FOURIER BASED FAST MULTIPOLE METHOD FOR THE HELMHOLTZ EQUATION

CRIS CECKA* AND ERIC DARVE†

Abstract. The fast multipole method (FMM) has had great success in reducing the computational complexity of solving the boundary integral form of the Helmholtz equation. We present a formulation of the Helmholtz FMM that uses Fourier basis functions rather than spherical harmonics. By modifying the transfer function in the precomputation stage of the FMM, time-critical stages of the algorithm are accelerated by causing the interpolation operators to become straightforward applications of fast Fourier transforms, retaining the diagonality of the transfer function, and providing a simplified error analysis. 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

AMS subject classifications. 31B10, 42B10, 65D05, 65R20, 65Y20, 65T40, 70F10, 78M15, 78M16

1. Introduction.

Since the development of the fast multipole method (FMM) for the wave equation in [22, 5, 23, 12, 21], 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 equation, although our results can be immediately extended to the vector case as described in [3, 7]. The application of the boundary element method to solve the integral form of the Helmholtz equation results in a dense linear system which can be solved by iterative methods such as GMRES or BCGSTAB. These methods require computing dense matrix-vector products which, using a direct implementation, are performed in $\mathcal{O}(N^2)$ floating-point operations. The FMM uses an approximation of the dense matrix to perform the product in $\mathcal{O}(N \log N)$ or $\mathcal{O}(N \log^2 N)$ operations. This approximation is constructed from close-pair interactions and far-field approximations represented by spherical integrals that are accumulated and distributed through the domain via an octree.

There are a number of difficulties in implementing the FMM, each of which must be carefully considered and optimized to achieve the improved complexity. The most significant complication in the Helmholtz FMM is that the quadrature sampling rate must increase with the size of the box in the octree, requiring interpolation and anterpolation algorithms to transform the data between spherical quadratures of different levels of the tree. Local algorithms such as Lagrange interpolation and techniques which sparsify interpolant matrices are fast, but incur significant errors [18, 7]. Spherical harmonic transforms are global interpolation schemes and are exact but require fast versions for efficiency of the FMM. Many of these fast spherical transform algorithms are only approximate, complicated to implement, and not always stable [10, 15, 26].

In this paper, we use a multipole expansion which allows the use of 2D fast Fourier transforms (FFT) in the spherical coordinate system (ϕ, θ). The main advantages

*Institute for Computational and Mathematical Engineering, Stanford University (ccecka@stanford.edu).

†Institute for Computational and Mathematical Engineering, Mechanical Engineering Department, Stanford University (darve@stanford.edu)

Notation	Description
κ	wavenumber, $2\pi/\lambda$ with wavelength λ
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, $\boldsymbol{x}\cdot\boldsymbol{y} = \lvert\boldsymbol{x}\rvert\,\lvert\boldsymbol{y}\rvert\cos(\varphi_{\boldsymbol{x},\boldsymbol{y}})$
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
$\mathcal{F}(m; f)$	mth frequency of the Fourier transform of f

TABLE 1.1
Table of notations

are two fold: i) high performance libraries are available for FFTs on practically all computer platforms, resulting in accurate, robust, and fast interpolation algorithms; ii) the resulting error analysis is simplified and leads to sharp, a priori error bounds on the FMM. One of the difficulties in using FFTs is that we are forced to use a uniform distribution of points along ϕ and θ in the spherical quadrature. Naively, this leads to a much increased quadrature size for a given accuracy compared to the original spherical harmonics-based FMM. The reason is as follows. The multipole expansion in the high frequency regime is derived from:

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

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

We propose to use a variant of the scheme by J. Sarvas in [25] whereby the integration is extended from 0 to 2π and the integrand modified:

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

We will describe in more details how an efficient scheme can be derived from this equation. The key property is that $e^{\imath\kappa\hat{\boldsymbol{s}}\cdot\boldsymbol{r}}$ 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 [25] is similar. However, rather than smoothing $T_{\ell,\boldsymbol{r}_0}(\hat{\boldsymbol{s}})\,|\sin(\theta)|$ once during the precomputation phase as we detail in this paper, Sarvas instead incorporates the $|\sin(\theta)|$ factor during the run-time phase of the FMM after the application of the transfer function. Although a detailed analysis is required for accurately assessing the relative cost of the two approaches, the technique in [25] requires approximately 1.5 times more sample points in the time-critical transfer pass of the algorithm, and requires an extra anterpolation step after applying T with about twice more sample points than used for the method in this paper. We note that the error analysis for the two methods is different, and is easier to carry out with the approach in this paper.

We derive a new a priori error analysis which incorporates both effects from truncation of the Gegenbauer series (a problem well analyzed [3]) and the numerical quadrature. Our algorithm to predict the error is very sharp. The sharp bounds allow the method to choose a 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 overestimated 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 novel contributions of this paper can be summarized as follows:

- Development of an efficient Helmholtz multi-level FMM which uses FFTs in the inter/anterpolation steps while retaining diagonal transfer and translation functions. The use of FFTs allows leveraging high performance FFT libraries available for most machines, sequential and parallel.
- An error analysis that accounts for all error in the method and yields constructive algorithms to choose optimal method parameters.
- Details of various optimizations to reduce the computational cost (use of symmetries in the precomputation of the transfer functions, use of symmetries on the unit sphere for the inter/anterpolation steps, optimization of the quadrature points near the poles of the unit sphere, ...).
- Pseudo-codes are provided to clarify the method and help with an implementation by the reader.
- Demonstration of the sharpness of the error bound and the asymptotic computational cost.

The paper is organized as follows. In section 2, we introduce the critical parts of the classical Helmholtz FMM 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, motivating the use of FFT interpolation. 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 analyzes the integration error to derive an algorithm which determines a quadrature 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 1.1 lists the notations used in this paper.

2. The Multilevel Fast Multipole Method. The FMM reduces the computational complexity of the matrix-vector multiplication

$$\sigma_i = \sum_{j \neq i} \frac{e^{\imath\kappa|\boldsymbol{x}_i - \boldsymbol{x}_j|}}{|\boldsymbol{x}_i - \boldsymbol{x}_j|} \psi_j \tag{2.1}$$

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

$$\frac{e^{\imath\kappa|\boldsymbol{r}+\boldsymbol{r}_0|}}{|\boldsymbol{r}+\boldsymbol{r}_0|} = \imath\kappa \sum_{n=0}^{\infty} (-1)^n (2n+1) h_n^{(1)}(\kappa\,|\boldsymbol{r}_0|) j_n(\kappa\,|\boldsymbol{r}|) P_n(\hat{\boldsymbol{r}} \cdot \hat{\boldsymbol{r}}_0) \tag{2.2}$$

The series converges absolutely and uniformly for $|\boldsymbol{r}_0| \geq \frac{2}{\sqrt{3}}\,|\boldsymbol{r}|$ and has been studied extensively in [2, 6].

Truncating the Gegenbauer series at ℓ and using the identity

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

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

$$\frac{e^{\imath\kappa|\boldsymbol{r}+\boldsymbol{r}_0|}}{|\boldsymbol{r}+\boldsymbol{r}_0|} = \int_{S^2} e^{\imath\kappa\hat{\boldsymbol{s}}\cdot\boldsymbol{r}}\, T_{\ell,\boldsymbol{r}_0}(\hat{\boldsymbol{s}})\, \mathrm{d}S(\hat{\boldsymbol{s}}) + \varepsilon_G$$

where ε_G is the Gegenbauer series truncation error and the transfer function, $T_{\ell,\boldsymbol{r}_0}(\hat{\boldsymbol{s}})$, is defined as

$$T_{\ell,\boldsymbol{r}_0}(\hat{\boldsymbol{s}}) = \frac{\imath\kappa}{4\pi} \sum_{n=0}^{\ell} \imath^n (2n+1) h_n^{(1)}(\kappa\,|\boldsymbol{r}_0|) P_n(\hat{\boldsymbol{s}} \cdot \hat{\boldsymbol{r}}_0). \tag{2.3}$$

The reduced computational complexity of the FMM is achieved by constructing a tree of nodes, typically an octree, over the domain of the source and field points. Let $M_\alpha^l(\hat{\boldsymbol{s}})$ be the outgoing field for B_α^l, the box α of the tree in level $l \in [0, L]$ with center $\boldsymbol{c}_\alpha^l$.

Initialization: The method is initialized by computing the outgoing plane-wave expansions for each cluster contained in a leaf of the tree:

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

Upward Pass (M2M): These outgoing expansions are then aggregated upward through the tree by accumulating the product of the child cluster expansions with the plane-wave translation function:

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

Transfer Pass (M2L): Incoming plane-wave expansions, $I_\alpha^l(\hat{\boldsymbol{s}})$ of box B_α^l, are computed from the outgoing by multiplication with the transfer function:

$$I_\alpha^l(\hat{\boldsymbol{s}}) = \sum_{\beta \in \mathcal{I}(B_\alpha^l)} M_\beta^l(\hat{\boldsymbol{s}}) T_{\ell,\boldsymbol{c}_\beta^l - \boldsymbol{c}_\alpha^l}(\hat{\boldsymbol{s}})$$

where $\mathcal{I}(B_\alpha^l)$ is the interaction list of box B_α^l, defined as all boxes of level l which are not neighbors of B_α^l, but whose parent is a neighbor of the parent of B_α^l.

Downward Pass (L2L): The incoming plane-waves are then disaggregated downward through the tree to compute the local field $L_\alpha^l(\hat{\boldsymbol{s}})$ of box B_α^l:

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

where $B_\alpha^l \subset B_\beta^{l-1}$.

Field Computation: 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:

$$\sigma_i = \int_{S^2} L_\alpha^L(\hat{\boldsymbol{s}}) e^{\imath\kappa\hat{\boldsymbol{s}}\cdot(\boldsymbol{c}_\alpha^L - \boldsymbol{x}_i)}\,\mathrm{d}S(\hat{\boldsymbol{s}}) \quad + \sum_{\substack{j,\, j\neq i,\\ \boldsymbol{x}_j \in \mathcal{N}(B_\alpha^L)}} \frac{e^{\imath\kappa|\boldsymbol{x}_i - \boldsymbol{x}_j|}}{|\boldsymbol{x}_i - \boldsymbol{x}_j|}\psi_j$$

where $\boldsymbol{x}_i \in B_\alpha^L$ and $\mathcal{N}(B_\alpha^L)$ is the neighbor list of B_α^L, defined as B_α^L and all neighbor boxes of B_α^L.

2.1. Truncation Parameter in the FMM. The truncation parameter ℓ must be chosen so that the Gegenbauer series (2.2) 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 of the Gegenbauer series is achieved, 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 [5]. This was considered and revised by Darve [6] using a detailed asymptotic analysis of the Gegenbauer series.

The excess bandwidth formula (EBF) is derived from the convergence of the plane-wave spectrum as presented in [3] and reproduced in appendix A. The EBF chooses ℓ as

$$\ell \approx \kappa|\boldsymbol{r}| + C(\kappa|\boldsymbol{r}|)^{1/3} \tag{2.4}$$

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 [14].

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

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

As Carayol and Collino showed in [2], an upper bound of this error for large 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

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

which they showed can be computed in closed form

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

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

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

Carayol and Collino in [1] and [2] present an in-depth analysis of the Jacobi-Anger series and the Gegenbauer series. They find the asymptotic formula

$$\ell \approx \kappa |\boldsymbol{r}| - \frac{1}{2} + \left(\frac{1}{2}\right)^{5/3} W^{2/3}\left(\frac{\kappa |\boldsymbol{r}|}{4\varepsilon^6} \left(\frac{1 + |\boldsymbol{r}_0| / |\boldsymbol{r}|}{1 - |\boldsymbol{r}_0| / |\boldsymbol{r}|} \right)^{3/2} \right)$$

where $W(x)$ is the Lambert function defined as the solution to

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

This appears to be near optimal for large box sizes.

The errors introduced by this truncation have been investigated in other papers including [18, 2, 6].

2.2. Spherical Quadrature in the FMM. The error analysis is simplified if a scheme is used which exactly integrates spherical harmonics, Y_n^m, up to some degree. Below, we enumerate a number of choices that have previously been studied.

1. The most common choice of quadrature uses uniform sample points in ϕ and Gauss-Legendre sample points in $z(\theta)$. 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, $-n \le m \le n$, $0 \le n \le N$ are integrated exactly [7, 18].
2. The simplest choice is a quadrature with uniform sample points 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 [7].
3. McLaren in [20] 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 translation operator, $e^{\imath\kappa \hat{\boldsymbol{s}} \cdot \boldsymbol{r}}$. Its coefficient in the spherical harmonic expansion decreases super-exponentially roughly for $n \gtrsim \kappa |\boldsymbol{r}|$. Therefore, as fields are aggregated in the upward pass and $|\boldsymbol{r}|$ becomes

larger, a larger quadrature is required to resolve 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.

Similarly, as fields are disaggregated in the downward pass, $|\boldsymbol{r}|$ becomes smaller and the higher modes of the incoming field make vanishingly small contributions to the integral as a consequence of Parseval's theorem. Thus, as the incoming field is disaggregated down the tree, a smaller quadrature can be used to resolve it. This makes the integration faster and is actually required to achieve an optimal asymptotic running time. See section 2.4 and appendix B.

There have been several approaches to performing the interpolation and anterpolation between levels in the FMM. Below, we enumerate a number of options that have previously been studied.

1. General, local interpolation methods like Lagrange interpolation, Gaussian interpolation, and B-splines are fast and provide for simple error analysis. In [18] it is shown that the error induced from Lagrange interpolation decreases exponentially as the number of interpolation points is increased for a given function of finite bandwidth. Local interpolation methods almost always have a trade-off between error and speed.
2. For a set of quadrature points (ϕ_k, θ_k), $k = 1, \ldots, K$ with respective weights ω_k and corresponding function value f_k, a spherical harmonic transform maps f_k to a new quadrature $(\phi'_{k'}, \theta'_{k'})$, $k' = 1, \ldots, K'$ via the linear transformation

$$f_{k'} = \sum_{m,l \leq K} Y_l^m(\phi'_{k'}, \theta'_{k'}) \sum_k \omega_k \overline{Y_l^m(\phi_k, \theta_k)} f_k = \sum_k A_{k'k} f_k \tag{2.5}$$

   This transform has nice properties analogous to those of the Fourier transform. A direct computation requires $\mathcal{O}(KK')$ operations which would result in an $\mathcal{O}(N^{4/3})$ or $\mathcal{O}(N^2)$ FMM (see appendix B). Fast spherical transforms (FST) have been developed in [10, 15, 26, 24] and applied to the FMM in [4]. Using the FST reduces the interpolation and anterpolation procedures to $\mathcal{O}(K \log^2 K)$, which results in an $\mathcal{O}(N \log^2 N)$ FMM. However, the accuracy and stability of these algorithms remain in question.
3. Approximations of the spherical transform have also been investigated in [16, 7]. The interpolation matrix $A_{k'k}$ in (2.5) can be sparsified in a number of ways to provide an interpolation/anterpolation method that scales as $\mathcal{O}(K)$ with controllable relative error.
4. Many other interpolation schemes exist with varying running times and errors. Rokhlin presents a fast polynomial interpolator based on the fast multipole method in [11]. See also [17].

2.4. 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.

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

In modeling a uniform distribution of point scatterers over a volumetric domain, we take $\mathcal{O}(N) = \mathcal{O}((\kappa a_0)^3)$ and $L \sim \log(N^{1/3})$ to achieve a total algorithmic complex-

ity of $\mathcal{O}(N \log N)$ using fast global interpolation methods. By using local interpolation methods, this can be reduced to $\mathcal{O}(N)$.

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

It should be noted again that for a given, fixed κa_0 there is a minimum size for the leaves of the tree. Below this critical size, $h_n^{(1)}$ oscillates wildly causing numerical instability in the transfer function. Therefore, when N is very large and the number of levels is saturated, this analysis fails and the algorithm is dominated by the $\mathcal{O}(N^2)$ computation of the close-field contribution, albeit with orders of magnitude speedup over a direct method. In the case when κa_0 is too small, broadband FMMs have been developed as detailed in [9, 8]. However, in many applications κa_0 is large enough to allow for all practical L and N. Furthermore, by keeping the number of points per wavelength constant, the $\mathcal{O}(N \log N)$ behavior can always be achieved.

3. Fourier Based Multilevel Fast Multipole Method. The Fourier based fast multipole method is based on the identity

$$\int_{S^2} e^{\imath\kappa\hat{\boldsymbol{s}}\cdot\boldsymbol{r}}\, T_{\ell,\boldsymbol{r}_0}(\hat{\boldsymbol{s}})\, \mathrm{d}S(\hat{\boldsymbol{s}}) = \int_0^{2\pi}\int_0^{2\pi} e^{\imath\kappa\hat{\boldsymbol{s}}\cdot\boldsymbol{r}}\, T^{\mathrm{s}}_{\ell,\boldsymbol{r}_0}(\hat{\boldsymbol{s}})\, \mathrm{d}\phi\, \mathrm{d}\theta \tag{3.1}$$

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

$$T^{\mathrm{s}}_{\ell,\boldsymbol{r}_0}(\hat{\boldsymbol{s}}) = \frac{1}{2}\, T_{\ell,\boldsymbol{r}_0}(\hat{\boldsymbol{s}})\, |\sin(\theta)| \tag{3.2}$$

is the modified transfer function with $T_{\ell,\boldsymbol{r}_0}(\hat{\boldsymbol{s}})$ defined in equation (2.3). Noting that the integrand is continuous and periodic, this formulation of the problem suggests the use of the Fourier functions $\{e^{\imath n\phi} e^{\imath m\theta}\}$ which form an orthonormal basis of $L^2([0, 2\pi] \times [0, 2\pi])$. This allows i) using two dimensional uniform quadratures; ii) fast Fourier transforms in the interpolation and anterpolation steps; and 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 truncation of the Gegenbauer series and the finite quadrature 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 B). 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

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

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

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

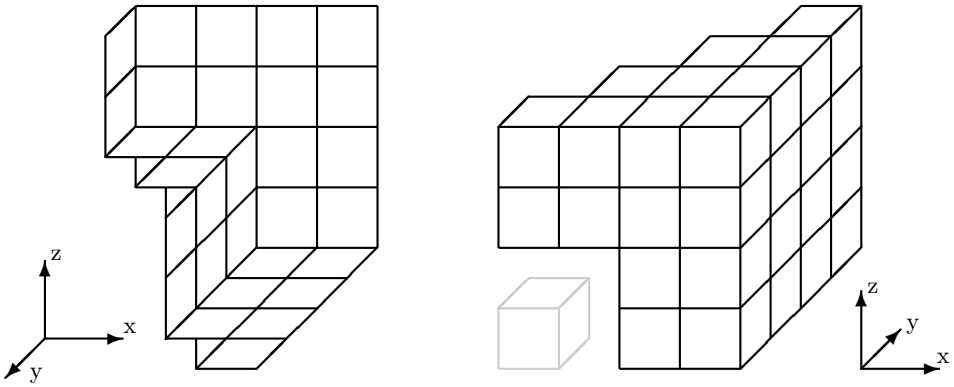

FIG. 3.1. *The center of each box represents one transfer vector $\boldsymbol{r}_0$ which must be computed. The pictures on the left and right panels represent the same set of boxes viewed under two different angles. Due to the symmetries of the quadrature, we need only compute transfer vectors with $x, y, z \geq 0$ 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.*

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 Figure 3.1.

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

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

Since the spectrum of the plane-wave,

$$e^{\imath\kappa\hat{\boldsymbol{s}}\cdot\boldsymbol{r}} = e^{\imath\kappa|\boldsymbol{r}|\cos(\varphi_{\hat{\boldsymbol{s}},\boldsymbol{r}})} = \sum_{n=-\infty}^{\infty} \imath^n J_n(\kappa\,|\boldsymbol{r}|) e^{\imath n\varphi_{\hat{\boldsymbol{s}},\boldsymbol{r}}} \tag{3.3}$$

where $\varphi_{\hat{\boldsymbol{s}},\boldsymbol{r}}$ is the angle between $\hat{\boldsymbol{s}}$ and $\boldsymbol{r}$, decays very 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 to the final integral as a result of Parseval's theorem. By removing these frequencies from the modified transfer function, a smaller quadrature can be used without affecting the final result.

The idea of accurately calculating the integral of a product of two functions by analytically removing high frequencies in one of the two functions can be found in many other papers dealing with the fast multipole method for the Helmholtz kernel in the high frequency regime, e.g. [7, 6]. In the context of these papers, the smoothing

operation (removal of high frequencies) is often termed anterpolation or subsampling. A similar idea is found in Sarvas et al. [25]. In McKay Hyde et al. [19] (Appendix A, p. 254–257), this idea is used in the more general context of calculating the integral of the product of a discontinuous function with a C^1 piecewise-smooth and periodic function.

Suppose we have chosen a quadrature characterized by (N_θ, N_ϕ). With this quadrature we are able to exactly resolve the frequencies in $e^{\imath\kappa\hat{\boldsymbol{s}}\cdot\boldsymbol{r}}$ between $-N_\theta/2+1$ and $N_\theta/2-1$ to the integral in equation (3.1). Consequently, we need to exactly calculate a band-limited approximation of $T^{\mathrm{s}}_{\ell,\boldsymbol{r}_0}$, called $T^{\mathrm{s},L}_{\ell,\boldsymbol{r}_0}$, such that:

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

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

for ϕ_i, $0 \le i < N_\phi/2$, **do**
  $T_k \leftarrow \frac{1}{2} T_{\ell,\boldsymbol{r}_0}(\phi_i, \frac{2\pi k}{2\ell+1})$, $k = 0, \ldots, 2\ell$;
  $\widetilde{T}_m \leftarrow \mathcal{F}(m, T)$;
  $\widetilde{s}_m \leftarrow \mathcal{F}(|m| \le N_\theta/2-1+\ell; |\sin(\theta)|)$;
  $\widetilde{T}^{\mathrm{s},L}_n \leftarrow \widetilde{s} \otimes \widetilde{T}$ convolution of Fourier series;
  $\widetilde{T}^{\mathrm{s},L}_n \leftarrow$ truncate to frequencies $|n| \le N_\theta/2-1$;
  $T^{\mathrm{s},L}(\theta_j, \phi_i) \leftarrow$ inverse transform of $\widetilde{T}^{\mathrm{s},L}_n$;

Algorithm 1: Compute the low-pass modified transfer function.

This algorithm yields the low-pass modified transfer function at (ϕ_i, θ_j), $0 \le i < N_\phi/2$, $0 \le j < N_\theta$ 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 3.2.

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 [13] 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

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

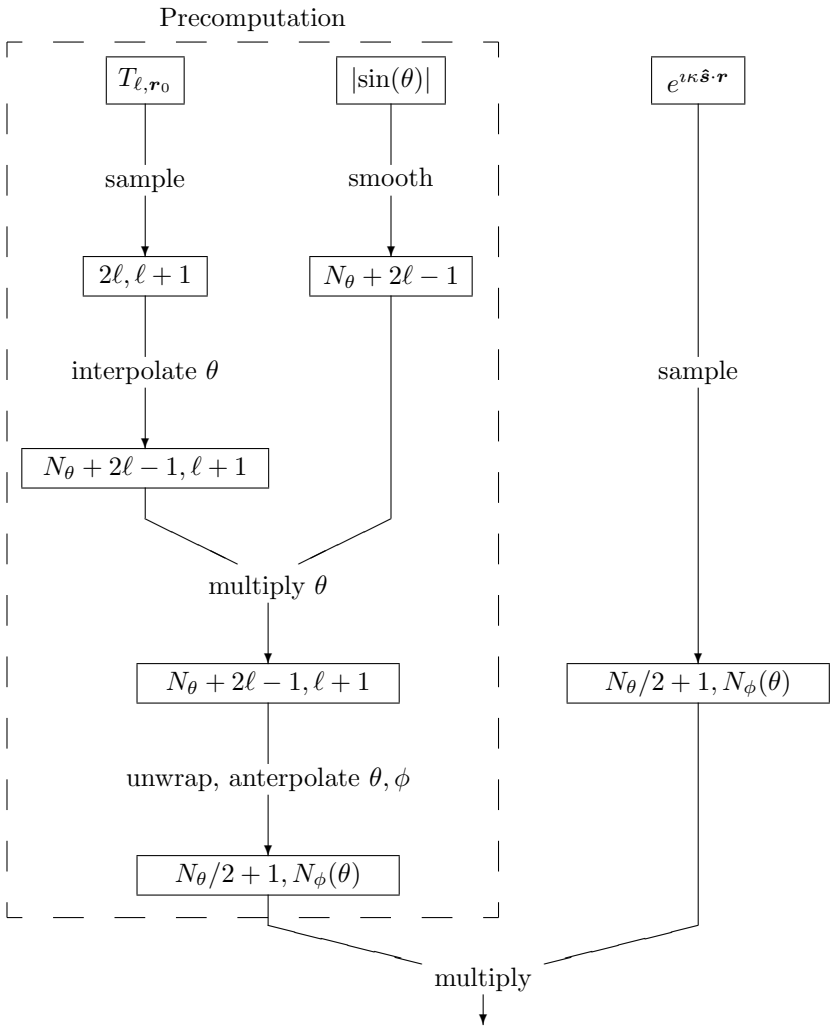

FIG. 3.2. *Procedure for precomputing the low-pass modified transfer function and its application to an outgoing field. The boxed numbers (e.g.,* $\boxed{2\ell, \ell+1}$ *) give the numbers of quadrature points for* θ *and* ϕ *(*N_θ *and* N_ϕ*) at each stage.*

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

$$
\begin{aligned}
&= \Big| \int_0^{2\pi}\int_0^{2\pi} \left[E^L_{\kappa\boldsymbol{r}}(\hat{\boldsymbol{s}}) + E^H_{\kappa\boldsymbol{r}}(\hat{\boldsymbol{s}})\right] \left[T^{\mathrm{s},L}_{\ell,\boldsymbol{r}_0}(\hat{\boldsymbol{s}}) + T^{\mathrm{s},H}_{\ell,\boldsymbol{r}_0}(\hat{\boldsymbol{s}})\right] \mathrm{d}\phi\,\mathrm{d}\theta \\
&\quad - \sum_{n=1}^{N_\theta}\sum_{m=1}^{N_\phi} \omega_{n,m} \left[E^L_{\kappa\boldsymbol{r}}(\hat{\boldsymbol{s}}_{n,m}) + E^H_{\kappa\boldsymbol{r}}(\hat{\boldsymbol{s}}_{n,m})\right] T^{\mathrm{s},L}_{\ell,\boldsymbol{r}_0}(\hat{\boldsymbol{s}}_{n,m}) \Big|
\end{aligned}
$$

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

$$
\begin{aligned}
&= \left| \int_0^{2\pi} \int_0^{2\pi} E^H_{\kappa \boldsymbol{r}}(\hat{\boldsymbol{s}}) T^{\mathrm{s},H}_{\ell,\boldsymbol{r}_0}(\hat{\boldsymbol{s}}) \, \mathrm{d}\phi \, \mathrm{d}\theta - \sum_{n=1}^{N_\theta} \sum_{m=1}^{N_\phi} \omega_{n,m} E^H_{\kappa \boldsymbol{r}}(\hat{\boldsymbol{s}}_{n,m}) T^{\mathrm{s},L}_{\ell,\boldsymbol{r}_0}(\hat{\boldsymbol{s}}_{n,m}) \right| \\
&= \left| \int_0^{2\pi} \int_0^{2\pi} E^H_{\kappa \boldsymbol{r}}(\hat{\boldsymbol{s}}) T^{\mathrm{s},H}_{\ell,\boldsymbol{r}_0}(\hat{\boldsymbol{s}}) - E^{AH}_{\kappa \boldsymbol{r}}(\hat{\boldsymbol{s}}) T^{\mathrm{s},L}_{\ell,\boldsymbol{r}_0}(\hat{\boldsymbol{s}}) \, \mathrm{d}\phi \, \mathrm{d}\theta \right|
\end{aligned}
$$

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

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

In Fourier space, this becomes

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

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,

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

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 3.3.

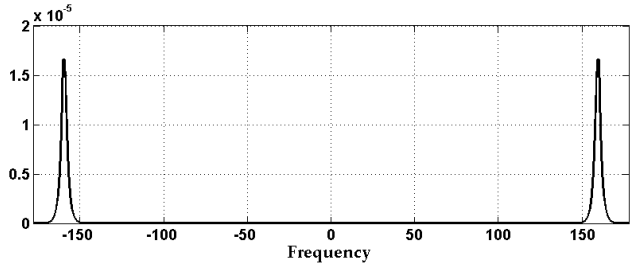

FIG. 3.3. *The value of* $\left| \widetilde{E}^H_{\kappa \boldsymbol{r}}(0,m) - \widetilde{E}^{AH}_{\kappa \boldsymbol{r}}(0,m) \right|$ *for* $\kappa \left| \boldsymbol{r} \right| = 0.8\sqrt{3} \cdot 100$ *and* $N_\phi = 318$.

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

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

Using the plane wave spectrum (3.3) with $\hat{\boldsymbol{r}} = \hat{\boldsymbol{z}}$, this is approximately simplified to

$$|\varepsilon_I| \leq 4\pi^2 \sum_{m=-\infty}^{\infty} \left|J_{M(N_\theta,m)}(\kappa\,|\boldsymbol{r}|)\right| \left|\widetilde{T}^{\mathrm{s}}_{\ell,\boldsymbol{r}_0}(0,m)\right| \tag{3.4}$$

where

$$M(N_\theta, m) = \begin{cases} N_\theta - |m| & |m| \leq N_\theta/2 - 1 \\ |m| & |m| > N_\theta/2 - 1 \end{cases} \tag{3.5}$$

This is an approximation because $\widetilde{E}^{AH}_{\kappa \boldsymbol{r}}$ would, in principle, contribute an infinite sum to equation (3.4) (to include the Bessel functions of order $m + pN_\theta$ for all $p \in \mathbb{Z}\backslash\{0\}$) rather than the single term ($p = \pm 1$) used. However, given the exponential decay of the Jacobi-Anger series, the difference is extremely negligible. Equation (3.4) can be used to search for a value N_θ via the algorithm sketched below:

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

Algorithm 2: Compute N_θ.

Since N^n_θ 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 in terms of N_ϕ 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 fixed θ_j and note that the plane-wave (3.3) can be expressed as

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

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 EBF in [3]. However, we find that the EBF generated quadrature typically overestimates the sampling rate.

To accurately compute $N_\phi(\theta_j)$ the same procedure as in section 3.2.1 is applied but with $\boldsymbol{r}$ and $\boldsymbol{r}_0$ in the xy-plane, 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

$$|\varepsilon_I| \leq 4\pi^2 \sum_{n=-\ell}^{\ell} \left|J_{M(N_\phi(\theta_j),n)}(\kappa\,|\boldsymbol{r}|\sin(\theta_j))\right| \left|\widetilde{T}^{\mathrm{s}}_{\ell,\boldsymbol{r}_0}(n;\theta_j)\right| \tag{3.6}$$

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

```
Set N_φ at the poles: N_φ(θ_0) = N_φ(θ_{N_θ/2}) = 1;
Choose N_φ^n sufficiently larger than 2ℓ + 1;
for θ_j, j = 1, ..., N_θ/2 − 1 do
    T_k ← T^{s,L}_{ℓ,|r_0|x̂}(2πk/(2ℓ+1), θ_j), k = 0, ..., 2ℓ;
    T̃_m ← |F(m; T_k)|;
    Ẽ_m ← |J_m(κ|r| sin(θ_j))|;
    for N_φ(θ_j) from 2 to N_φ^n by 2 do
        Ẽ*_m ← Ẽ_{M(N_φ(θ_j),m)};
        if Ẽ* · T̃ < ε/4π² then
            Save N_φ(θ_j)
```

Algorithm 3: Compute $N_\phi(\theta_j)$.

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 section 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 $|\boldsymbol{r}|$ and $|\boldsymbol{r}_0|$. The previous algorithms require representative values of $|\boldsymbol{r}|$ and $|\boldsymbol{r}_0|$ for each level of the tree. The worst-case transfer vectors, $\boldsymbol{r}_0$, are those with smallest length. If a_l is the box size at level l, then $|\boldsymbol{r}_0| = 2a_l$ is the smallest transfer vector length in the common one buffer-box case.

The worst-case value of $|\boldsymbol{r}|$ is the largest. For a box of size a_l, $|\boldsymbol{r}| \leq a_l\sqrt{3}$. However, using $|\boldsymbol{r}| = a_l\sqrt{3}$ in the previous methods is often too conservative. This case only occurs when two points are located in the exact corners of the boxes – a rare case indeed. See Figure 3.4. 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

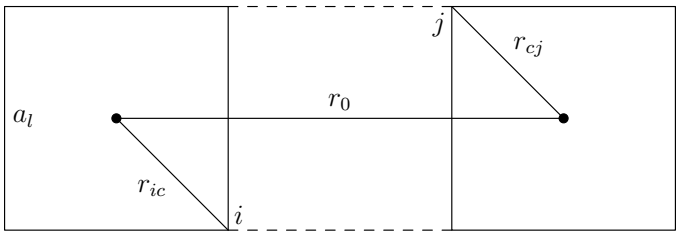

FIG. 3.4. *The worst-case $\boldsymbol{r}$ and $\boldsymbol{r}_0$, projected from the 3D box. Here, $|\boldsymbol{r}_0| = 2a_l$ and i and j are on the opposite corners of the box so that $|\boldsymbol{r}| = |\boldsymbol{r}_{ic}| + |\boldsymbol{r}_{cj}| = a_l\sqrt{3}$.*

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, $-n \le m \le n$, $0 \le n \le N$ are integrated exactly. In [18], 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

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

where $C_1, C_2 \ge 1$ are small integers dependent on ℓ, numerically computed from the methods in section 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 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 efficient and easy to implement while retaining the exactness of global methods.

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

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

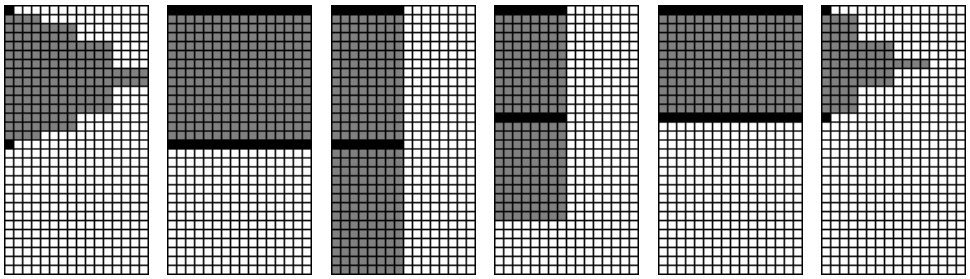

FIG. 3.5. *The data profile at each step in an anterpolation from a large quadrature Q with $N_\theta = 30$ to a smaller quadrature Q' with $N'_\theta = 24$. The angle ϕ is in the x direction while the angle θ is in the y direction. The data corresponding to a pole has been darkened for clarity.*

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 3.5, perform an exact interpolation/anterpolation using only FFTs.

1. For each θ_j, $0 \le j \le N_\theta/2$, Fourier interpolate the data $[F(\phi_{i=0,\ldots,N_\phi(\theta_j)-1}, \theta_j)]$ 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 $[F(\phi_i, \theta_{j=0,\ldots,N_\theta/2}), F(\phi_{i+\mathcal{N}_\phi/2}, \theta_{j=N_\theta/2-1,\ldots,1})]$.
3. For each ϕ_i, $0 \le i < \mathcal{N}_\phi/2$, Fourier interpolate the data $[F(\phi_i, \theta_{j=0,\ldots,N_\theta-1})]$ from length N_θ to N'_θ.
4. For each ϕ_i, $0 \le i < \mathcal{N}_\phi/2$, unwrap the data $[F(\phi_i, \theta_{j=0,\ldots,N'_\theta-1})]$ to construct the sequences $[F(\phi_i, \theta_{j=0,\ldots,N'_\theta/2})]$ and $[F(\phi_{i+\mathcal{N}_\phi/2}, \theta_{j=0,\ldots,N'_\theta/2})]$.
5. For each θ_j, $0 \le j \le N'_\theta/2$, Fourier anterpolate the data $[F(\phi_{i=0,\ldots,\mathcal{N}_\phi(\theta_j)-1}, \theta_j)]$ 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 (3.1). For a given box size a, the quadrature 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

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

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 3.6, 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 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 efficient.

To show that the optimal asymptotic running time is achieved, Figure 3.7 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. The target error is set to 10^{-4} and is achieved in every case. 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 $L = 7$.

A performance comparison of our method with other published FMMs is beyond the scope of this paper. As of now, there is no reference implementation of the Helmholtz FMM that is recognized as a benchmark. In addition, the Helmholtz FMM is a complex scheme with many steps and requires complex error analysis to find the optimal parameters and ensure a given target error. Developing another FMM that could be considered a reference benchmark FMM would be a significant study in itself.

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 ap-

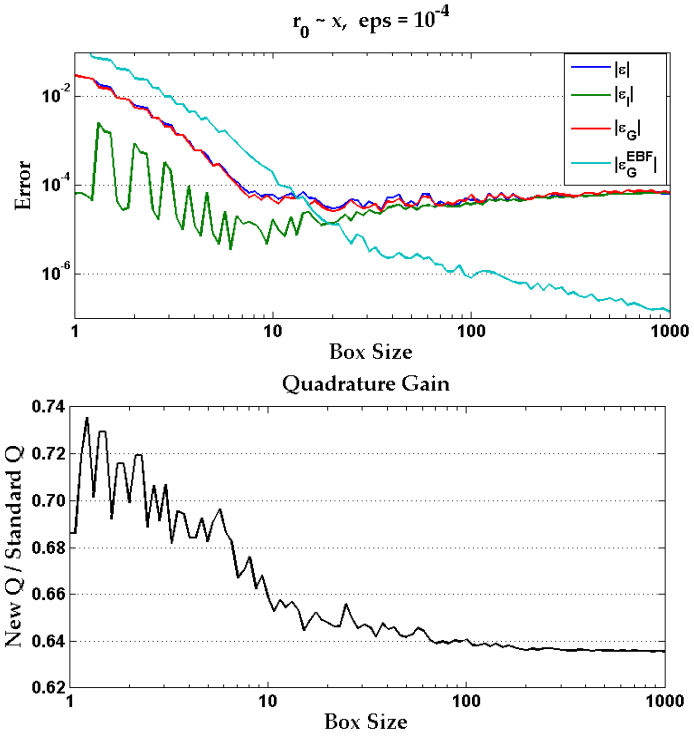

FIG. 3.6. *Top plots: error of the FMM integral using a direct computation of ℓ as described in section 2.1 and the choice of quadrature described in section 3.2. All the errors fall very close to the target error of 10^{-4}. The standard EBF overestimates ℓ and will result in a suboptimal quadrature. 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 from section 3.2.4.*

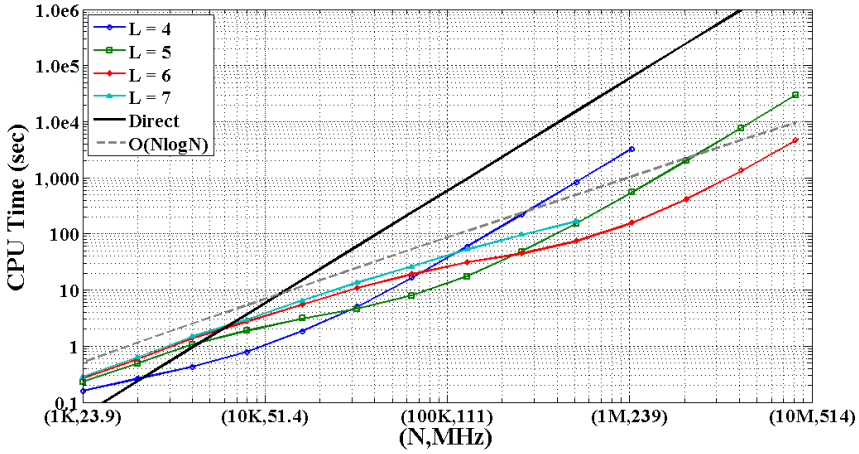

FIG. 3.7. *Average running times of the Fourier based FMM for constant number of volumetric points per wavelength.*

proximating the Helmholtz kernel with

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

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. This is realized by removing the high frequency components of the modified transfer function, $T^{\mathrm{s}}_{\ell,\boldsymbol{r}_0}(\hat{\boldsymbol{s}})$, during the precomputation phase 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 the sampling rate in the θ-direction. The truncation error ε_G 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 finite sampling 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 highly efficient FFT algorithms are available in virtually every computing environment, the time-critical interpolation stages of the algorithm are much easier to implement efficiently.

Acknowledgments. This work was supported by grants from the Stanford School of Engineering, the Army High Performance Computing and Research Center at Stanford University, and the King Abdullah University of Science and Technology.

Appendix A. Excess Bandwidth Formula.

The EBF is derived in [3] and a short explanation is reproduced here. To determine an appropriate truncature, the spectrum of a plane-wave (3.3) is used to estimate how many terms in the series are needed before the error 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

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

Appendix B. 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. 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. **Field Computation:** 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.

REFERENCES

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

[2] ———, *Error estimates in the fast multipole method for scattering problems. Part 2: Truncation of the Gegenbauer series*, ESAIM: M2NA, 39 (2004), pp. 183–221.

[3] W.C. Chew, E. Michielssen, J. M. Song, and J. M. Jin, eds., *Fast and Efficient Algorithms in Computational Electromagnetics*, Artech House, Inc., Norwood, MA, USA, 2001.

[4] Indranil Chowdhury and Vikram Jandhyala, *Integration and interpolation based on fast spherical transforms for the multilevel fast multipole method*, Microwave and Optical Technology Letters, 48 (2006), pp. 1961–1964.

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

[6] Eric Darve, *The fast multipole method I: Error analysis and asymptotic complexity*, SIAM J. Numer. Anal., 38 (2000), pp. 98–128.

[7] ———, *The fast multipole method: Numerical implementation*, J. Comput. Phys., 160 (2000), pp. 195–240.

[8] Eric Darve and Pascal Havé, *Efficient fast multipole method for low-frequency scattering*, J. Comput. Phys., 197 (2004), pp. 341–363.
[9] ———, *A fast multipole method for Maxwell equations stable at all frequencies*, Phil. Trans. R. Soc. Lond. A, 362 (2004), pp. 603–628.
[10] James R. Driscoll and D.M. Healy, *Computing Fourier transforms and convolutions on the 2-sphere*, Adv. Appl. Math., 15 (1994), pp. 202–250.
[11] A. Dutt, M. Gu, and V. Rokhlin, *Fast algorithms for polynomial interpolation, integration, and differentiation*, SIAM J. Numer. Anal., 33 (1996), pp. 1689–1711.
[12] N. Engheta, W. D. Murphy, V. Rokhlin, and M. S. Vassiliou, *The fast multipole method (FMM) for electromagnetic scattering problems*, IEEE Trans. Antennas Propag., 40 (1992), p. 634.
[13] O. Ergul and L. Gurel, *Optimal interpolation of translation operator in multilevel fast multipole algorithm*, IEEE Trans. Antennas Propag., 54 (2006), pp. 3822–3826.
[14] M. L. Hastriter, S. Ohnuki, and W. C. Chew, *Error control of the translation operator in 3D MLFMA*, Microwave and Optical Technology Letters, 37 (2003), pp. 184–188.
[15] D.M. Healy, D. Rockmore, P.J. Kostelec, and S. Moore, *FFTs for the 2-sphere - improvements and variations*, J. Fourier Analysis and Applications, 9 (2003), pp. 341–385.
[16] Rüdiger Jakob-Chien and Bradley K. Alpert, *A fast spherical filter with uniform resolution*, J. Comput. Phys., 136 (1997), pp. 580–584.
[17] J. Knab, *Interpolation of bandlimited functions using the approximate prolate series (corresp.)*, Information Theory, IEEE, 25 (1979), pp. 717–720.
[18] S. Koc, J. M. Song, and W. C. Chew, *Error analysis for the numerical evaluation of the diagonal forms of the scalar spherical addition theorem*, SIAM J. Numer. Anal., 36 (1999), pp. 906–921.
[19] E. McKay Hyde and OP Bruno, *A fast higher-order solver for scattering by penetrable bodies in three dimension*, J. Comp. Phys., 202 (2005), pp. 236–261.
[20] A. D. McLaren, *Optimal numerical integration on a sphere*, Mathematics of Computation, 17 (1963), pp. 361–383.
[21] J. Rahola, *Diagonal forms of the translation operators in the fast multipole algorithm for scattering problems*, BIT, 36 (1996), pp. 333–358.
[22] V. Rokhlin, *Rapid solution of integral equations of scattering theory in two dimensions*, J. Comput. Phys., 86 (1990), pp. 414–439.
[23] ———, *Diagonal forms of translation operators for the Helmholtz equation in three dimensions*, Applied and Computational Harmonic Analysis, 1 (1993), pp. 82–93.
[24] V. Rokhlin and M. Tygert, *Fast algorithms for spherical harmonic expansions*, SIAM J. Scientific Computing, 27 (2006), pp. 1903–1928.
[25] Jukka Sarvas, *Performing interpolation and anterpolation entirely by fast Fourier transform in the 3-D multilevel fast multipole algorithm*, SIAM J. Numer. Anal., 41 (2003), pp. 2180–2196.
[26] Reiji Suda and Masayasu Takami, *A fast spherical harmonics transform algorithm*, Math. of Comp., 71 (2001), pp. 709–715.