{QBx}} = 7$	$p_{\text{QBx}} = 9$
0	(direct)	4.35×10^{-6}	6.21×10^{-7}	1.05×10^{-7}	5.71×10^{-8}
6×10^{-2}	3	5.16×10^{-3}	6.35×10^{-3}	6.33×10^{-3}	6.34×10^{-3}
2×10^{-2}	5	3.83×10^{-4}	5.95×10^{-4}	5.95×10^{-4}	5.93×10^{-4}
5×10^{-4}	10	4.35×10^{-6}	4.82×10^{-6}	6.94×10^{-6}	9.30×10^{-6}
2×10^{-5}	15	4.35×10^{-6}	6.21×10^{-7}	1.05×10^{-7}	1.76×10^{-7}
5×10^{-7}	20	4.35×10^{-6}	6.21×10^{-7}	1.05×10^{-7}	5.71×10^{-8}

Table 4

Analogous data to Table 3 for the conventional QBx FMM algorithm of [46].

$(1/2)^{p_{\text{FMM}}+1}$	p_{FMM}	$p_{\text{QBx}} = 3$	$p_{\text{QBx}} = 5$	$p_{\text{QBx}} = 7$	$p_{\text{QBx}} = 9$
0	(direct)	4.35×10^{-6}	6.21×10^{-7}	1.05×10^{-7}	5.71×10^{-8}
6×10^{-2}	3	2.55×10^{-2}	2.96×10^{-2}	4.07×10^{-2}	5.77×10^{-2}
2×10^{-2}	5	6.94×10^{-3}	1.61×10^{-2}	2.29×10^{-2}	3.10×10^{-2}
5×10^{-4}	10	4.95×10^{-4}	1.75×10^{-3}	5.80×10^{-3}	9.48×10^{-3}
2×10^{-5}	15	1.58×10^{-5}	1.85×10^{-4}	6.40×10^{-4}	3.17×10^{-3}
5×10^{-7}	20	4.35×10^{-6}	1.31×10^{-5}	8.99×10^{-5}	5.01×10^{-4}

5.1. Accuracy

We test the accuracy of the algorithm of this paper in a sequence of experiments. To assess accuracy, we employ *Green's formula* on the source geometry. Let Γ be the boundary of the domain. Let u be a harmonic function defined inside the domain and extending smoothly to the boundary. Because u extends smoothly to the boundary of the domain, the normal derivative $\partial_n u$ at the boundary is well-defined. Then Green's formula (e.g. [35, Theorem 6.5]) states that for $x \in \Gamma$,

$$\mathcal{S}(\partial_n u)(x) - \mathcal{D}(u)(x) = \frac{u(x)}{2}.$$

We use the residual in Green's formula as a convenient proxy for the accuracy attained in the evaluation of the layer potential evaluation as well as the overall accuracy attainable in application problems, in particular in the context of the solution of boundary value problems. We conducted two sets of experiments in this section. The first set is summarized in Table 5, which presents data to support the assertion that the residual in Green's formula is a reasonable proxy for the error in the solution of boundary value problems. In the second set of experiments, we measured the residual in the evaluation of Green's formula for a complicated geometry. We let u be the potential due to a charge located outside Γ at $(2, 1)$. We compute approximations to $\mathcal{S}(\partial_n u) - \mathcal{D}(u)$ and report the error in the discrete infinity norm. The error reported is the absolute error scaled by $1/\|u\|_\infty$ (so that it is relative to the magnitude of u). We use the starfish curve with $n = 65$ and test with various combinations of QBx order p_{QBx} and FMM order p_{FMM} . γ_{65} was discretized with 3250 Gauss–Legendre panels with 9 nodes oversampled to panels with 33 nodes (cf. Section 2.1). The curve was subsequently refined according to the refinement criteria of Section 2.2.1.

Table 3 shows the results of these experiments for the GIGAQBx FMM, varying p_{QBx} across columns and p_{FMM} across rows. The error incurred in unaccelerated QBx is shown in the first row of results. This value represents a lower bound on the accuracy of the scheme for errors of a given QBx order (as shown within a column); no error obtained with acceleration (as shown in the remaining rows) will be meaningfully smaller. Any error beyond the value in the first row is necessarily attributable to the effects of acceleration. We show table entries in bold if they do not significantly exceed the error value for unaccelerated QBx, indicating that the error contribution of FMM acceleration is negligible.

We choose the target confinement factor as $t_f = 0.9$. For that value of t_f , Theorem 7 roughly establishes $\|u\|_\infty (1/2)^{p_{\text{FMM}}+1}$ as a bound on the absolute error incurred by acceleration, neglecting a factor of $(p_{\text{QBx}} + 1)$ and a number of other factors that do not vary across the entries of the table. We show $(1/2)^{p_{\text{FMM}}+1}$ in the left column of the table. We find that the results support $\|u\|_\infty (1/2)^{p_{\text{FMM}}+1}$ as an asymptotic upper bound on the error, and, in turn, the assertion that the error in the potential computed via the GIGAQBx FMM is bounded simply by

$$|\text{unaccelerated QBx error}| + \|u\|_\infty (1/2)^{p_{\text{FMM}}+1}, \tag{23}$$

consistent with (5). In addition, we observe that the bound (23) lends itself to the simple interpretation that the additional error in the potential incurred from due to GIGAQBx FMM acceleration is asymptotically (in p_{FMM}) the same as the error incurred in the evaluation of a point potential in the adaptive FMM of [11].

Table 3 also allows us to assess the sharpness of the analysis underpinning Theorem 7. In the regime where the error is dominated by the contributions of FMM acceleration (the upper part of the table), we observe a match between the

Table 5

A comparison of relative ℓ^∞ errors attained at a set of target points in the solution of an exterior Neumann boundary value problem (3) using the integral equation (4) (top row of each segment) with errors attained in the residual of Green's formula on γ_{25} (bottom row of each segment). Discretization parameters for both problem types as well as the procedure for obtaining the residual in Green's formula are as in Table 3. Iteration counts for unpreconditioned GMRES are shown in the columns labeled '#it'. The discrete linear system used the weighting technique of [8]. To 'manufacture' a reference solution of the BVP, point potentials were evaluated originating from sources at locations $0.75[\cos\alpha_i, \sin\alpha_i]^T$ with 'charges' randomly assigned according to a standard normal distribution. The angles α_i are given by $\alpha_i = \pi/2 + 2\pi i/25$ ($i \in \{0, \dots, 24\}$). The ℓ^∞ norm of the vector of differences between the 'manufactured' potentials and the potentials from the BVP solve at the target points at locations $1.5[\cos(\pi + \alpha_i), \sin(\pi + \alpha_i)]^T$ was computed and is shown in the table.

$(1/2)^{p_{\text{FMM}}+1}$	p_{FMM}	$p_{\text{QBX}} = 3$	#it	$p_{\text{QBX}} = 5$	#it	$p_{\text{QBX}} = 7$	#it	$p_{\text{QBX}} = 9$	#it
6×10^{-2}	3	5.33×10^{-3} 2.78×10^{-3}	195	4.86×10^{-3} 2.86×10^{-3}	208	4.87×10^{-3} 2.84×10^{-3}	206	4.87×10^{-3} 2.84×10^{-3}	207
2×10^{-2}	5	2.57×10^{-3} 2.72×10^{-4}	189	2.32×10^{-4} 4.87×10^{-4}	188	3.13×10^{-4} 4.67×10^{-4}	180	3.02×10^{-4} 4.66×10^{-4}	182
5×10^{-4}	10	2.74×10^{-3} 8.26×10^{-5}	192	5.45×10^{-5} 1.88×10^{-5}	184	1.58×10^{-5} 5.83×10^{-6}	179	2.70×10^{-6} 5.10×10^{-6}	183
2×10^{-5}	15	2.74×10^{-3} 8.26×10^{-5}	192	5.46×10^{-5} 1.88×10^{-5}	184	1.53×10^{-5} 5.11×10^{-6}	179	9.34×10^{-7} 2.07×10^{-6}	183
5×10^{-7}	20	2.74×10^{-3} 8.26×10^{-5}	192	5.46×10^{-5} 1.88×10^{-5}	184	1.53×10^{-5} 5.11×10^{-6}	179	9.26×10^{-7} 2.07×10^{-6}	183

Table 6

Values of the parameter M_C for various starfish geometries γ_n parametrized by n . N_S denotes the number of source quadrature points. Shown here are percentiles for the distribution of the number of particles in a square of radius $8/t_f$ around each QBX center. Here, $t_f = 0.9$. M_C denotes the empirical mean of the distribution.

n	N_S	M_C	Percentiles				
			20%	40%	60%	80%	100%
5	9735	572.5	349.8	414.0	518.0	810.0	1696.0
15	52965	1001.2	567.0	696.0	913.0	1271.0	7833.0
25	122925	1030.2	620.0	696.0	842.0	1257.0	6775.0
35	255255	905.6	621.0	678.0	777.0	1141.0	3809.0
45	392040	924.0	624.0	684.0	777.0	1183.0	4229.0
55	555390	976.7	630.0	699.0	871.0	1247.0	4459.0
65	789360	964.3	625.0	681.0	825.0	1240.0	3909.0

bound and the behavior of the error in asymptotic behavior, although concrete error values are overestimated by around two orders of magnitude.

Table 4 shows an analogous set of results for the conventional QBX FMM of [46]. We find that the conventional QBX FMM is also able to match the error achieved by unaccelerated QBX, albeit at considerably higher p_{FMM} than our scheme. The relationship between p_{QBX} , p_{FMM} and the error is more complicated than the simple bound of (23). As a matter of fact, only empirical error data were shown in [46]. Most poignantly perhaps, for the conventional QBX FMM, the error contribution due to acceleration is *not* bounded by the FMM error incurred in a corresponding point FMM, and the error appears to degrade with increasing QBX order as p_{FMM} is held fixed. Our scheme exhibits neither of these two issues.

The difference in behavior between the two schemes is easily explained. The proofs of Lemma 3 and Lemma 4 give bounds on the error of individual expansion coefficients. Careful study shows that in a generic FMM translation operator, the higher order coefficients are approximated less accurately than the lower order coefficients, e.g. in formula (17). While this issue in principle applies to both versions of the scheme, the additional geometric restrictions in our version mitigate the impact of this phenomenon by controlling the amplification of this error by a geometric condition.

5.2. Cost and scalability

Having established that the accuracy of layer potentials evaluated GIGAQBX FMM can be understood with the help of easy-to-use estimates and that high levels of accuracy can be achieved, we seek to evaluate several aspects of the computational cost of our algorithm. First and foremost, we examine the scaling behavior of the scheme to large problem sizes. Next, we briefly highlight the cost-accuracy trade-off encountered. Lastly, since our scheme competes with the conventional QBX FMM, we give a cost comparison between the two approaches.

For the remainder of this section, we use the same family of 'starfish' geometries from (22) already familiar to the reader from our accuracy experiments. More specifically, for a fixed value of the 'arm count' n , we begin with the curve γ_n discretized into $50n$ panels equispaced in the parameter domain with 9 nodes per panel, which was upsampled to 33 nodes per panel. We use values of n ranging from 5 to 65 in increments of 10. Additional refinement in accordance with Section 2.2.1 was applied if necessary. Ultimately, this family of geometries ranged in size from about $1.7 \cdot 10^4$ to

Table 7

Cost per interaction list entry modeled in Figs. 14 and 15, i.e. for a single (source box, target box) interaction list pair. p_{FMM} = FMM order and p_{QBX} = QBX order. n_s = number of sources in the source box and n_t = number of QBX centers in the target box.

List	Cost
U_b	$p_{\text{QBX}}n_s n_t$
V_b	p_{FMM}^2
W_b^{close}	$p_{\text{QBX}}n_s n_t$
W_b^{far}	$p_{\text{FMM}}p_{\text{QBX}}n_t$
X_b^{close}	$p_{\text{QBX}}n_s n_t$
X_b^{far}	$p_{\text{FMM}}n_s$

about $1.4 \cdot 10^6$ particles, where by ‘particle’ we mean a class of entities including QBX centers, source quadrature nodes, and targets. We choose to employ this family of geometries with increasing complexity over, say, a simpler, growing grid of identical geometries because we expect the resulting scalability data to be credibly applicable to most other scenarios, including those of the growing grid.

5.2.1. Factors influencing computational scalability

Ideally, we would like to retain linear scaling of computational cost with the size of the geometries, as measured in the number of source quadrature points. Following the discussion of Section 4.6, it is not obvious that such scaling necessarily occurs. Recall the definition of the model parameter M_C , the main use of which is to provide a worst case bound on the number of direct interactions between source particles and *suspended* QBX centers. The cost of these interactions in the GIGAQBX FMM is always bounded from above by $O(N_C M_C)$, where N_C is the number of centers. For worst-case particle distributions, this cost is unavoidably quadratic, because M_C can be as large as N_C . Linear scaling of the method will only be seen if M_C does not change substantially across different-sized geometries. We expect particle distributions to which our method is applied to originate from discretizations of smooth, non-self-intersecting curves, and these are significantly more regular than an arbitrary particle distribution. Consequently it is conceivable that we will observe behavior considerably more benign than the worst case.

Although M_C does not depend on the tree, it is nevertheless not immediately obvious how one might derive a meaningful a-priori bound for M_C for general geometries that may ‘loop back’ on themselves in the way that (say) the starfish geometries do, bringing QBX centers into the proximity of source geometry non-adjacent to their ‘parent’ geometry. To empirically determine the behavior of M_C , we wrote a program that counts the number of source particles within $\overline{B}_\infty(8r_c/t_f, c)$ for each QBX center c . M_C is the mean of these counts. Recall Proposition 8, which states that this region is a superset of the 1-near neighborhood of b_c , which in turn represents the region with which a center may need to interact directly. The results are presented in Table 6, including means and percentiles for the distribution of source particle counts.

As can be seen in the table, the distribution of particles seems to be heavy-tailed, but with a mean value (M_C) of at most 1030 particles, which does not appear to be growing as the number of source particles increases. These data are consistent with the observation that M_C should not depend on the number of particles for smooth geometries of adequate refinement.

5.2.2. Experimental results on scaling and comparative cost

In this section, we illustrate the cost of the algorithm with the help of operation counts. To give a machine-independent understanding of the computational cost of our algorithm, we modeled computational cost by attributing an operation count to each entry in the interaction lists. The cost we attributed to an entry in each type of interaction lists is summarized in Table 7. Similarly to the approach of Section 4.6, the operation counts thus obtained are intended to roughly correspond with the number of floating point operations required.

We chose two FMM/QBX order pairs at which to gather this data for the GIGAQBX FMM, namely $(p_{\text{QBX}}, p_{\text{FMM}}) = (3, 10)$ and $(p_{\text{QBX}}, p_{\text{FMM}}) = (7, 15)$. These values yield an average of roughly five and six digits of accuracy, respectively. We show modeled operation counts across a number of ‘arm counts’ of the ‘starfish’ geometries, as described. The results are shown in graphical form in Fig. 14. In addition to the cost for each type interaction list, we also show an overall operation count summing the other contributions, labeled ‘all’.

The costs in Fig. 14 include the performance optimization mentioned in Remark 6. Our implementation used a W_b^{far} interaction only with source boxes having a cumulative source particle count of 15 or more. In every case, the improvement in cumulative operation counts due to this optimization was no more than 1%.

Before entering into a discussion of this data, we introduce a second set of data for comparison, based on the conventional QBX FMM of [46]. We applied the same cost model to the QBX FMM in order to perform an approximate comparison of the cost of the two algorithms. In order to make the comparison meaningful, we compare the computational cost of the FMMs for achieving a similar level of accuracy on the Green’s identity test (Section 5.1) with a fixed QBX order. Ex-

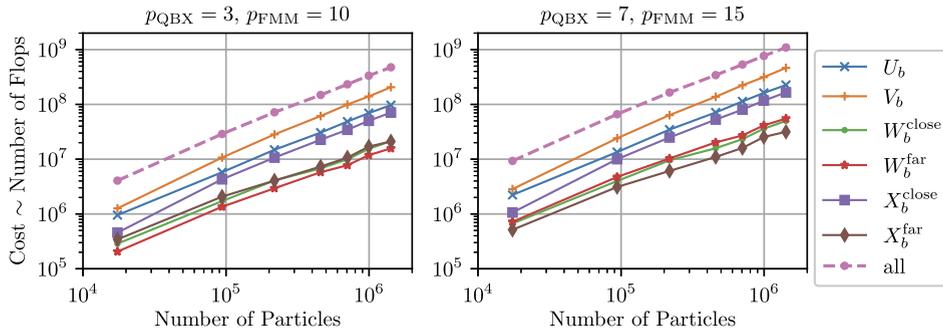


Fig. 14. Modeled operation counts for the GIGAQBXM FMM for evaluating the single layer potential on a sequence of ‘starfish’ geometries of increasing particle count. The operations are counted according to the model presented in Table 7. Here, $n_{\max} = 64$ and $t_f = 0.9$. The mean ℓ^∞ error in Green’s identity across all runs, scaled by $1/\|u\|_\infty$, was 4.97×10^{-5} for $p_{\text{QBXM}} = 3$ and 2.97×10^{-6} for $p_{\text{QBXM}} = 7$.

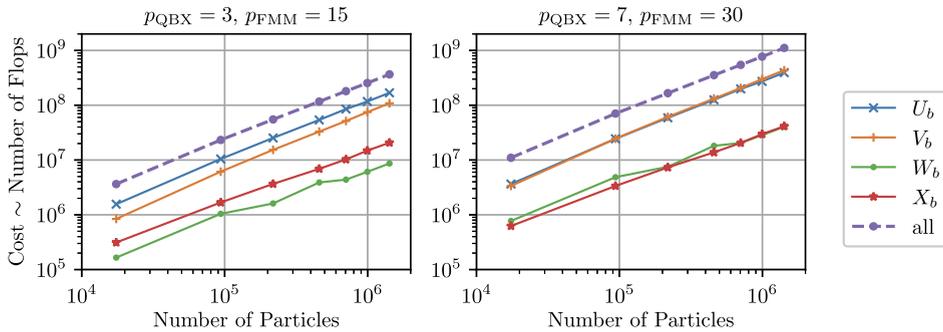


Fig. 15. Modeled operation counts for the conventional QBXM FMM of [46] for evaluating the single layer potential on a sequence of ‘starfish’ geometries of increasing particle count. The operations are counted according to the model presented in Table 7. Here, $n_{\max} = 128$. The mean ℓ^∞ error in Green’s identity across all runs, scaled by $1/\|u\|_\infty$, was 5.37×10^{-5} for $p_{\text{QBXM}} = 3$ and 3.08×10^{-6} for $p_{\text{QBXM}} = 7$.

periments showed that the (higher) FMM order values of $(p_{\text{QBXM}}, p_{\text{FMM}}) = (3, 15)$ and $(p_{\text{QBXM}}, p_{\text{FMM}}) = (7, 30)$ resulted in accuracies matching the above for the conventional QBXM FMM. We show graphs of computational cost across geometry sizes analogous to the earlier ones for this data set in Fig. 15.

We have tuned the user-chosen parameters for both algorithms to minimize their cost as measured by our model. (This process is also known as ‘balancing’ an FMM, since it tends to balance various contributions to the cost.) The main parameter amenable to such optimization is n_{\max} , the maximum number of particles per box. We observed that n_{\max} has different impact on the performance for the two algorithms. Roughly, the GIGAQBXM FMM will benefit from a smaller n_{\max} , as this can potentially decrease the number of direct interactions. In contrast, the conventional QBXM FMM benefits from a larger n_{\max} . The main reason this is the case is that this reduces the number of boxes/levels in the tree, and hence the number of multipole-to-local translations. We also observed a noticeable degradation of accuracy for small n_{\max} in the conventional QBXM FMM, which we believe may be related to the effect of n_{\max} on source/target separation. As a result, we found $n_{\max} = 64$ for the GIGAQBXM FMM and $n_{\max} = 128$ for the conventional QBXM FMM to yield near-minimal modeled cost. We used $t_f = 0.9$ for the GIGAQBXM FMM.

The linear scaling of both schemes is evident from the slope of the graphs, with one decade of geometry growth (indicated by the vertical grid lines) leading to one decade of cost growth (indicated by horizontal grid lines). As is typical for schemes based on the FMM, the overall cost is dominated by multipole-to-local translations (List 2/ V_b) and direct interactions (List 1/ U_b). Additionally, in the GIGAQBXM FMM, List 4 close (X_b^{close} ; which consists of direct interactions just like List 1) is also a significant contributor to the cost.

The overall operation counts for the two schemes are roughly comparable, with $p_{\text{QBXM}} = 7$ more closely matching than $p_{\text{QBXM}} = 3$. For $p_{\text{QBXM}} = 3$, the GIGAQBXM FMM has on average $1.3\times$ as many modeled operations as the QBXM FMM, but for $p_{\text{QBXM}} = 7$, it has about $0.96\times$ as many. In terms of actual wall times for evaluating the single layer potential, our implementation of the GIGAQBXM FMM is on average 21% slower than conventional QBXM FMM for $p_{\text{QBXM}} = 3$ and 17% slower for $p_{\text{QBXM}} = 7$. As the QBXM order increases, we expect the GIGAQBXM FMM to maintain its competitiveness, particularly considering the rapid growth of the FMM orders required to maintain accuracy in the conventional QBXM FMM.

Another factor worth highlighting is that the GIGAQBXM FMM is composed of a larger number of simpler operations, typically, lower-order expansions (thus with relatively short chains of dependent computations within one translation operation), while the conventional QBXM FMM uses fewer higher-complexity operations to translate expansions of higher order. While, according to our results above, these costs are similar for sequential execution, we expect that the GIGAQBXM FMM will be able to make better use of massively parallel computational resources.

A number of limitations of this study are evident. Both schemes stand to benefit from standard FMM optimizations that have not been applied, such as, for instance, using translation operators with asymptotically improved costs [28]. Additionally, tuning for the specific hardware was not applied when measuring the wall time. Nevertheless, the results in this section suggest that the two schemes are competitive in terms of cost.

6. Conclusions and future work

In this paper, we have presented a fast algorithm for Quadrature by Expansion for which we have also supplied analytical accuracy estimates. This algorithm is compatible with (and builds upon) previous work designed to control for truncation and quadrature error in QBX [46]. Unlike this and other work on accelerated global QBX, we are able to prove *strong accuracy guarantees* for our fast high order QBX scheme while retaining a cost comparable with or cheaper than previous schemes. We have demonstrated the viability of our approach through numerical experiments. Lastly, we have provided a set of sufficient conditions under which the algorithm exhibits linear scaling and also shown that in practice the algorithm scales linearly on complicated geometries.

Traditional hierarchical algorithms developed for n -body problems have considered *point* (i.e., zero-dimensional) sources and targets. An important feature of our work is the recognition that local expansions behave like ‘targets with extent’ from the point of view of the accuracy of translation operators. It is possible to view this work purely in this context, removed from QBX: As a fast algorithm that permits targets with extent.

Many exciting avenues for future work open up building upon this contribution: First, the cost estimates of Section 4.6 are inherently pessimistic because they do not leverage very much information about the particle distribution. Furthermore, they are also conservative when it comes to constants. Sharpening these estimates, perhaps with a more detailed understanding of the typical particle distribution of a source curve, will help provide a better understanding of the cost of our algorithm along with ideas to reduce said cost. Another direction of work is to apply the techniques of the error analysis used in this paper in order to understand analytically the accuracy behavior for the global QBX FMM of [46]. Preliminary results along these lines are encouraging. Lastly, we are in the process of extending the FMM developed in this paper to more kernels in two and three dimensions, with the goal of providing black-box, fast, and accurate layer potential evaluation for any kernel for which FMM translation infrastructure is available.

Acknowledgements

The authors’ research was supported by the National Science Foundation under grants DMS-1418961 and DMS-1654756 and through start-up funding provided by the University of Illinois. Part of the work was performed while the authors were participating in the HKUST-ICERM workshop ‘Integral Equation Methods, Fast Algorithms and Their Applications to Fluid Dynamics and Materials Science’ held in 2017.

References

- [1] K.E. Atkinson, D. Chien, Piecewise polynomial collocation for boundary integral equations, *SIAM J. Sci. Comput.* 16 (3) (1995) 651–681.
- [2] Kendall Atkinson, A Survey of Numerical Methods for the Solution of Fredholm Integral Equations of the Second Kind, *Soc. for Industrial and Applied Mathematics*, 1976.
- [3] Josh Barnes, Piet Hut, A hierarchical $O(N \log N)$ force-calculation algorithm, *Nature* 324 (6096) (1986) 446–449.
- [4] Alexander H. Barnett, Evaluation of layer potentials close to the boundary for Laplace and Helmholtz problems on analytic planar domains, *SIAM J. Sci. Comput.* 36 (2) (2014) A427–A451, <https://doi.org/10.1137/120900253>.
- [5] J.T. Beale, M.-C. Lai, A method for computing nearly singular integrals, *SIAM J. Sci. Comput.* 38 (6) (2001) 1902–1925, <https://doi.org/10.1137/S0036142999362845>.
- [6] George Biros, Lexing Ying, Denis Zorin, An embedded boundary integral solver for the unsteady incompressible Navier–Stokes equations, *J. Comput. Phys.* (2004) 121–141.
- [7] J. Bremer, Z. Gimbutas, V. Rokhlin, A nonlinear optimization procedure for generalized Gaussian quadratures, *SIAM J. Sci. Comput.* 32 (2010) 1761–1788, <https://doi.org/10.1137/080737046>.
- [8] James Bremer, On the Nyström discretization of integral equations on planar curves with corners, *Appl. Comput. Harmon. Anal.* 32 (1) (2012) 45–64, <https://doi.org/10.1016/j.acha.2011.03.002>.
- [9] O.P. Bruno, 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) 80–110, <https://doi.org/10.1006/jcph.2001.6714>.
- [10] M. Carley, Numerical quadratures for singular and hypersingular integrals in boundary element methods, *SIAM J. Sci. Comput.* 29 (3) (2007) 1207–1216.
- [11] J. Carrier, L. Greengard, V. Rokhlin, A fast adaptive multipole algorithm for particle simulations, *SIAM J. Sci. Stat. Comput.* 9 (4) (1988) 669–686, <https://doi.org/10.1137/0909044>.
- [12] Roman Chapko, Rainer Kress, Lars Mönch, On the numerical solution of a hypersingular integral equation for elastic scattering from a planar crack, *IMA J. Numer. Anal.* 20 (4) (Oct. 2000) 601–619, <https://doi.org/10.1093/imanum/20.4.601>.
- [13] Andrew J. Christlieb, et al., Grid-free plasma simulation techniques, *IEEE Trans. Plasma Sci.* 34 (2) (2006) 149–165, <https://doi.org/10.1109/TPS.2006.871104>.
- [14] John B. Conway, *Functions of One Complex Variable*, second ed., *Graduate Texts in Mathematics*, vol. 11, Springer-Verlag, New York–Berlin, 1978, pp. xiii+317.
- [15] P.J. Davis, P. Rabinowitz, *Methods of Numerical Integration*, Academic Press, San Diego, 1984.
- [16] Walter Dehnen, A hierarchical $O(N)$ force calculation algorithm, *J. Comput. Phys.* 179 (1) (2002) 27–42, <https://doi.org/10.1006/jcph.2002.7026>.
- [17] Charles L. Epstein, Leslie Greengard, Andreas Klöckner, On the convergence of local expansions of layer potentials, *SIAM J. Numer. Anal.* 51 (5) (2013) 2660–2679, <https://doi.org/10.1137/120902859>.

- [18] L. Farina, Evaluation of single layer potentials over curved surfaces, *SIAM J. Sci. Comput.* 23 (1) (2001) 81–91, <https://doi.org/10.1137/S1064827599363393>.
- [19] J. Goodman, T.Y. Hou, J. Lowengrub, Convergence of the point vortex method for the 2-D Euler equations, *Commun. Pure Appl. Math.* 43 (1990) 415–430, <https://doi.org/10.1002/cpa.3160430305>.
- [20] R.D. Graglia, G. Lombardi, Machine precision evaluation of singular and nearly singular potential integrals by use of Gauss quadrature formulas for rational functions, *IEEE Trans. Antennas Propag.* 56 (4) (2008) 981–998, <https://doi.org/10.1109/TAP.2008.919181>.
- [21] Leslie Greengard, *The Rapid Evaluation of Potential Fields in Particle Systems*, ACM Distinguished Dissertations, MIT Press, Cambridge, MA, 1988, pp. xiv+91.
- [22] Leslie Greengard, Vladimir Rokhlin, A fast algorithm for particle simulations, *J. Comput. Phys.* 73 (2) (1987) 325–348, [https://doi.org/10.1016/0021-9991\(87\)90140-9](https://doi.org/10.1016/0021-9991(87)90140-9).
- [23] W. Hackbusch, S.A. Sauter, On numerical cubatures of nearly singular surface integrals arising in BEM collocation, *Computing* 52 (2) (1994) 139–159, <https://doi.org/10.1007/BF02238073>.
- [24] S. Hao, et al., High-order accurate methods for Nyström discretization of integral equations on smooth curves in the plane, *Adv. Comput. Math.* 40 (1) (Feb. 2014) 245–272, <https://doi.org/10.1007/s10444-013-9306-3>.
- [25] D.J. Haroldsen, D.I. Meiron, Numerical calculation of three-dimensional interfacial potential flows using the point vortex method, *Commun. Pure Appl. Math.* 43 (1990) 415–430, <https://doi.org/10.1137/S1064827596302060>.
- [26] J. Helsing, R. Ojala, Corner singularities for elliptic problems: integral equations, graded meshes, quadrature, and compressed inverse preconditioning, *J. Comput. Phys.* 227 (2008) 8820–8840, <https://doi.org/10.1016/j.jcp.2008.06.022>.
- [27] J. Helsing, R. Ojala, On the evaluation of layer potentials close to their sources, *J. Comput. Phys.* 227 (2008) 2899–2921, <https://doi.org/10.1016/j.jcp.2007.11.024>.
- [28] Tomasz Hrycak, Vladimir Rokhlin, An improved fast multipole algorithm for potential fields, *SIAM J. Sci. Comput.* 19 (6) (1998) 1804–1826, <https://doi.org/10.1137/S106482759630989X>.
- [29] S. Jarvenpää, M. Taskinen, P. Ylä-Oijala, Singularity extraction technique for integral equation methods with higher order basis functions on plane triangles and tetrahedra, *Int. J. Numer. Methods Eng.* 58 (2003) 1149–1165, <https://doi.org/10.1002/nme.810>.
- [30] C.G.L. Johnson, L.R. Scott, An analysis of quadrature errors in second-kind boundary integral methods, *SIAM J. Numer. Anal.* 26 (6) (1989) 1356–1382, <https://doi.org/10.1137/0726079>.
- [31] M.A. Khayat, D.R. Wilton, Numerical evaluation of singular and near-singular potential integrals, *IEEE Trans. Antennas Propag.* 53 (10) (2005) 3180–3190, <https://doi.org/10.1109/TAP.2005.856342>.
- [32] Ludvig af Klinteberg, Anna-Karin Tornberg, Adaptive quadrature by expansion for layer potential evaluation in two dimensions, *SIAM J. Sci. Comp.* 40 (3) (2018) A1225–A1249, <https://doi.org/10.1137/17M1121615>.
- [33] Ludvig af Klinteberg, Anna-Karin Tornberg, Error estimation for quadrature by expansion in layer potential evaluation, *Adv. Comput. Math.* 43 (1) (2017) 195–234, <https://doi.org/10.1007/s10444-016-9484-x>.
- [34] Andreas Klöckner, et al., Quadrature by expansion: a new method for the evaluation of layer potentials, *J. Comput. Phys.* 252 (2013) 332–349, <https://doi.org/10.1016/j.jcp.2013.06.027>.
- [35] Rainer Kress, *Linear Integral Equations*, third ed., Applied Mathematical Sciences, vol. 82, Springer, New York, 2014, pp. xvi+412.
- [36] Rainer Kußmaul, Ein numerisches Verfahren zur Lösung des Neumannschen Außenraumproblems für die Helmholtzsche Schwingungsgleichung, *Computing* 4 (3) (1969) 246–273.
- [37] J. Lowengrub, M. Shelley, B. Merriman, High-order and efficient methods for the vorticity formulation of the Euler equations, *SIAM J. Sci. Comput.* 14 (1993) 1107–1142, <https://doi.org/10.1137/0914067>.
- [38] J.N. Lyness, L.M. Delves, On numerical contour integration round a closed contour, *Math. Comput.* 21 (100) (Oct. 1967) 561–577, <https://doi.org/10.2307/2005000>.
- [39] Erich Martensen, Über eine Methode zum räumlichen Neumannschen Problem mit einer Anwendung für torusartige Berandungen, *Acta Math.* 109 (1) (1963) 75–135.
- [40] A. Mayo, Fast, high-order accurate solution of Laplace's equation on irregular regions, *SIAM J. Sci. Comput.* 20 (1998) 648–683, <https://doi.org/10.1137/0906012>.
- [41] Doug Moore, The cost of balancing generalized quadtrees, in: *Proceedings of the Third ACM Symposium on Solid Modeling and Applications, SMA '95*, Salt Lake City, Utah, USA, ACM, 1995, pp. 305–312.
- [42] Henrik G. Petersen, et al., Error estimates for the fast multipole method. I. The two-dimensional case, *Proc. Math. Phys. Sci.* 448 (1934) (1995) 389–400, <https://doi.org/10.1098/rspa.1995.0023>.
- [43] Hadi Pouransari, Eric Darve, Optimizing the adaptive fast multipole method for fractal sets, *SIAM J. Sci. Comput.* 37 (2) (2015) A1040–A1066, <https://doi.org/10.1137/140962681>.
- [44] Manas Rachh, *Integral Equation Methods for Problems in Electrostatics, Elastostatics and Viscous Flow*, PhD thesis, New York University, 2015.
- [45] Manas Rachh, Andreas Klöckner, Leslie Greengard, Fast Algorithms for 'Quadrature by Expansion' II: Unidirectional Interactions, in preparation.
- [46] Manas Rachh, Andreas Klöckner, Michael O'Neil, Fast algorithms for quadrature by expansion I: globally valid expansions, *J. Comput. Phys.* 345 (2017) 706–731, <https://doi.org/10.1016/j.jcp.2017.04.062>.
- [47] Abtin Rahimian, Alex Barnett, Denis Zorin, Ubiquitous evaluation of layer potentials using quadrature by kernel-independent expansion, *BIT Numer. Math.* (ISSN 1572-9125) (2017).
- [48] Vladimir Rokhlin, Rapid solution of integral equations of classical potential theory, *J. Comput. Phys.* 60 (2) (1985) 187–207, [https://doi.org/10.1016/0021-9991\(85\)90002-6](https://doi.org/10.1016/0021-9991(85)90002-6).
- [49] Youcef Saad, Martin H. Schultz, GMRES: a generalized minimal residual algorithm for solving nonsymmetric linear systems, *SIAM J. Sci. Stat. Comput.* 7 (3) (July 1986) 856–869, <https://doi.org/10.1137/0907058>.
- [50] C. Schwab, W.L. Wendland, On numerical cubatures of singular surface integrals in boundary element methods, *Numer. Math.* 62 (1992) 342–369, <https://doi.org/10.1007/BF01396234>.
- [51] A. Sidi, M. Israeli, Quadrature methods for periodic singular Fredholm integral equations, *J. Sci. Comput.* 3 (1988) 201–231, <https://doi.org/10.1007/BF01061258>.
- [52] Michael Siegel, Anna-Karin Tornberg, A local target specific quadrature by expansion method for evaluation of layer potentials in 3D, *J. Comput. Phys.* (ISSN 0021-9991) 364 (2018) 365–392, <https://doi.org/10.1016/j.jcp.2018.03.006>.
- [53] J. Strain, Locally-corrected multidimensional quadrature rules for singular functions, *SIAM J. Sci. Comput.* 16 (4) (1995) 992–1017, <https://doi.org/10.1137/0916058>.
- [54] C.A. White, et al., The continuous fast multipole method, *Chem. Phys. Lett.* 230 (Nov. 1994) 8–16, [https://doi.org/10.1016/0009-2614\(94\)01128-1](https://doi.org/10.1016/0009-2614(94)01128-1).
- [55] Nathan Yarvin, Vladimir Rokhlin, Generalized Gaussian quadratures and singular value decompositions of integral operators, *SIAM J. Sci. Comput.* 20 (2) (1998) 699–718, <https://doi.org/10.1137/S1064827596310779>.
- [56] L. Ying, G. Biros, D. Zorin, A high-order 3D boundary integral equation solver for elliptic PDEs in smooth domains, *J. Comput. Phys.* 219 (2006) 247–275, <https://doi.org/10.1016/j.jcp.2006.03.021>.
- [57] Lexing Ying, George Biros, Denis Zorin, A kernel-independent adaptive fast multipole algorithm in two and three dimensions, *J. Comput. Phys.* 196 (2) (2004) 591–626, <https://doi.org/10.1016/j.jcp.2003.11.021>.