Fourier Based Fast Multipole Method for the Helmholtz Equation

Cris Cecka [a] and Eric Darve [a,b],*

[a]*Institute for Computational and Mathematical Engineering, Stanford University*
[b]*Mechanical Engineering Department, Stanford University*

Abstract

The fast multipole method (FMM) has had great success in reducing the computational time required to solve the boundary integral form of the Helmholtz equation. We present a formulation of the Helmholtz FMM using Fourier basis functions rather than spherical harmonics that accelerates some of the time-critical stages of the algorithm. With modifications to the transfer function in the precomputation stage of the FMM, the interpolation and anterpolation operators become straightforward applications of fast Fourier transforms and the transfer operator remains diagonal. Using Fourier analysis, constructive algorithms are derived to a priori determine an integration quadrature for a given error tolerance. Sharp error bounds are derived and verified numerically. Various optimizations are considered to reduce the number of quadrature points and reduce the cost of computing the transfer function.

Key words: fast multipole method, fast Fourier transform, Fourier basis, interpolation, anterpolation, Helmholtz, Maxwell, integral equations, boundary element method

1 Introduction

Since the development of the fast multipole method (FMM) for the wave equation in [1–5], the FMM has proven to be a very effective tool for solving scalar acoustic and vector electromagnetic problems. In this paper, we consider the application of the FMM to the scalar Helmholtz equations, although our results can be immediately extended to the vector case as described in [6,7].

---

* Corresponding author
   Email addresses: `ccecka@stanford.edu` (Cris Cecka), `darve@stanford.edu`
(Eric Darve).

7 The application of the boundary element method to solve the integral form of
8 the Helmholtz equation results in a dense linear system which can be solved
9 by iterative methods such as GMRES or BCGSTAB. These methods require
10 computing dense matrix-vector products which, using a direct implementa-
11 tion, are performed in $\mathcal{O}(N^2)$ floating point operations. The FMM uses an
12 approximation of the dense matrix to perform the product in $\mathcal{O}(N \log N)$ or
13 $\mathcal{O}(N \log^2 N)$ operations. This approximation is constructed from close pair in-
14 teractions and far field approximations represented by spherical integrals that
15 are accumulated and distributed through the domain via an octree.

16 There are a number of difficulties in implementing the FMM, each of which
17 must be carefully considered and optimized to achieve the improved com-
18 plexity. The most significant complication is that the quadrature sampling
19 rate must increase with the size of the box in the octree, requiring interpo-
20 lation and anterpolation algorithms to transform the data between spherical
21 quadratures of different levels of the tree. Local algorithms such as Lagrange
22 interpolation and techniques which sparsify interpolant matrices are fast, but
23 incur significant errors [8,7]. Spherical harmonic transforms are global inter-
24 polation schemes and are exact but require fast versions for efficiency of the
25 FMM. Many of these fast spherical transform algorithms are only approxi-
26 mate, complicated to implement, and not always stable [9–11].

27 In this paper, we use a multipole expansion which allows the use of 2D fast
28 Fourier transforms (FFT) in the spherical coordinate system (ϕ, θ). The main
29 advantages are two fold: i) high performance libraries are available for FFTs
30 on practically all computer platforms, resulting in accurate, robust, and fast
31 interpolation algorithms; ii) the resulting error analysis is simplified leading to
32 sharp a priori error bounds on the calculation. One of the difficulties in using
33 FFTs is that we are forced to use a uniform distribution of points along ϕ
34 and θ in the spherical quadrature. This leads to a much increased quadrature
35 size for a given accuracy compared to the original spherical harmonics-based
36 FMM. The reason is as follows. The multipole expansion in the high frequency
37 regime is derived from:

$$\frac{e^{\iota \kappa |\boldsymbol{r} + \boldsymbol{r}_0|}}{|\boldsymbol{r} + \boldsymbol{r}_0|} = \int_{\phi=0}^{2\pi} \int_{\theta=0}^{\pi} e^{\iota \kappa \hat{\boldsymbol{s}} \cdot \boldsymbol{r}} \, T_{\ell, \boldsymbol{r}_0}(\hat{\boldsymbol{s}}) \, \sin(\theta) \, d\theta d\phi$$

38 where $\hat{\boldsymbol{s}} = [\cos(\phi)\sin(\theta), \sin(\phi)\sin(\theta), \cos(\theta)]$ is the spherical unit vector. It
39 is apparent that we are integrating along θ a function which has period 2π.
40 However the bounds of the integral are 0 to π, over which interval the func-
41 tion has a discontinuity in its derivative. This results in a slow decay of the
42 Fourier spectrum (essentially $1/\text{frequency}^2$) and consequently a large number
43 of quadrature points along θ are required.

44 We propose to use a variant of the scheme by J. Sarvas in [12] whereby the

2

integration is extended from 0 to 2π and the integrand modified:

$$\frac{e^{\iota\kappa|\boldsymbol{r}+\boldsymbol{r}_0|}}{|\boldsymbol{r}+\boldsymbol{r}_0|} = \frac{1}{2}\int_{\phi=0}^{2\pi}\int_{\theta=0}^{2\pi} e^{\iota\kappa\hat{\boldsymbol{s}}\cdot\boldsymbol{r}}\, T_{\ell,\boldsymbol{r}_0}(\hat{\boldsymbol{s}})\,|\sin(\theta)|\,d\theta d\phi$$

We will describe in more details how an efficient scheme can be derived from this equation. The key property is that $e^{\iota\kappa\hat{\boldsymbol{s}}\cdot\boldsymbol{r}}$ is approximately band-limited in θ and therefore it is possible to remove the high frequency components of $T_{\ell,\boldsymbol{r}_0}(\hat{\boldsymbol{s}})|\sin(\theta)|$ without affecting the accuracy of the approximation. Using this smooth transfer function, which is now band-limited in Fourier space, the number of quadrature points can be reduced dramatically. We show that the resulting number of quadrature points is reduced by about 40% compared to the original spherical harmonics-based FMM. Consequently, we now have a scheme which requires few quadrature points and enables the use of efficient FFT routines. The approach in [12] is similar. However, rather than smoothing $T_{\ell,\boldsymbol{r}_0}(\hat{\boldsymbol{s}})\,|\sin(\theta)|$ once during the precomputation phase as we will detail shortly, Sarvas instead incorporates the $|\sin(\theta)|$ factor during the run-time phase of the FMM in the application of the transfer pass. This requires extra interpolations/anterpolations and a quadrature approximately twice the size during the time-critical transfer pass of the algorithm. The method presented in this paper retains a diagonal transfer pass operator and provides improved control over the error.

We derive a new a priori error analysis which incorporates both effects from truncation of the Gegenbauer series (a problem well analyzed [6]) and the numerical quadrature. Our algorithm to predict the error is very sharp. This allows choosing the minimal number of quadrature points to guarantee a prescribed error. By comparison, the conventional approach can be shown to result in conservative error bounds. That is, the number of quadrature points is usually over-estimated and the error may be significantly below the target accuracy. Although not considered in this paper, our error analysis approach can also be applied to the spherical harmonics-based FMM to yield similarly accurate error bounds. This has practical importance since it allows guaranteeing the error in the calculation while reducing the computational cost.

The paper is organized as follows. In section 2, we introduce the critical parts of the classical FMM algorithm including the Gegenbauer series truncation (2.1), the spherical quadrature (2.2), and a short overview of interpolation/anterpolation strategies (2.3). In section 2.4, the asymptotic complexity of the FMM is discussed. Section 3 details the Fourier basis approach. The transfer function must be modified to lower the computational cost and obtain a competitive scheme, as detailed in section 3.1. Section 3.2 applies an analysis of the integration error to derive an algorithm that determines a quadrature to yield an FMM with a prescribed error tolerance. The FFT based interpolation and anterpolation algorithms are described in section 3.3 and numerical results are given in section 3.4. Table 1 lists the notations used in this paper.

Notation	Description
λ	wavelength
κ	wavenumber, $2\pi/\lambda$
θ	polar angle
ϕ	azimuthal angle
$\overline{a}$	complex conjugate of a
$\boldsymbol{x}$	vector in $\mathbb{R}^3$, $\boldsymbol{x} = \lvert\boldsymbol{x}\rvert\,\hat{\boldsymbol{x}}$
$\boldsymbol{x} \cdot \boldsymbol{y}$	inner product, $\boldsymbol{x} \cdot \boldsymbol{y} = \lvert\boldsymbol{x}\rvert\,\lvert\boldsymbol{y}\rvert\cos(\phi_{\boldsymbol{x},\boldsymbol{y}})$
$\boldsymbol{M}$	matrix in $\mathbb{R}^{n\times m}$ with elements M_{ij}
S^2	sphere, $\{\hat{\boldsymbol{s}} \in \mathbb{R}^3 : \lvert\hat{\boldsymbol{s}}\rvert = 1\}$
j_n	spherical Bessel function of the first kind
y_n	spherical Bessel function of the second kind
$h_n^{(1)}$	spherical Hankel function of the first kind
P_n	Legendre polynomial
Y_n^m	normalized spherical harmonic of degree n, order m
$\mathcal{F}(m; f)$	mth frequency of the Fourier transform of f

Table 1. Table of notations

2 The Multilevel Fast Multipole Method

The FMM reduces the computational complexity of the matrix-vector multi-plication

$$\sigma_i = \sum_{j\neq i} \frac{e^{\imath\kappa\lvert\boldsymbol{x}_i-\boldsymbol{x}_j\rvert}}{\lvert\boldsymbol{x}_i - \boldsymbol{x}_j\rvert}\psi_j = \sum_j M_{ij}\,\psi_j \tag{1}$$

for $i, j = 1, \ldots, N$ from $\mathcal{O}(N^2)$ to $\mathcal{O}(N\log^2 N)$. This improvement is based on the Gegenbauer series

$$\frac{e^{\imath\kappa\lvert\boldsymbol{r}+\boldsymbol{r}_0\rvert}}{\lvert\boldsymbol{r}+\boldsymbol{r}_0\rvert} = \imath\kappa\sum_{n=0}^{\infty}(-1)^n(2n+1)h_n^{(1)}(\kappa\,\lvert\boldsymbol{r}_0\rvert)j_n(\kappa\,\lvert\boldsymbol{r}\rvert)P_n(\hat{\boldsymbol{r}}\cdot\hat{\boldsymbol{r}}_0) \tag{2}$$

The series converges absolutely and uniformly for $\lvert\boldsymbol{r}_0\rvert \geq \frac{2}{\sqrt{3}}\lvert\boldsymbol{r}\rvert$ and has been studied extensively in [13,14].

4

92 Truncating the Gegenbauer series at ℓ and using the identity

$$\int_{S^2} e^{\imath\kappa\hat{\boldsymbol{s}}\cdot\boldsymbol{r}} P_n(\hat{\boldsymbol{s}}\cdot\hat{\boldsymbol{r}}_0)\, dS(\hat{\boldsymbol{s}}) = 4\pi\imath^n j_n(\kappa\,|\boldsymbol{r}|) P_n(\hat{\boldsymbol{r}}\cdot\hat{\boldsymbol{r}}_0)$$

93 where the integral is over the unit sphere S^2, then

$$\imath\kappa \sum_{n=0}^{\ell} (-1)^n (2n+1) h_n^{(1)}(\kappa\,|\boldsymbol{r}_0|) j_n(\kappa\,|\boldsymbol{r}|) P_n(\hat{\boldsymbol{r}}\cdot\hat{\boldsymbol{r}}_0) = \int_{S^2} e^{\imath\kappa\hat{\boldsymbol{s}}\cdot\boldsymbol{r}} T_{\ell,\boldsymbol{r}_0}(\hat{\boldsymbol{s}})\, dS(\hat{\boldsymbol{s}})$$

$$(3)$$

94 where the transfer function, $T_{\ell,\boldsymbol{r}_0}(\hat{\boldsymbol{s}})$, is defined as

$$T_{\ell,\boldsymbol{r}_0}(\hat{\boldsymbol{s}}) = \frac{\imath\kappa}{4\pi} \sum_{n=0}^{\ell} \imath^n (2n+1) h_n^{(1)}(\kappa\,|\boldsymbol{r}_0|) P_n(\hat{\boldsymbol{s}}\cdot\hat{\boldsymbol{r}}_0). \tag{4}$$

95 Consider two disjoint clusters of points $\{\boldsymbol{x}_i \mid i \in A\}$ and $\{\boldsymbol{x}_i \mid i \in B\}$ with radii
96 $r_A \geq r_B \geq 0$ and centers $\boldsymbol{c}_A$ and $\boldsymbol{c}_B$ respectively. If $|\boldsymbol{c}_A - \boldsymbol{c}_B| \geq \frac{2}{\sqrt{3}}(r_A + r_B)$,
97 then the matrix-vector product (1) is accelerated by using the approximation

$$\boldsymbol{M} \approx \begin{bmatrix} \boldsymbol{M}_{AA} & \widetilde{\boldsymbol{M}}_{AB} \\ \widetilde{\boldsymbol{M}}_{BA} & \boldsymbol{M}_{BB} \end{bmatrix}$$

98 where

$$\boldsymbol{M}_{AA} = [M_{ij}]_{i\in A, j\in A}$$
$$\widetilde{\boldsymbol{M}}_{AB} = \left[\int_{S^2} e^{\imath\kappa\hat{\boldsymbol{s}}\cdot(\boldsymbol{c}_A - \boldsymbol{x}_i)} T_{\ell,\boldsymbol{c}_B - \boldsymbol{c}_A}(\hat{\boldsymbol{s}}) e^{\imath\kappa\hat{\boldsymbol{s}}\cdot(\boldsymbol{x}_j - \boldsymbol{c}_B)}\, dS(\hat{\boldsymbol{s}}) \right]_{i\in A, j\in B}$$

99 For $i \in A$, this corresponds to computing

$$\sigma_i = \int_{S^2} e^{\imath\kappa\hat{\boldsymbol{s}}\cdot(\boldsymbol{c}_A - \boldsymbol{x}_i)} T_{\ell,\boldsymbol{c}_B - \boldsymbol{c}_A}(\hat{\boldsymbol{s}}) \left[\sum_{j\in B} e^{\imath\kappa\hat{\boldsymbol{s}}\cdot(\boldsymbol{x}_j - \boldsymbol{c}_B)} \psi_j \right] dS(\hat{\boldsymbol{s}})$$
$$+ \sum_{j\in A} \frac{e^{\imath\kappa|\boldsymbol{x}_i - \boldsymbol{x}_j|}}{|\boldsymbol{x}_i - \boldsymbol{x}_j|} \psi_j + \varepsilon_\ell$$

100 where ε_ℓ is the error introduced by the Gegenbauer series truncation.

101 The reduced computational complexity of the FMM is achieved by construct-
102 ing a tree of nodes, typically an octree, over the domain of the source and field
103 points. Let $U_\alpha^l(\hat{\boldsymbol{s}})$ be the outgoing field for B_α^l, the box α of the tree in level
104 $l \in [0, L]$ with center $\boldsymbol{c}_\alpha^l$. The method is initialized by computing the outgoing
105 plane-wave expansions for each cluster contained in a leaf of the tree:

$$U_\alpha^L(\hat{\boldsymbol{s}}) = \sum_{i,\, \boldsymbol{x}_i \in B_\alpha^l} e^{\imath\kappa\hat{\boldsymbol{s}}\cdot(\boldsymbol{x}_i - \boldsymbol{c}_\alpha^L)} \psi_i$$

These outgoing expansions are then aggregated upward through the tree by accumulating the product of the child cluster expansions with the diagonal plane-wave translation function:

$$U_\alpha^l(\hat{\boldsymbol{s}}) = \sum_{\beta,\, B_\beta^{l+1} \subset B_\alpha^l} U_\beta^{l+1}(\hat{\boldsymbol{s}}) e^{\imath \kappa \hat{\boldsymbol{s}} \cdot (\boldsymbol{c}_\beta^{l+1} - \boldsymbol{c}_\alpha^l)}$$

Incoming plane-wave expansions, $I_\alpha^l(\hat{\boldsymbol{s}})$ of box B_α^l, are computed from the outgoing by multiplication with the diagonal transfer function:

$$I_\alpha^l(\hat{\boldsymbol{s}}) = \sum_\beta U_\beta^l(\hat{\boldsymbol{s}}) T_{\ell, \boldsymbol{c}_\beta^l - \boldsymbol{c}_\alpha^l}(\hat{\boldsymbol{s}})$$

where the parent of B_β^l is a neighbor of the parent of B_α^l, and B_β^l is not a neighbor of B_α^l. The incoming plane-waves are then disaggregated downward through the tree to compute the local field $D_\alpha^l(\hat{\boldsymbol{s}})$:

$$D_\alpha^l(\hat{\boldsymbol{s}}) = D_\beta^{l-1}(\hat{\boldsymbol{s}}) e^{\imath \kappa \hat{\boldsymbol{s}} \cdot (\boldsymbol{c}_\beta^{l-1} - \boldsymbol{c}_\alpha^l)} + I_\alpha^l(\hat{\boldsymbol{s}})$$

where $B_\alpha^l \subset B_\beta^{l-1}$. At the finest level, the integration over the sphere is finally performed and added to the near field contribution to determine the field value at the N field points:

$$\sigma_i = \int_{S^2} D_\alpha^L(\hat{\boldsymbol{s}}) e^{\imath \kappa \hat{\boldsymbol{s}} \cdot (\boldsymbol{c}_\alpha^L - \boldsymbol{x}_i)} \, dS(\hat{\boldsymbol{s}}) + \sum_{j \neq i} \frac{e^{\imath \kappa |\boldsymbol{x}_i - \boldsymbol{x}_j|}}{|\boldsymbol{x}_i - \boldsymbol{x}_j|} \psi_j$$

where $\boldsymbol{x}_i \in B_\alpha^L$ and $\boldsymbol{x}_j \in B_\beta^L \cup B_\alpha^L$ where B_β^L is any neighbor of B_α^L.

2.1   Truncation Parameter in the FMM

The truncation parameter ℓ must be chosen so that the Gegenbauer series (2) is converged to a desired accuracy. However, for $n > x$, $j_n(x)$ decreases super-exponentially while $h_n^{(1)}(x)$ diverges. The divergence of the Hankel function causes the transfer function to oscillate wildly and become numerically unstable. Even though the expansion converges, roundoff errors will adversely affect the accuracy if ℓ is too large. Thus, while one must choose $\ell > \kappa |\boldsymbol{r}|$ so that sufficient convergence is achieved for the plane-wave, it must also be small enough to avoid the divergence of the transfer function.

The selection the truncation parameter ℓ has been studied extensively and and a number of procedures for selecting it have been proposed.

The empirical formula

$$\ell \approx \kappa |\boldsymbol{r}| + C(\varepsilon) \log(\pi + \kappa |\boldsymbol{r}|)$$

appears to have been first proposed by Rokhlin [2]. This was considered and revised by Darve [14] using a detailed asymptotic analysis of the Gegenbauer series.

The excess bandwidth formula (EBF) is derived from the convergence of the plane-wave expansion and is presented in [6]. To determine an appropriate truncature, the spectrum of a plane wave

$$e^{\imath\kappa\hat{\boldsymbol{s}}\cdot\boldsymbol{r}} = e^{\imath\kappa|\boldsymbol{r}|\cos(\phi_{\hat{\boldsymbol{s}},\boldsymbol{r}})} = \sum_{n=-\infty}^{\infty} \imath^n J_n(\kappa\,|\boldsymbol{r}|)e^{\imath n\phi_{\hat{\boldsymbol{s}},\boldsymbol{r}}} \tag{5}$$

where $\phi_{\hat{\boldsymbol{s}},\boldsymbol{r}}$ is the angle between $\hat{\boldsymbol{s}}$ and $\boldsymbol{r}$, is used to estimate how many terms in the series are needed before the error in the summation is exponentially small. It can be shown that when $n \to \infty$ and $x \sim \mathcal{O}(n)$,

$$J_n(x) \sim \frac{e^{\sqrt{n^2-x^2}-n\cosh^{-1}(n/x)}}{\sqrt{2\pi(n^2-x^2)}}$$

which decays exponentially fast when $n > x$. Let $n/x = 1 + \delta$ where $\delta \ll 1$. Then $\cosh^{-1}(n/x) \sim \sqrt{2\delta}$ and $\sqrt{n^2-x^2} \sim x\sqrt{2\delta}$. Thus, the above becomes

$$J_n(x) \sim \frac{e^{(x-n)\sqrt{2\delta}}}{2x\sqrt{\pi\delta}} = \frac{e^{-x\sqrt{2}\delta^{3/2}}}{2x\sqrt{\pi\delta}}$$

This expression is exponentially small when $x\delta^{3/2} \gg 1$, or $\delta = Cx^{-2/3}$, where $C \gg 1$. That is, when

$$\frac{n}{x} - 1 \approx Cx^{-2/3}$$

Therefore, the number of terms we need can be approximated as

$$\ell \approx \kappa\,|\boldsymbol{r}| + C(\kappa\,|\boldsymbol{r}|)^{1/3}$$

An empirically determined common choice is $C = 1.8(d_0)^{2/3}$, where d_0 is the desired number of digits of accuracy. The EBF is one of the most popular choices to select the truncation parameter.

The desired number of digits of accuracy cannot always be achieved due to the divergent nature of the transfer function. Thus, the EBF fails when high accuracy is desired or the box size and box separation are small. Some modifications in this regime are proposed in [15].

A direct numerical computation of the Gegenbauer truncation error ℓ could be computed or approximated

$$\varepsilon_G(\ell) = \imath\kappa \sum_{n=\ell+1}^{\infty} (-1)^n(2n+1)h_n^{(1)}(\kappa\,|\boldsymbol{r}_0|)j_n(\kappa\,|\boldsymbol{r}|)P_n(\hat{\boldsymbol{r}}\cdot\hat{\boldsymbol{r}}_0)$$

7

153 As Carayol and Collino showed in [13], an upper bound of this error for large
154 values of $|\boldsymbol{r}|$ is obtained when $P_n(\hat{\boldsymbol{r}} \cdot \hat{\boldsymbol{r}}_0) = P_n(\pm 1) = (\pm 1)^n$ so that

$$|\varepsilon_G| \lesssim \kappa \left| \sum_{n=\ell+1}^{\infty} (\mp 1)^n (2n+1) h_n^{(1)}(\kappa |\boldsymbol{r}_0|) j_n(\kappa |\boldsymbol{r}|) \right|$$

155 which they showed can be computed in closed form

$$= \kappa \frac{\sqrt{|\boldsymbol{r}| \, |\boldsymbol{r}_0|}}{|\boldsymbol{r}_0| \pm |\boldsymbol{r}|} \frac{\pi}{2} \left| H_{\ell+\frac{3}{2}}(\kappa |\boldsymbol{r}_0|) J_{\ell+\frac{1}{2}}(\kappa |\boldsymbol{r}|) \pm H_{\ell+\frac{1}{2}}(\kappa |\boldsymbol{r}_0|) J_{\ell+\frac{3}{2}}(\kappa |\boldsymbol{r}|) \right|$$

$$= \kappa^2 \frac{|\boldsymbol{r}| \, |\boldsymbol{r}_0|}{|\boldsymbol{r}_0| \pm |\boldsymbol{r}|} \left| h_{\ell+1}^{(1)}(\kappa |\boldsymbol{r}_0|) j_\ell(\kappa |\boldsymbol{r}|) \pm h_\ell^{(1)}(\kappa |\boldsymbol{r}_0|) j_{\ell+1}(\kappa |\boldsymbol{r}|) \right|$$

156 This fails for small $|\boldsymbol{r}|$ when the upper bound is instead given by the $\hat{\boldsymbol{r}} \cdot \hat{\boldsymbol{r}}_0$
157 which causes the oscillation of $P_n(\hat{\boldsymbol{r}} \cdot \hat{\boldsymbol{r}}_0)$ to compensate for the oscillation of
158 $(-1)^n h_n^{(1)}(\kappa |\boldsymbol{r}_0|) j_n(\kappa |\boldsymbol{r}|)$.

159 Using the EBF as an initial guess for ℓ and refining the choice using the above
160 closed form when $|\boldsymbol{r}|$ is sufficiently large is a simple algorithm which yields an
161 accurate value for ℓ.

162 Carayol and Collino in [16] and [13] present an in-depth analysis of the Jacobi-
163 Anger series and the Gegenbauer series. They find the asymptotic formula

$$\ell \approx v - \frac{1}{2} + \left(\frac{1}{2}\right)^{5/3} W^{2/3} \left(\frac{v}{4\varepsilon^6} \left(\frac{1+\alpha}{1-\alpha}\right)^{3/2}\right)$$

164 where $v = \kappa |\boldsymbol{r}|$, $\alpha = u/v$, $u = \kappa |\boldsymbol{r}_0|$, and $W(x)$ is the Lambert function
165 defined as the solution to

$$W(x) e^{W(x)} = x \qquad x > 0$$

166 This appears to be near optimal for large box sizes.

167 The errors introduced by this truncation have been investigated in other pa-
168 pers including [8,13,14].

169 *2.2   Spherical Quadrature in the FMM*

170 With the expansion formula

$$(2n+1) P_n(\hat{\boldsymbol{p}} \cdot \hat{\boldsymbol{q}}) = 4\pi \sum_{m=-n}^{n} \overline{Y_n^m(\hat{\boldsymbol{p}})} \, Y_n^m(\hat{\boldsymbol{q}})$$

8

the transfer function (4) can be expressed as

$$T_{\ell,\boldsymbol{r}_0}(\hat{\boldsymbol{s}}) = \imath\kappa \sum_{n=0}^{\ell} \imath^n h_n^{(1)}(\kappa\,|\boldsymbol{r}_0|) \sum_{m=-n}^{n} \overline{Y_n^m(\hat{\boldsymbol{r}}_0)}\, Y_n^m(\hat{\boldsymbol{s}}) = \sum_{n=0}^{\ell} \sum_{m=-n}^{n} t_n^m Y_n^m(\hat{\boldsymbol{s}}) \quad (6)$$

Similarly, the Jacobi-Anger series,

$$e^{\imath\kappa\hat{\boldsymbol{s}}\cdot\boldsymbol{r}} = \sum_{n=0}^{\infty} \imath^n(2n+1)j_n(\kappa\,|\boldsymbol{r}|)P_n(\hat{\boldsymbol{s}}\cdot\hat{\boldsymbol{r}}) \quad (7)$$

becomes the spherical harmonic series

$$e^{\imath\kappa\hat{\boldsymbol{s}}\cdot\boldsymbol{r}} = 4\pi \sum_{n=0}^{\infty} \imath^n j_n(\kappa\,|\boldsymbol{r}|) \sum_{m=-n}^{n} \overline{Y_n^m(\hat{\boldsymbol{r}})} Y_n^m(\hat{\boldsymbol{s}}) = \sum_{n=0}^{\infty} \sum_{m=-n}^{n} e_n^m Y_n^m(\hat{\boldsymbol{s}}) \quad (8)$$

The error analysis simplifies if we use a scheme which exactly integrates spherical harmonics, Y_l^m, up to some order. Below, we enumerate a number of choices that have previously been studied.

(1) The simplest choice are sample points chosen uniformly in ϕ and θ. However, this choice does not accurately integrate the spherical harmonics and requires approximately twice as many points as the Gauss-Legendre quadrature below [7].

(2) The most common choice of sample points are uniform points for ϕ and Gauss-Legendre points for θ. With $N+1$ uniform points in the ϕ direction and $\frac{N+1}{2}$ Gauss-Legendre points in the θ direction, all Y_n^m, $-n \leq m \leq n$, $0 \leq n \leq N$ are integrated exactly [7,8].

(3) McLaren in [17] developed optimal choices of samples for general functions on S^2 based on equally spaced points and derived from invariants of finite groups of rotations. He also proposes a method for constructing equally weighted integration formulas on sets of any desired number of points by taking the union of icosahedral configurations.

2.3 Interpolation and Anterpolation in the FMM

The quadrature sampling rate depends on the spectral content of the diagonalized translation operator, $e^{\imath\kappa\hat{\boldsymbol{s}}\cdot\boldsymbol{r}}$. These plane-waves contain more oscillatory modes as we go up in the tree. Its coefficient in the spherical harmonic expansion e_n^m [equation (8)] decreases super-exponentially roughly for $n \gtrsim \kappa\,|\boldsymbol{r}|$. Therefore, as we go up the tree in the aggregation step and $|\boldsymbol{r}|$ becomes larger, a larger quadrature is required to resolve these higher modes. These modes must be resolved since they interact with the modes in the transfer function, which do not significantly decay as ℓ increases.

₁₉₉ Similarly, as we go down the tree in the disaggregation step, $|\boldsymbol{r}|$ becomes
₂₀₀ smaller and the higher modes of the incoming field make vanishingly small
₂₀₁ contributions to the integral as a consequence of Parseval's theorem. Thus, as
₂₀₂ the incoming field is disaggregated down the tree, a smaller quadrature can be
₂₀₃ used to resolve it. This makes the integration faster and is actually required to
₂₀₄ achieve an optimal asymptotic running time. See section 2.4 and appendix A.

₂₀₅ There have been several approaches to performing the interpolation and an-
₂₀₆ terpolation between levels in the FMM. Below, we enumerate a number of
₂₀₇ choices that have previously been studied.

₂₀₈ (1) General interpolation algorithms like Lagrange interpolation or B-splines
₂₀₉ are fast and provide for simple error analysis. In [8] it is shown that the
₂₁₀ error induced from Lagrange interpolation decreases exponentially as the
₂₁₁ number of interpolation points is increased for a given function of finite
₂₁₂ bandwidth. Thus, there is a trade-off between error and speed.
₂₁₃ (2) For a set of quadrature points (ϕ_k, θ_k), $k = 1, \ldots, K$ with respective
₂₁₄ weights ω_k and corresponding function value f_k, a spherical harmonic
₂₁₅ transform maps f_k to a new quadrature $(\phi'_{k'}, \theta'_{k'})$, $k' = 1, \ldots, K'$ via the
₂₁₆ linear transformation

$$f_{k'} = \sum_{m, l \leq K} Y_l^m(\phi'_{k'}, \theta'_{k'}) \sum_k \omega_k \overline{Y_l^m(\phi_k, \theta_k)} f_k = \sum_k A_{k'k} f_k \qquad (9)$$

₂₁₇ This transform has nice properties analogous to those of the Fourier trans-
₂₁₈ form. A direct computation requires $\mathcal{O}(KK')$ operations which would
₂₁₉ result in an $\mathcal{O}(N^2)$ FMM (see appendix A). Fast spherical transforms
₂₂₀ (FST) have been developed in [9–11,18] and applied to the FMM in [19].
₂₂₁ Using the FST reduces the interpolation and anterpolation procedures
₂₂₂ to $\mathcal{O}(K \log^2 K)$, which results in an $\mathcal{O}(N \log^2 N)$ FMM. However, the
₂₂₃ accuracy and stability of these algorithms remain in question.
₂₂₄ (3) Approximations of the spherical transform have also been investigated in
₂₂₅ [20,7]. The interpolation matrix $A_{k'k}$ in (9) can be sparsified in a number
₂₂₆ of ways to provide an interpolation/anterpolation method that scales as
₂₂₇ $\mathcal{O}(K)$ with controllable relative error.
₂₂₈ (4) Many other interpolation schemes exist with varying running times and
₂₂₉ errors. Rokhlin presents a fast polynomial interpolator based on the fast
₂₃₀ multipole method in [21]. See also [22].

₂₃₁ *2.4 Asymptotic Complexity*

₂₃₂ In order to resolve a sufficient number of spherical harmonics, the number
₂₃₃ of points in the ϕ and θ directions must be $\mathcal{O}(\ell) = \mathcal{O}(\kappa a)$, where a is the
₂₃₄ side length of the box. Therefore, the total number of quadrature points is

235 $\mathcal{O}(\ell^2) = \mathcal{O}((\kappa a)^2)$. If a_0 is the side length of a box at the root of the octree
236 (level $l = 0$), then the side length of a box at level l is $a_l = 2^{-l} a_0$ and has
237 $\mathcal{O}((\kappa 2^{-l} a_0)^2)$ quadrature points.

238 With these parameters defined, the asymptotic complexity of the FMM can
239 be determined by carefully counting the number of operations required in each
240 step. The is done in detail in Appendix A and the results are discussed below.

241 In modeling a uniform distribution of point scatterers over a volumetric do-
242 main, we take $\mathcal{O}(N) = \mathcal{O}((\kappa a_0)^3)$ and $L \sim \log(N^{1/3})$ to achieve a total algo-
243 rithmic complexity of $\mathcal{O}(N \log N)$ using fast global interpolation methods. By
244 using local methods, this can be reduced to $\mathcal{O}(N)$.

245 In modeling the scattering from the surface of an object using a uniform dis-
246 tribution of basis functions, we take $\mathcal{O}(N) = \mathcal{O}((\kappa a_0)^2)$ and $L \sim \log(N^{1/2})$ to
247 achieve a total algorithmic complexity of $\mathcal{O}(N \log^2(N))$ using fast global inter-
248 polation methods. By using local methods, this can be reduced to $\mathcal{O}(N \log N)$.

249 It should be noted again that for a given, fixed κa_0 there is a minimum size
250 for the leaves of the tree. Below this critical size, $h_n^{(1)}$ oscillates wildly causing
251 numerical instability in the transfer function. Therefore, when N is very large
252 and the number of levels is saturated, this analysis fails and the algorithm is
253 dominated by the $\mathcal{O}(N^2)$ computation of the close field contribution, albeit
254 with orders of magnitude speedup over a direct method. In the case when κa_0
255 is too small, broadband FMMs have been developed as detailed in [23,24].
256 However, in many applications κa_0 is large enough to allow for all practi-
257 cal L and N. Furthermore, by keeping the number of points per wavelength
258 constant, the $\mathcal{O}(N \log N)$ behavior can always be achieved.

259 ## 3   Fourier Based Multilevel Fast Multipole Method

260 The Fourier based fast multipole method is based on the identity

$$\int_{S^2} e^{\imath \kappa \hat{\boldsymbol{s}} \cdot \boldsymbol{r}} \, T_{\ell, \boldsymbol{r}_0}(\hat{\boldsymbol{s}}) \, dS(\hat{\boldsymbol{s}}) = \int_0^{2\pi} \int_0^{2\pi} e^{\imath \kappa \hat{\boldsymbol{s}} \cdot \boldsymbol{r}} \, T_{\ell, \boldsymbol{r}_0}^{\mathrm{s}}(\hat{\boldsymbol{s}}) \, d\phi d\theta \qquad (10)$$

261 where $\hat{\boldsymbol{s}} = [\cos(\phi)\sin(\theta), \, \sin(\phi)\sin(\theta), \, \cos(\theta)]$ and

$$T_{\ell, \boldsymbol{r}_0}^{\mathrm{s}}(\hat{\boldsymbol{s}}) = \frac{1}{2} \, T_{\ell, \boldsymbol{r}_0}(\hat{\boldsymbol{s}}) \, |\sin(\theta)| \qquad (11)$$

262 is the modified transfer function. Noting that the integrand is continuous and
263 periodic, this formulation of the problem suggests the use of the Fourier func-
264 tions $\{e^{\imath n\phi} e^{\imath m\theta}\}$ which form an orthonormal basis of $L^2([0, 2\pi] \times [0, 2\pi])$. This

11

allows i) using two dimensional uniform quadratures; ii) fast Fourier transforms in the interpolation and anterpolation steps; iii) spectral arguments in the error analysis. Of these advantages, the most important is that the FFT interpolations and anterpolations are fast and exact. Since there is no interpolation error, only the finite quadrature and the truncation of the Gegenbauer series introduce error to the final solution. Thus, the error analysis is simplified and we will determine in this paper precise bounds on the final error. In fact, our error analysis is fairly general and can be extended to the classical FMM with schemes that exactly integrate spherical harmonics (see direct and fast global methods in section 2.3 and appendix A). The result is a fast, easy to implement, and controllable version of the FMM, which we detail in the following sections.

3.1   Computing the Modified Transfer Function

Select a uniform quadrature with points (ϕ_i, θ_j) defined by

$$\phi_i = 2\pi \frac{i}{N_\phi} \qquad \theta_j = 2\pi \frac{j}{N_\theta}$$

Noting that the plane wave $e^{\iota \kappa \hat{\boldsymbol{s}} \cdot \boldsymbol{r}}$ and $T^{\mathrm{s}}_{\ell, \boldsymbol{r}_0}(\hat{\boldsymbol{s}})$ both have spherical symmetry,

$$f(\hat{\boldsymbol{s}})\Big|_{\hat{\boldsymbol{s}}(\phi, \theta)} = f(\hat{\boldsymbol{s}})\Big|_{\hat{\boldsymbol{s}}(\pi + \phi, 2\pi - \theta)},$$

the computational and memory cost are reduced by computing and storing only half of the quadrature points.

Additionally, in an FMM with a single octree, there are 316 distinct transfer vectors $\boldsymbol{r}_0$ per level. By enforcing symmetries in the quadrature, the number of modified transfer functions that must be precomputed is reduced. Specifically, by requiring N_θ to be a multiple of 2 and N_ϕ to be a multiple of 4, we enforce reflection symmetries in the $z = 0$, $x = 0$, $y = 0$, $x = y$, and $x = -y$ planes. This reduces the number of modified transfer functions that need to be precomputed from 316 per level to 34 – saving a factor of 9.3 in memory and costing a negligible permutation of the values of a computed modified transfer function. See Fig. 1.

A key step to constructing a fast algorithm is to remove the high frequencies in $T^{\mathrm{s}}_{\ell, \boldsymbol{r}_0}(\hat{\boldsymbol{s}})$ whose contribution to the final result is negligible. This reduces the number of needed quadrature points considerably. If $T^{\mathrm{s}}_{\ell, \boldsymbol{r}_0}(\hat{\boldsymbol{s}})$ were simply sampled, significant aliasing would occur unless we used an unreasonably large
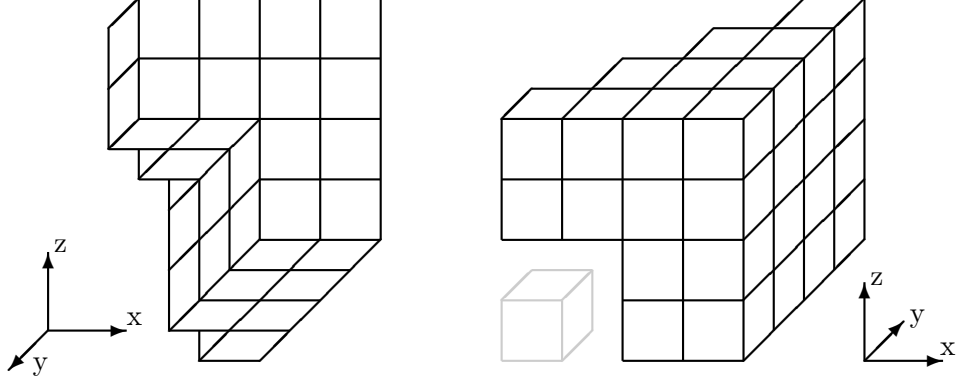
12

Fig. 1. The center of each box represents one transfer vector $\boldsymbol{r}_0$ which must be computed. The pictures on the left and right panels represent the same set of boxes viewed under two different angles. Due to the symmetries of the quadrature, we need only compute transfer vectors with $x, y, z \geq 0$ and $x \geq y$. We therefore end up with essentially half of an octant. Specifically, 34 transfer vectors are required; they can be reflected into any of the 316 needed.

295   quadrature. This is due to the slow decay of the Fourier series of $|\sin(\theta)|$,

$$\mathcal{F}(m; |\sin(\theta)|) = \frac{(-1)^m + 1}{\pi(1 - m^2)} = \begin{cases} \frac{2}{\pi} \frac{1}{1-m^2} & \text{if } m \text{ even} \\ 0 & \text{if } m \text{ odd} \end{cases}$$

296   Since the spectrum of the plane-wave function in equation (5) decays very
297   rapidly for $n \gtrsim \kappa |\boldsymbol{r}|$, the high frequencies in θ of $T^{\mathrm{s}}_{\ell,\boldsymbol{r}_0}(\hat{\boldsymbol{s}})$ do not contribute
298   to the final integral as a result of Parseval's theorem. By removing these
299   frequencies from the modified transfer function, a smaller quadrature can be
300   used without affecting the final result.

301   Suppose we have chosen a quadrature characterized by (N_θ, N_ϕ). With this
302   quadrature we are able to exactly resolve the frequencies in $e^{\iota \kappa \hat{\boldsymbol{s}} \cdot \boldsymbol{r}}$ between
303   $-N_\theta/2 + 1$ and $N_\theta/2 - 1$ to the integral in equation (10). Consequently, we
304   need to exactly calculate a band limited approximation of $T^{\mathrm{s}}_{\ell,\boldsymbol{r}_0}$, called $T^{\mathrm{s},L}_{\ell,\boldsymbol{r}_0}$,
305   such that:

$$\mathcal{F}(m; T^{\mathrm{s},L}_{\ell,\boldsymbol{r}_0}) = \begin{cases} \mathcal{F}(m; T^{\mathrm{s}}_{\ell,\boldsymbol{r}_0}), & \text{if } -N_\theta/2 + 1 \leq m \leq N_\theta/2 - 1 \\ 0, & \text{otherwise} \end{cases}$$

306   Since $T_{\ell,\boldsymbol{r}_0}$ is bandlimited in θ with bandwidth $2\ell + 1$, only the frequencies
307   $|m| \leq N_\theta/2 - 1 + \ell$ of $|\sin(\theta)|$ contribute to the $N_\theta - 1$ frequencies of $T^{\mathrm{s},L}_{\ell,\boldsymbol{r}_0}$.
308   Therefore, the low-pass modified transfer function $T^{\mathrm{s},L}_{\ell,\boldsymbol{r}_0}$ can be computed using
309   the following pseudo-code:

1 **for** ϕ_i, $0 \leq i < N_\phi/2$, **do**

2 $\quad$ $T_k \leftarrow \frac{1}{2}T_{\ell,\boldsymbol{r}_0}(\phi_i, \frac{2\pi k}{2\ell+1})$, $k = 0, \ldots, 2\ell$;

3 $\quad$ $\widetilde{T}_m \leftarrow \mathcal{F}(m, T)$;

4 $\quad$ $\widetilde{s}_m \leftarrow \mathcal{F}(|m| \leq N_\theta/2 - 1 + \ell; |\sin(\theta)|)$;

5 $\quad$ $\widetilde{T}_n^{\text{s},L} \leftarrow \widetilde{s} \otimes \widetilde{T}$ convolution of Fourier series;

6 $\quad$ $\widetilde{T}_n^{\text{s},L} \leftarrow$ truncate to frequencies $|n| \leq N_\theta/2 - 1$;

7 $\quad$ $T^{\text{s},L}(\theta_j, \phi_i) \leftarrow$ inverse transform of $\widetilde{T}_n^{\text{s},L}$;

This algorithm yields the low-pass modified transfer function at (ϕ_i, θ_j), $0 \leq i < N_\phi/2$, $0 \leq j < N_\theta$ which can be unwrapped to the remaining points by using the spherical symmetry $(\phi_i, \theta_j) = (\phi_{N_\phi/2+i}, \theta_{N_\theta-j})$. Note that this calculation can also be performed in the real space. It is equivalent to making a Fourier interpolation of T_k from $2\ell+1$ points to $N_\theta+2\ell-1$ points, multiplying by a low-pass $|\sin(\theta)|$, and performing a Fourier anterpolation back to N_θ points, as shown in Figure 2.

Because sampling the transfer function at a single point is an $\mathcal{O}(\ell)$ operation, the algorithm as presented is $\mathcal{O}(\ell^3)$. The computation of the transfer function at all sample points can be accelerated to $\mathcal{O}(\ell^2)$ as in [25] by taking advantage of its symmetry about the $\hat{\boldsymbol{r}}_0$ axis and using interpolation algorithms, but at the cost of introducing additional error.

3.2 *Choice of Quadrature*

The quadrature parameters can be constructively computed by determining the maximum error they incur. The error in computing the integral with a finite uniform quadrature is

$$|\varepsilon_I| = \left| \int_0^{2\pi} \int_0^{2\pi} e^{\imath\kappa\hat{\boldsymbol{s}}\cdot\boldsymbol{r}} T_{\ell,\boldsymbol{r}_0}^{\text{s}}(\hat{\boldsymbol{s}}) \, d\phi d\theta - \sum_{n=1}^{N_\theta} \sum_{m=1}^{N_\phi} \omega_{n,m} e^{\imath\kappa\hat{\boldsymbol{s}}_{n,m}\cdot\boldsymbol{r}} T_{\ell,\boldsymbol{r}_0}^{\text{s},L}(\hat{\boldsymbol{s}}_{n,m}) \right|$$

where $\hat{\boldsymbol{s}}_{n,m} = [\cos(\phi_m)\sin(\theta_n), \sin(\phi_m)\sin(\theta_n), \cos(\theta_n)]$ and $T_{\ell,\boldsymbol{r}_0}^{\text{s},L}(\hat{\boldsymbol{s}}_{n,m})$ is the low-pass modified transfer function described in Section 3.1. This can be further expanded as:

$$= \left| \int_0^{2\pi} \int_0^{2\pi} \left[E_{\kappa\boldsymbol{r}}^L(\hat{\boldsymbol{s}}_{n,m}) + E_{\kappa\boldsymbol{r}}^H(\hat{\boldsymbol{s}}_{n,m}) \right] \left[T_{\ell,\boldsymbol{r}_0}^{\text{s},L}(\hat{\boldsymbol{s}}) + T_{\ell,\boldsymbol{r}_0}^{\text{s},H}(\hat{\boldsymbol{s}}) \right] \, d\phi d\theta \right.$$

$$\left. - \sum_{n=1}^{N_\theta} \sum_{m=1}^{N_\phi} \omega_{n,m} \left[E_{\kappa\boldsymbol{r}}^L(\hat{\boldsymbol{s}}_{n,m}) + E_{\kappa\boldsymbol{r}}^H(\hat{\boldsymbol{s}}_{n,m}) \right] T_{\ell,\boldsymbol{r}_0}^{\text{s},L}(\hat{\boldsymbol{s}}_{n,m}) \right|$$
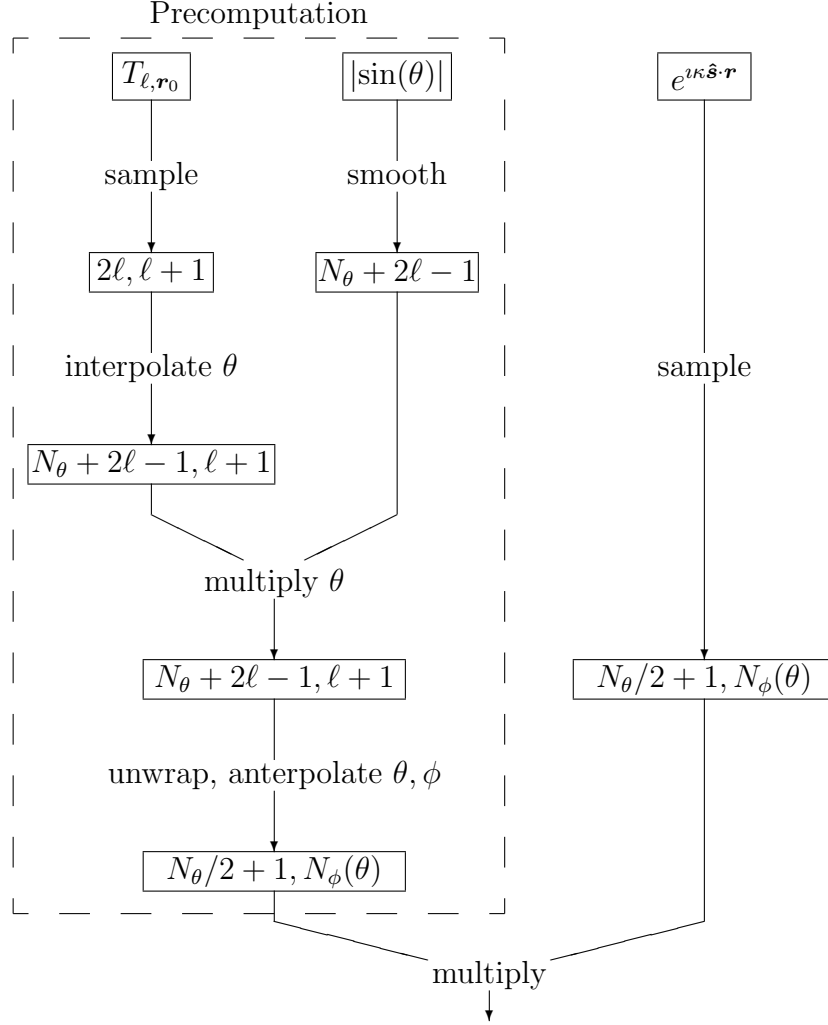
14

Fig. 2. Procedure for precomputing the low-pass modified transfer function and its application to an outgoing field. The boxed numbers (e.g., $\boxed{2\ell, \ell+1}$) give the numbers of quadrature points for θ and ϕ (N_θ and N_ϕ) at each stage.

where $E_{\kappa r}^L$ consists of the low frequencies of $e^{\iota\kappa\hat{\boldsymbol{s}}\cdot\boldsymbol{r}}$ which are resolved by the quadrature and $E_{\kappa r}^H$ consists of the high frequencies that are aliased by the quadrature. Similary, $T_{\ell,\boldsymbol{r}_0}^{\mathrm{s},H}$ are the high frequencies of $T_{\ell,\boldsymbol{r}_0}^{\mathrm{s}}$ which we removed in section 3.1. Since $E_{\kappa r}^L T_{\ell,\boldsymbol{r}_0}^{\mathrm{s},L}$ is integrated exactly by a uniform quadrature, and $E_{\kappa r}^H T_{\ell,\boldsymbol{r}_0}^{\mathrm{s},L}$ does not contribute to the integral by Parseval's theorem, we get

$$= \left| \int_0^{2\pi} \int_0^{2\pi} E_{\kappa r}^H(\hat{\boldsymbol{s}}) T_{\ell,\boldsymbol{r}_0}^{\mathrm{s},H}(\hat{\boldsymbol{s}}) \, d\phi d\theta - \sum_{n=1}^{N_\theta} \sum_{m=1}^{N_\phi} \omega_{n,m} E_{\kappa r}^H(\hat{\boldsymbol{s}}_{n,m}) T_{\ell,\boldsymbol{r}_0}^{\mathrm{s},L}(\hat{\boldsymbol{s}}_{n,m}) \right|$$

$$= \left| \int_0^{2\pi} \int_0^{2\pi} E_{\kappa r}^H(\hat{\boldsymbol{s}}) T_{\ell,\boldsymbol{r}_0}^{\mathrm{s},H}(\hat{\boldsymbol{s}}) - E_{\kappa r}^{AH}(\hat{\boldsymbol{s}}) T_{\ell,\boldsymbol{r}_0}^{\mathrm{s},L}(\hat{\boldsymbol{s}}) \, d\phi d\theta \right|$$

15

where we have denoted the aliased high frequencies of $e^{\iota \kappa \hat{\boldsymbol{s}} \cdot \boldsymbol{r}}$ as $E_{\kappa \boldsymbol{r}}^{AH}(\hat{\boldsymbol{s}})$,

$$= \left| \int_0^{2\pi} \int_0^{2\pi} \left(E_{\kappa \boldsymbol{r}}^H(\hat{\boldsymbol{s}}) - E_{\kappa \boldsymbol{r}}^{AH}(\hat{\boldsymbol{s}}) \right) T_{\ell, \boldsymbol{r}_0}^{\mathrm{s}}(\hat{\boldsymbol{s}}) \, d\phi d\theta \right|$$

In Fourier space, this becomes

$$= 4\pi^2 \left| \sum_{n=-\ell}^{\ell} \sum_{m=-\infty}^{\infty} \left(\widetilde{E}_{\kappa \boldsymbol{r}}^H(n, m) - \widetilde{E}_{\kappa \boldsymbol{r}}^{AH}(n, m) \right) \widetilde{T}_{\ell, \boldsymbol{r}_0}^{\mathrm{s}}(-n, -m) \right|$$

Choosing N_θ using this expression for the error with a representative $\boldsymbol{r}$ and $\boldsymbol{r}_0$ leads to unpredictable cancelation effects and may result in a poor choice. Instead, we apply the triangle inequality,

$$\leq 4\pi^2 \sum_{n=-\ell}^{\ell} \sum_{m=-\infty}^{\infty} \left| \widetilde{E}_{\kappa \boldsymbol{r}}^H(n, m) - \widetilde{E}_{\kappa \boldsymbol{r}}^{AH}(n, m) \right| \left| \widetilde{T}_{\ell, \boldsymbol{r}_0}^{\mathrm{s}}(-n, -m) \right|$$

This remains an accurate upper bound due to the fast decay of $\widetilde{E}$ for sufficiently large values of N_θ and N_ϕ. See Figure 3.
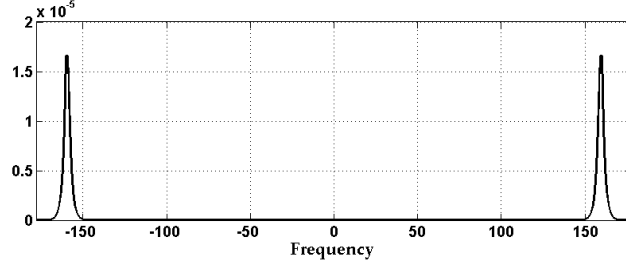
Fig. 3. The value of $\left| \widetilde{E}_{\kappa \boldsymbol{r}}^H(0, m) - \widetilde{E}_{\kappa \boldsymbol{r}}^{AH}(0, m) \right|$ for $\kappa \left| \boldsymbol{r} \right| = 0.8\sqrt{3} \cdot 100$ and $N_\phi = 318$.

3.2.1 Choosing N_θ

The worst case for ε_I in terms of N_θ occurs when $\boldsymbol{r}$ and $\boldsymbol{r}_0$ are aligned with the z-axis. This causes all spectral information to be contained in the θ-direction and makes ε_I a function of N_θ only. It leads to

$$|\varepsilon_I| \leq 4\pi^2 \sum_{m=-\infty}^{\infty} \left| \widetilde{E}_{\kappa \boldsymbol{r}}^H(0, m) - \widetilde{E}_{\kappa \boldsymbol{r}}^{AH}(0, m) \right| \left| \widetilde{T}_{\ell, \boldsymbol{r}_0}^{\mathrm{s}}(0, -m) \right|$$

Using the plane wave spectrum (5) with $\hat{\boldsymbol{r}} = \hat{\boldsymbol{z}}$, this is approximately simplified to

$$|\varepsilon_I| \leq 4\pi^2 \sum_{m=-\infty}^{\infty} \left| J_{M(N_\theta, m)}(\kappa \left| \boldsymbol{r} \right|) \right| \left| \widetilde{T}_{\ell, \boldsymbol{r}_0}^{\mathrm{s}}(0, m) \right| \tag{12}$$

16

where

$$M(N_\theta, m) = \begin{cases} N_\theta - |m| & |m| \leq N_\theta/2 - 1 \\ |m| & |m| > N_\theta/2 - 1 \end{cases} \qquad (13)$$

This is an approximation because $\widetilde{E}_{\kappa\boldsymbol{r}}^{AH}$ would in principle contribute an infinite sum to equation (12) rather than the single term used. However, given the exponential decay of the Jacobi-Anger series, the difference is negligible. Equation (12) can be used to search for a value N_θ via the algorithm sketched below:

1 Choose N_θ^n sufficiently larger than $2\ell + 1$;
2 $T_k \leftarrow T_{\ell,|\boldsymbol{r}_0|\hat{\boldsymbol{z}}}^{\mathrm{s},L}(0, \frac{2\pi k}{N_\theta^n})$, $k = 0, \ldots, N_\theta^n - 1$;
3 $\widetilde{T}_m \leftarrow |\mathcal{F}(m; T_k)|$;
4 $\widetilde{E}_m \leftarrow |J_m(\kappa |\boldsymbol{r}|)|$;
5 **for** N_θ *from* 2ℓ *to* N_θ^n *by* 2 **do**
6 $\quad \widetilde{E}_m^* \leftarrow \widetilde{E}_{M(N_\theta, m)}$;
7 $\quad$ **if** $\widetilde{E}^* \cdot \widetilde{T} < \varepsilon/4\pi^2$ **then**
8 $\quad\quad$ return N_θ

Since N_θ^n is typically only a small constant larger than $2\ell + 1$, the algorithm as presented is dominated by the computation of the $\mathcal{O}(\ell)$ modified transfer function values and requires $\mathcal{O}(\ell^2)$ operations. Important optimizations include using more advanced searching methods (such as bisection), applying the symmetries $\widetilde{E}_m^* = \widetilde{E}_{-m}^*$ and $\widetilde{T}_m = \widetilde{T}_{-m}$, and taking advantage of the very fast decay of J_n to neglect very small terms in the dot product.

3.2.2 Choosing N_ϕ

After determining an appropriate N_θ, letting N_ϕ be a function of θ allows reducing the number of quadrature points without affecting the error. The worst case for the integration error occurs when $\boldsymbol{r}$ and $\boldsymbol{r}_0$ are in the xy-plane. Without loss of generality, suppose $\hat{\boldsymbol{r}} = \hat{\boldsymbol{x}}$. Consider a constant $\theta = \theta_j$ and note that the plane wave can be expressed as

$$e^{\iota\kappa\hat{\boldsymbol{s}}\cdot\boldsymbol{r}} = \sum_{n=-\infty}^{\infty} i^n J_n(\kappa |\boldsymbol{r}| \sin(\theta_j)) e^{in\phi}$$

Since $J_n(\kappa |\boldsymbol{r}| \sin(\theta_j))$ is exponentially small when $n \gtrsim \kappa |\boldsymbol{r}| \sin(\theta_j)$, the series can be truncated at $N_\phi(\theta_j) \sim \kappa |\boldsymbol{r}| \sin(\theta_j)$ without incurring any appreciable error. Estimates of $N_\phi(\theta_j)$ can be developed by determining when $J_n(\kappa |\boldsymbol{r}| \sin(\theta_j))$ becomes exponentially small, as in the computation of the excess bandwidth formula (EBF) in [6]. However, we find that the EBF generated quadrature typically overestimates the sampling rate.

17

To accurately compute $N_\phi(\theta_j)$ the same procedure as in Sec. 3.2.1 is applied but with r and r_0 in the xy-plane. This represents the worst case for the integration error as a function of N_ϕ. For a given θ_j, we search for a $N_\phi(\theta_j)$ such that

$$|\varepsilon_I| \leq 4\pi^2 \sum_{n=-\ell}^{\ell} \left| J_{M(N_\phi(\theta_j),n)}(\kappa\,|r|\sin(\theta_j)) \right| \left| \widetilde{T}^{s}_{\ell,r_0}(n;\theta_j) \right| \qquad (14)$$

is bounded by a prescribed error. This is accomplished via the following sketched algorithm.

1 Set N_ϕ at the poles: $N_\phi(\theta_0) = N_\phi(\theta_{N_\theta/2}) = 1$;
2 Choose N_ϕ^n sufficiently larger than $2\ell + 1$;
3 for $\theta_j,\ j = 1, \ldots, N_\theta/2 - 1$ **do**
4 $\quad$ $T_k \leftarrow T^{s,L}_{\ell,|r_0|\hat{x}}(\frac{2\pi k}{2\ell+1}, \theta_j),\ k = 0, \ldots, 2\ell$;
5 $\quad$ $\widetilde{T}_m \leftarrow |\mathcal{F}(m; T_k)|$;
6 $\quad$ $\widetilde{E}_m \leftarrow |J_m(\kappa\,|r|\sin(\theta_j))|$;
7 $\quad$ **for** $N_\phi(\theta_j)$ *from* 2 *to* N_ϕ^n *by* 2 **do**
8 $\quad\quad$ $\widetilde{E}^*_m \leftarrow \widetilde{E}_{M(N_\phi(\theta_j),m)}$;
9 $\quad\quad$ **if** $\widetilde{E}^* \cdot \widetilde{T} < \varepsilon/4\pi^2$ **then**
10 $\quad\quad\quad$ Save $N_\phi(\theta_j)$

Since N_ϕ^n is only a small constant larger than $2\ell+1$, the algorithm as presented is dominated by the computation of the modified transfer function and requires $\mathcal{O}(\ell^3)$ operations. Optimizations similar to those presented in Sec. 3.2.1 can be applied. Using the EBF as an initial guess in the search for $N_\phi(\theta_j)$ further improves the searching speed. Additionally, only half of the $N_\phi(\theta_j)$'s may be computed due to symmetry about the $z = 0$ plane.

We finally note that letting N_ϕ be a function of θ_j requires an additional step in the computation of the modified transfer function. Section 3.1 computed the transfer function on a $N_\theta/2 + 1 \times N_\phi$ grid. With $N_\phi \to N_\phi(\theta_j)$, the data computed for each θ_j must be Fourier anterpolated from length N_ϕ to length $N_\phi(\theta_j)$.

3.2.3 Choosing $|r|$ and $|r_0|$

The previous algorithms require representative values of $|r|$ and $|r_0|$ for each level of the tree. The worst-case transfer vectors, r_0, are those with smallest length. If a_l is the box size at level l, then $|r_0| = 2a_l$ is the smallest transfer vector length in the common one buffer box case.

The worst case value of $|r|$ is the largest. For a box of size a_l, $|r| \leq a_l\sqrt{3}$. However, using $|r| = a_l\sqrt{3}$ in the previous methods is often too conservative.

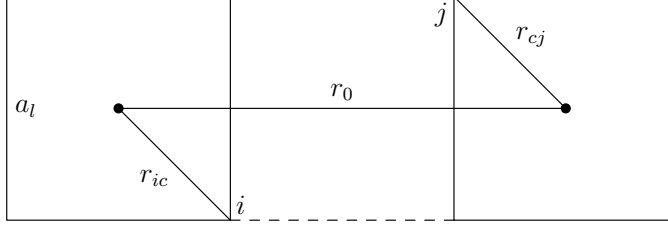Fig. 4. The worst case $\boldsymbol{r}$ and $\boldsymbol{r}_0$, projected from the 3D box. Here, $|\boldsymbol{r}_0| = 2a_l$ and i and j are on the opposite corners of the box so that $|\boldsymbol{r}| = |\boldsymbol{r}_{ic}| + |\boldsymbol{r}_{cj}| = a_l\sqrt{3}$.

This case only occurs when two points are located in the exact corners of the boxes – a rare case indeed. See Figure 4. Instead, we let $|\boldsymbol{r}| = \alpha a_l\sqrt{3}$ for some $\alpha \in [0, 1]$. A high α guarantees an upper bound on the error generated by the quadrature, but the points which actually generate this error become increasingly rare. A lower value of α will yield a smaller quadrature, but more points may fall outside the radius $|\boldsymbol{r}|$ for which the upper bound on the error is guaranteed.

3.2.4 Number of Quadrature Points

Recall from section 2.2 that the typical approach in the FMM is to use $N + 1$ uniform points in the ϕ direction and $\frac{N+1}{2}$ Gauss-Legendre points in the θ direction so that all Y_n^m, $-n \leq m \leq n$, $0 \leq n \leq N$ are integrated exactly. In [8], Chew et al. takes $\frac{N+1}{2} = \ell + 1$, which is an approximate choice based on the rapid decay of the coefficients in the spherical harmonics expansion of a plane wave. This results in approximately

$$M_g = 2(\ell + 1)^2 \approx 2\ell^2$$

quadrature points.

For a given Gegenbauer series truncation ℓ, the total number of quadrature points required in the Fourier based FMM is approximately

$$\begin{aligned} M_f &\approx \frac{N_\theta}{2} \frac{1}{\pi} \int_0^\pi N_\phi(\theta)\, d\theta \\ &\approx (\ell + C_1)\frac{1}{\pi} \int_0^\pi (2\ell + C_2(\theta)) \sin(\theta)\, d\theta \end{aligned}$$

where $C_1, C_2 \geq 1$ are small integers dependent on ℓ, numerically computed from the methods in Sec. 3.2.1, 3.2.2. Keeping only the leading term in ℓ:

$$M_f \approx \frac{4}{\pi}\ell^2 \approx 1.3\,\ell^2$$

Thus, the method presented in this paper uses approximately 0.64 times the number of quadrature points in the standard FMM. However, it is possible

19

that the same N_ϕ optimization can be applied to the standard FMM for the same reasons it was applied in section 3.2.2 to reduce the standard quadrature to a comparable size.

3.3 Interpolation and Anterpolation

Most importantly, the Fourier based FMM directly uses FFTs in the interpolation and anterpolation steps. This makes the time critical upward pass and downward pass especially fast and easy to implement while retaining the exactness of global methods.

Characterize a quadrature by an array of length $N_\theta/2 + 1$,

$$Q = [1, N_\phi(\theta_1), \dots, N_\phi(\theta_{N_\theta/2-1}), 1]$$

noting that $N_\phi(\theta_j) = N_\phi(\theta_{N_\theta/2+j})$ and $N_\phi(\theta_j) = N_\phi(\theta_{N_\theta/2-j})$. The data $F(\phi_i, \theta_j)$ sampled on a quadrature Q is transformed to a another quadrature Q' by performing a sequence of Fourier interpolations and anterpolations. Let

$$\mathcal{N}_\phi = \max\left[\max_{0 \le j \le N_\theta/2} N_\phi(\theta_j),\ \max_{0 \le j \le N'_\theta/2} N'_\phi(\theta_j) \right]$$

Then, the following steps, as illustrated in Figure 5, perform an exact interpolation/anterpolation using only FFTs.

(1) For each θ_j, $0 \le j \le N_\theta/2$, Fourier interpolate the data $[F(\phi_{i=0,\dots,N_\phi(\theta_j)-1}, \theta_j)]$ from length $N_\phi(\theta_j)$ to $\mathcal{N}_\phi$.

(2) For each ϕ_i, $0 \le i < \mathcal{N}_\phi/2$, wrap the data to construct the periodic sequence from the rest of the line $[F(\phi_i, \theta_{j=0,\dots,N_\theta/2}), F(\phi_{i+\mathcal{N}_\phi/2}, \theta_{j=N_\theta/2-1,\dots,1})]$.

(3) For each ϕ_i, $0 \le i < \mathcal{N}_\phi/2$, Fourier interpolate the data $[F(\phi_i, \theta_{j=0,\dots,N_\theta-1})]$ from length N_θ to N'_θ.

(4) For each ϕ_i, $0 \le i < \mathcal{N}_\phi/2$, unwrap the data $[F(\phi_i, \theta_{j=0,\dots,N'_\theta-1})]$ to construct the sequences $[F(\phi_i, \theta_{j=0,\dots,N'_\theta/2})]$ and $[F(\phi_{i+\mathcal{N}_\phi/2}, \theta_{j=0,\dots,N'_\theta/2})]$.

(5) For each θ_j, $0 \le j \le N'_\theta/2$, Fourier anterpolate the data $[F(\phi_{i=0,\dots,\mathcal{N}_\phi(\theta_j)-1}, \theta_j)]$ from length $\mathcal{N}_\phi$ to $N'_\phi(\theta_j)$.

3.4 Numerical Results

3.4.1 Error

A direct computation was used to compute the optimal Gegenbauer truncation ℓ and the methods described in section 3.2 were used to construct a quadrature for use in computing the integral (10). For a given box size a, the quadrature
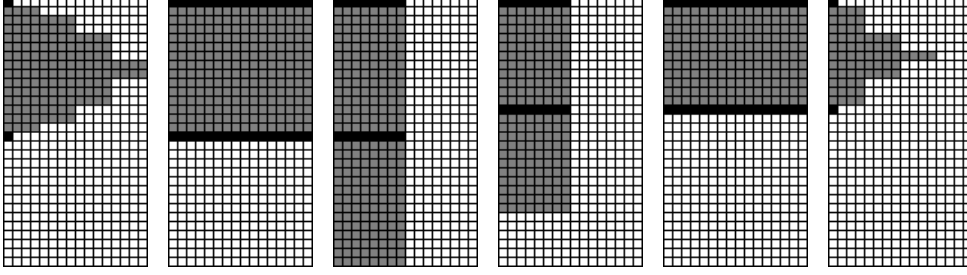
20

Fig. 5. The data profile at each step in an anterpolation from a large quadrature Q with $N_\theta = 30$ to a smaller quadrature Q' with $N'_\theta = 24$. The angle ϕ is in the x direction while the angle θ is in the y direction. The data corresponding to a pole has been darkened for clarity.

and truncation are constructed with $|\boldsymbol{r}| = 0.8a\sqrt{3}$, $|\boldsymbol{r}_0| = 2a$, and target error eps. The total measured error, ε, is defined as

$$\varepsilon = \frac{e^{\imath\kappa|\boldsymbol{r}+\boldsymbol{r}_0|}}{|\boldsymbol{r}+\boldsymbol{r}_0|} - \sum_{m=1}^{N_\theta} \sum_{n=1}^{N_\phi(\theta_m)} \omega_{n,m} \, e^{\imath\kappa\hat{\boldsymbol{s}}_{n,m}\cdot\boldsymbol{r}} \, T_{\ell,\boldsymbol{r}_0}^{\mathrm{s},L}(\hat{\boldsymbol{s}}_{n,m})$$

The total Gegenbauer truncation error, ε_G, is

$$\varepsilon_G = \frac{e^{\imath\kappa|\boldsymbol{r}+\boldsymbol{r}_0|}}{|\boldsymbol{r}+\boldsymbol{r}_0|} - \imath\kappa \sum_{n=0}^{\ell} (-1)^n (2n+1) h_n^{(1)}(\kappa\,|\boldsymbol{r}_0|) j_n(\kappa\,|\boldsymbol{r}|) P_n(\hat{\boldsymbol{r}}\cdot\hat{\boldsymbol{r}}_0)$$

The total integration error ε_I is

$$\varepsilon_I = \varepsilon - \varepsilon_G$$

In Figure 6, the plotted errors represent the maximum found over many directions $\hat{\boldsymbol{r}}$ and magnitudes $|\boldsymbol{r}| \leq 0.8s\sqrt{3}$. As is evident, as the box size increases, the target error eps is accurately achieved. The increase in error for small box sizes corresponds to the low frequency breakdown when the transfer function has very large amplitude and roundoff errors become dominant. In this regime the quadrature target error bound is also relaxed to improve efficiency - it is inefficient to have a large quadrature that provides a small integration error when the transfer function cannot provide comparable accuracy.

On the same plot we show ε_G^{EBF}, the Gegenbauer series error resulting from choosing the truncation with the EBF from section 2.1. Clearly, the EBF is overestimating ℓ, which causes the Gegenbauer error to fall far below the target error and will force the quadrature to be larger and less efficient.

The bottom plot shows the ratio of the number of points in the quadrature presented in this paper to the number of quadrature points that would be used in a typical spherical harmonics based FMM. Each of these quadratures were computed for the same Gegenbauer truncation ℓ chosen by the direct calculation. The procedures presented in this paper result in a quadrature

21

which is substantially smaller than what would typically be used. Notably, the analysis in Section 3.2.4 is supported.

Together, these results demonstrate that by choosing ℓ and the quadrature as presented in this paper, the error is better controlled and the quadrature size at each level in the tree is reduced. Improved error control means that we can provide a sharp bound of the total final error of the method and optimize the running time of the method for that prescribed error. A reduction in the quadrature size improves memory usage and suggests an improved running time over similar algorithms.

3.4.2 Speed

As discussed in Section 3.3, the Fourier based FMM uses only FFTs in the upward pass and downward pass to perform the interpolations and anterpolations. FFTs make these steps easier to implement and very fast.

Figure 7 shows the recorded running times of the Fourier based FMM and the direct matrix-vector product on a Intel Core 2 Quad CPU Q9450 2.66GHz with 4GB of RAM. For $N = 8.2 \cdot 10^6$ the points are uniformly distributed in a cube with side length 80λ and the wave number κ is scaled with $N^{1/3}$. This provides a nearly constant density of points per wavelength as N is varied. As expected, by choosing the correct number of levels the running time is asymptotically $\mathcal{O}(N \log N)$ as N is increased with a constant number of points per wavelength. Note that the cross-over point is less than $N = 4,000$. The code used to produce these results was not optimized for memory usage, preventing results for $N \gtrsim 10^6$ when $L = 7$.

4   Conclusion

We have proposed using the Fourier basis $e^{\imath p\phi}e^{\imath q\theta}$ in the spherical variables ϕ and θ to represent the far field approximation in the FMM. By approximating the Helmholtz kernel with

$$\frac{e^{\imath\kappa|\boldsymbol{r}+\boldsymbol{r}_0|}}{|\boldsymbol{r}+\boldsymbol{r}_0|} \approx \int_0^{2\pi} \int_0^{2\pi} e^{\imath\kappa\hat{\boldsymbol{s}}\cdot\boldsymbol{r}}\, T_{\ell,\boldsymbol{r}_0}^{\mathrm{s}}(\hat{\boldsymbol{s}})\, d\phi d\theta, \qquad T_{\ell,\boldsymbol{r}_0}^{\mathrm{s}}(\hat{\boldsymbol{s}}) = \frac{1}{2}T_{\ell,\boldsymbol{r}_0}(\hat{\boldsymbol{s}})\,|\sin(\theta)|\,,$$

and using a uniform quadrature we can take advantage of very fast, exact, and well-known FFT interpolation/anterpolation methods. By exploiting symmetries and a scheme to reduce the number of points in the ϕ direction, the total number of uniform quadrature points required is smaller than the number of Gauss-Legendre quadrature points typically used with spherical harmonics.
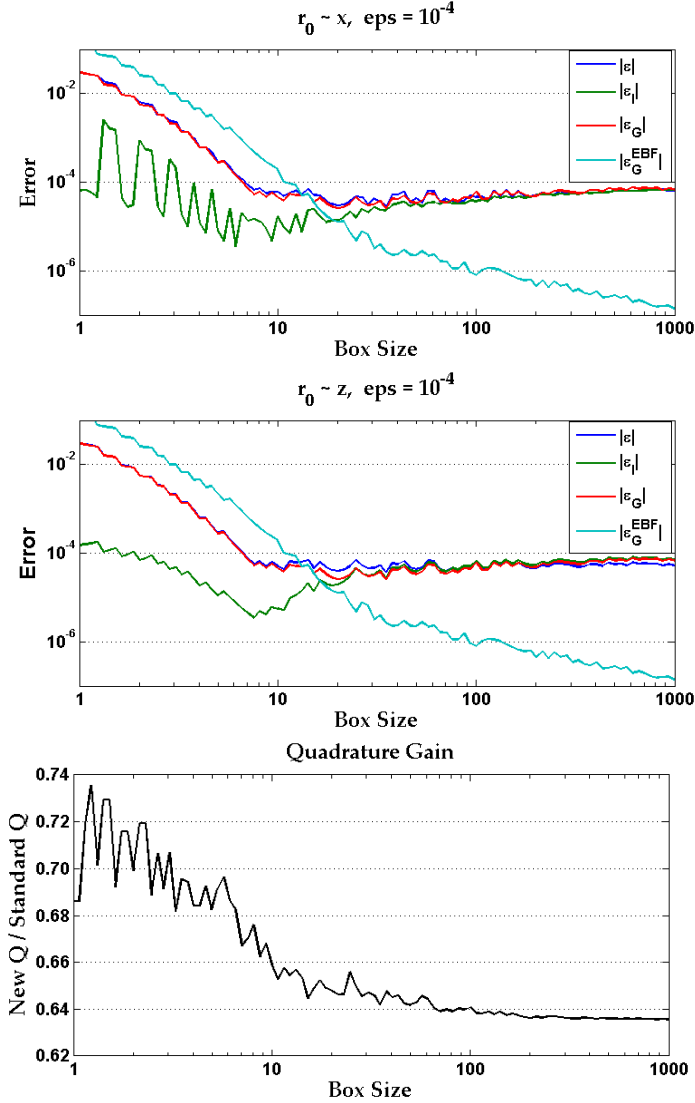
22

Fig. 6. Top plots: error of the FMM integral using a direct computation of ℓ as described in section 2.1 and the choice of quadrature described in section 3.2. All the errors fall very close to the target error of 10^{-4}. The standard EBF overestimates ℓ and will result in a suboptimal quadrature. Bottom plot: ratio of the number of quadrature points required in the Fourier based FMM with what would be used in a typical spherical harmonics based FMM for the same ℓ. The curve asymptotes close to $2/\pi \approx 0.64$ as expected.

This is realized by correcting the transfer function $T^s_{\ell, r_0}(\hat{s})$ during the precomputation phase to remove high frequency terms which do not significantly contribute to the final integration.

The Fourier based FMM approach has a number of advantages. Since the interpolation and anterpolation algorithms are exact, the error analysis is simplified; we establish a sharp upper bound for the error. The key parameters are the Gegenbauer truncation parameter ℓ and the quadrature size, in particular
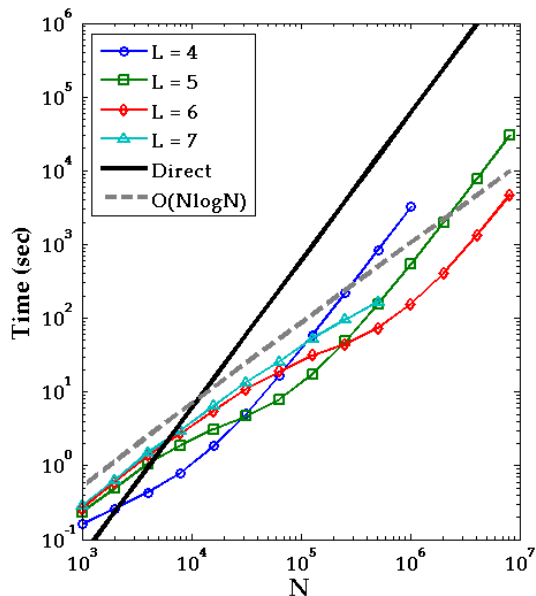
23

Fig. 7. Average running times of the Fourier based FMM for constant number of points per wavelength.

the bandwidth in the θ-direction. The truncation error ε_ℓ has been extensively studied by other authors and is well understood. The integration error ε_I accounts for the low-pass approximation of the modified transfer function and the aliasing of the plane waves. This error can be accounted for a priori during the precomputation stage. Numerical tests have confirmed that this error analysis is quite sharp. Constructive algorithms to find nearly optimal parameters were proposed.

Since efficient FFT algorithms are available in virtually every computing environment, the time-critical stages of the algorithm are much easier to implement. Furthermore, although the asymptotic arithmetic complexity is unchanged, the smaller constant in the FFT's $\mathcal{O}(N \log N)$ complexity yields improved running times.

Acknowledgements

24

References

[1] V. Rokhlin, Rapid solution of integral equations of scattering theory in two dimensions, J. Comput. Phys. 86 (2) (1990) 414–439.

[2] R. Coifman, V. Rokhlin, S. Wandzura, The fast multipole method for the wave equation: A pedestrian prescription, Antennas and Propag. Magazine, IEEE 35 (3) (1993) 7–12.

[3] V. Rokhlin, Diagonal forms of translation operators for the Helmholtz equation in three dimensions, Applied and Computational Harmonic Analysis 1 (1993) 82–93.

[4] N. Engheta, W. D. Murphy, V. Rokhlin, M. S. Vassiliou, The fast multipole method (FMM) for electromagnetic scattering problems, IEEE Trans. Antennas Propag. 40 (1992) 634.

[5] J. Rahola, Diagonal forms of the translation operators in the fast multipole algorithm for scattering problems, BIT 36 (1996) 333–358.

[6] W. Chew, E. Michielssen, J. M. Song, J. M. Jin (Eds.), Fast and Efficient Algorithms in Computational Electromagnetics, Artech House, Inc., Norwood, MA, USA, 2001.

[7] E. Darve, The fast multipole method: Numerical implementation, J. Comput. Phys. 160 (1) (2000) 195–240.

[8] S. Koc, J. M. Song, W. C. Chew, Error analysis for the numerical evaluation of the diagonal forms of the scalar spherical addition theorem, SIAM J. Numer. Anal. 36 (3) (1999) 906–921.

[9] J. R. Driscoll, D. Healy, Computing Fourier transforms and convolutions on the 2-sphere, Adv. Appl. Math. 15 (2) (1994) 202–250.

[10] D. Healy, D. Rockmore, P. Kostelec, S. Moore, FFTs for the 2-sphere - improvements and variations, J. Fourier Analysis and Applications 9 (4) (2003) 341–385.

[11] R. Suda, M. Takami, A fast spherical harmonics transform algorithm, Math. of Comp. 71 (238) (2001) 709–715.

[12] J. Sarvas, Performing interpolation and anterpolation entirely by fast Fourier transform in the 3-D multilevel fast multipole algorithm, SIAM J. Numer. Anal. 41 (6) (2003) 2180–2196.

[13] Q. Carayol, F. Collino, Error estimates in the fast multipole method for scattering problems. Part 2: Truncation of the Gegenbauer series, ESAIM: M2NA 39 (1) (2004) 183–221.

[14] E. Darve, The fast multipole method I: Error analysis and asymptotic complexity, SIAM J. Numer. Anal. 38 (1) (2000) 98–128.

25

[15] M. L. Hastriter, S. Ohnuki, W. C. Chew, Error control of the translation operator in 3D MLFMA, Microwave and Optical Technology Letters 37 (3) (2003) 184–188.

[16] Q. Carayol, F. Collino, Error estimates in the fast multipole method for scattering problems. Part 1: Truncation of the Jacobi-Anger series, ESAIM: M2NA 38 (2) (2004) 371–394.

[17] A. D. McLaren, Optimal numerical integration on a sphere, Mathematics of Computation 17 (84) (1963) 361–383.

[18] V. Rokhlin, M. Tygert, Fast algorithms for spherical harmonic expansions, SIAM J. Scientific Computing 27 (6) (2006) 1903–1928.

[19] I. Chowdhury, V. Jandhyala, Integration and interpolation based on fast spherical transforms for the multilevel fast multipole method, Microwave and Optical Technology Letters 48 (10) (2006) 1961–1964.

[20] R. Jakob-Chien, B. K. Alpert, A fast spherical filter with uniform resolution, J. Comput. Phys. 136 (2) (1997) 580–584.

[21] A. Dutt, M. Gu, V. Rokhlin, Fast algorithms for polynomial interpolation, integration, and differentiation, SIAM J. Numer. Anal. 33 (5) (1996) 1689–1711.

[22] J. Knab, Interpolation of bandlimited functions using the approximate prolate series (corresp.), Information Theory, IEEE 25 (6) (1979) 717–720.

[23] E. Darve, P. Havé, A fast multipole method for Maxwell equations stable at all frequencies, Phil. Trans. R. Soc. Lond. A 362 (1816) (2004) 603–628.

[24] E. Darve, P. Havé, Efficient fast multipole method for low-frequency scattering, J. Comput. Phys. 197 (1) (2004) 341–363.

[25] O. Ergul, L. Gurel, Optimal interpolation of translation operator in multilevel fast multipole algorithm, IEEE Trans. Antennas Propag. 54 (12) (2006) 3822–3826.

A   Asymptotic Complexity

In order to resolve a sufficient number of spherical harmonics, the number of points in the ϕ and θ directions must be $\mathcal{O}(\ell) = \mathcal{O}(\kappa a)$, where a is the side length of the box. Therefore, the total number of quadrature points is $\mathcal{O}(\ell^2) = \mathcal{O}((\kappa a)^2)$. If a_0 is the side length of a box at the root of the octree (level $l = 0$), then the side length of a box at level l is $a_l = 2^{-l}a_0$ and has $\mathcal{O}((\kappa 2^{-l}a_0)^2)$ quadrature points.

Suppose the octree has levels $l = 0, \ldots, L$ with boxes of side length $a_l = 2^{-l}a_0$. A box at level l will contain a field approximation defined over a quadrature

26

of size $\mathcal{O}((\kappa 2^{-l} a_0)^2)$. We now determine the number of operations required for each step.

(1) **Initialization/Collection:** This step requires sampling $e^{\imath \kappa \hat{\boldsymbol{s}} \cdot \boldsymbol{r}}$ for each of the N source points at the leaves of the tree. Thus, this step is

$$\mathcal{O}(N(\kappa 2^{-L} a_0)^2) = \mathcal{O}(N 2^{-2L}(\kappa a_0)^2)$$

(2) **Upward Pass:** This step requires aggregating and interpolating each outgoing field upward through the tree. The type of interpolation algorithm is key to the running time of this step.

At level l in the tree, the number of interpolations that must be performed is equal to the number of boxes at that level in the tree. The number of boxes depends on the distribution of source points. If the source points are uniformly distributed over a volumetric domain then the asymptotic number of boxes at level l in the octree is $\mathcal{O}(8^l)$. However, if the source points are uniformly distributed over the surface of an object then the asymptotic number of boxes at level l is $\mathcal{O}(4^l)$.

- **Direct method:** Each direct interpolation requires $\mathcal{O}(K_L K_{L-1}) = \mathcal{O}((\kappa 2^{-l} a_0)^2 (\kappa 2^{-(l-1)} a_0)^2)$ operations. Thus, the direct method has complexity

$$\text{Volume:} \quad \sum_{l=3}^{L} \mathcal{O}(8^l (\kappa 2^{-l} a_0)^2 (\kappa 2^{-(l-1)} a_0)^2) = \mathcal{O}((\kappa a_0)^4)$$

$$\text{Surface:} \quad \sum_{l=3}^{L} \mathcal{O}(4^l (\kappa 2^{-l} a_0)^2 (\kappa 2^{-(l-1)} a_0)^2) = \mathcal{O}((\kappa a_0)^4)$$

- **Fast global methods:** By using a fast interpolation method, the complexity for an individual interpolation is reduced to $\mathcal{O}(K_l \log(K_l))$. The upward pass complexity then becomes

$$\text{Volume:} \quad \sum_{l=3}^{L} \mathcal{O}(8^l (\kappa 2^{-l} a_0)^2 \log(\kappa 2^{-l} a_0)) = \mathcal{O}(2^L (\kappa a_0)^2 (\log(\kappa a_0) - L))$$

$$\text{Surface:} \quad \sum_{l=3}^{L} \mathcal{O}(4^l (\kappa 2^{-l} a_0)^2 \log(\kappa 2^{-l} a_0)) = \mathcal{O}(L(\kappa a_0)^2 (\log(\kappa a_0) - L))$$

- **Local methods:** Local methods use a stencil of some given size to compute the interpolated values. Although these methods introduce additional error, they have fast execution times with $\mathcal{O}(K_l)$ operations. The upward pass complexity then becomes

$$\text{Volume:} \quad \sum_{l=3}^{L} \mathcal{O}(8^l (\kappa 2^{-l} a_0)^2) = \mathcal{O}(2^L (\kappa a_0)^2)$$

$$\text{Surface:} \quad \sum_{l=3}^{L} \mathcal{O}(4^l (\kappa 2^{-l} a_0)^2) = \mathcal{O}(L(\kappa a_0)^2)$$

(3) **Transfer Pass:** Since each box has a maximum of 189 transfers and the transfer function is diagonal, the running time of this step is

$$\text{Volume:} \qquad \sum_{l=2}^{L} \mathcal{O}(8^l (\kappa 2^{-l} a_0)^2) = \mathcal{O}(2^L (\kappa a_0)^2)$$

$$\text{Surface:} \qquad \sum_{l=2}^{L} \mathcal{O}(4^l (\kappa 2^{-l} a_0)^2) = \mathcal{O}(L (\kappa a_0)^2)$$

(4) **Downward Pass:** The downward pass is the adjoint operation of the Upward Pass and has the same asymptotic complexity.

(5) **Finalization:** For each of the N field points, we integrate the spherical function at the leaf and compute the close contributions from neighboring boxes:

$$\mathcal{O}(N 2^{-2L} (\kappa a_0)^2) + \mathcal{O}(\text{close})$$

In the worst case, the close interaction is $\mathcal{O}(N^2)$ which occurs when there is an accumulation of points somewhere in the domain. In that case a different scheme is required since the expansions used in this paper are unstable at low frequency. When the field points are distributed roughly uniformly, then

$$\text{Volume:} \qquad \mathcal{O}(\text{close}) = \mathcal{O}((N/8^L)^2) = \mathcal{O}(2^{-6L} N^2)$$

$$\text{Surface:} \qquad \mathcal{O}(\text{close}) = \mathcal{O}((N/4^L)^2) = \mathcal{O}(2^{-4L} N^2)$$

The total asymptotic running time then depends on the scaling of the number of points with a_0, the scaling of L with N or a_0, and the interpolation methods that are used.

For a volume of scatters, we have $\mathcal{O}(N) = \mathcal{O}((\kappa a_0)^3)$ and let $L \sim \log(N^{1/3})$ to achieve a total algorithmic complexity of $\mathcal{O}(N \log N)$ using fast global interpolation methods and $\mathcal{O}(N)$ by using approximate, local methods. For a surface of scatters, we have $\mathcal{O}(N) = \mathcal{O}((\kappa a_0)^2)$ and let $L \sim \log(N^{1/2})$ to achieve a total algorithmic complexity of $\mathcal{O}(N \log^2 N)$ using fast global interpolation methods and $\mathcal{O}(N \log N)$ by using approximate, local methods.