Parallel Wideband MLFMA for Analysis of Electrically Large, Nonuniform, Multiscale Structures

Stephen Hughey, *Student Member, IEEE,* H. M. Aktulga, Melapudi Vikram, Mingyu Lu, *Senior Member, IEEE,* Balasubramaniam Shanker, *Fellow, IEEE,* and Eric Michielssen, *Fellow, IEEE*

***Abstract*— Electromagnetic scattering from electrically large objects with multiscale features is an increasingly important problem in computational electromagnetics. A conventional approach is to use an integral equation-based solver that is then augmented with an accelerator, a popular choice being a parallel multilevel fast multipole algorithm (MLFMA). One consequence of multiscale features is locally dense discretization, which leads to low-frequency breakdown and requires nonuniform trees. To the authors' knowledge, the literature on parallel MLFMA for such multiscale distributions capable of arbitrary accuracy is sparse; this paper aims to fill this niche. We prescribe an algorithm that overcomes this bottleneck. We demonstrate the accuracy (with respect to analytical data) and performance of the algorithm for both PEC scatterers and point clouds as large as 755λ with several hundred million unknowns and nonuniform trees as deep as 16 levels.**

***Index Terms*— Adaptive algorithms, computational electromagnetics, method of moments (MoM), multilevel fast multipole algorithm (MLFMA), parallel algorithms.**

I. Introduction

OVER the past several decades, computer simulation has become an indispensable tool for designing and prototyping in electromagnetic (EM) engineering. This includes designing microwave/RF circuits and antennas, EM interference mitigation, stealth aircraft profile reduction, and THz devices. Many of these problems, however, pose significant challenges for existing computational methods. The confluence of several factors, namely, steady increase in frequencies of engineering interest, the widespread availability of powerful computational resources, and the demand for realism in computer models is driving a massive increase in the number of degrees of freedom, N_s, required to solve these problems.

While the literature is filled with means to tackle these problems, over the years, surface integral equations have emerged as a principal tool of analysis [1]–[3]. The main bottleneck to widespread use of these methods was their computational complexity; since the 1990s, this has largely been overcome thanks to accelerators such as the multilevel fast multipole algorithm (MLFMA) that have reduced the computational complexity from $\mathcal{O}(N_s^2)$ to $\mathcal{O}(N_s \log N_s)$ for surfaces. As a result of these significant benefits, research into advancing nuances and application of these approaches is extensive and perhaps too numerous to enumerate; a partial list can be found in [4] and [5] and references therein. As is evident from these citations, there is still considerable interest in exploring and expanding the range of these techniques at either end, i.e., extending these methods to problems with high disparity in discretization scales that are then embedded in electrically large objects. At the same time, there has been a concerted effort to develop parallel algorithms to further exploit the capabilities of these algorithms. Both these issues are explored in more detail next.

The literature is rich with efforts to bridge the length scale between the Helmholtz regime and the low-frequency or Laplace regime. The principal challenge is the low-frequency breakdown of the classical MLFMA. The earliest work on this problem was done by Greengard and Rokhlin [6] and Greengard *et al.* [7]. Soon after, Zhao and Chew [8] introduced a modification of the original addition theorem to stabilize the MLFMA at low frequencies. Since these, other methods have emerged, including [9]–[13]. Another method introduced in 2007 based on accelerated Cartesian expansions (ACEs) [14], [15] blends seamlessly and intuitively with the MLFMA and is error-controllable to arbitrary accuracy with no bound on the discretization density.

It is apparent that to extend the reach of fast multipole methods (FMMs), the development of parallel algorithms is a necessity [16]. The earliest efforts were based on extending methods developed for the Laplace (electrostatic) FMM to MLFMA with little success. Recent effort focused

Manuscript received April 3, 2018; revised September 3, 2018; accepted October 19, 2018. Date of publication November 21, 2018; date of current version February 5, 2019. This work was supported in part by NSF under Grant ECCS-1408115 and in part by the Department of Defense High Performance Computing and Modernization Program. *(Corresponding author: Stephen Hughey.)*

S. Hughey and B. Shanker are with the Department of Electrical and Computer Engineering, Michigan State University, East Lansing, MI 48824 USA (e-mail: hugheyst@msu.edu).

H. M. Aktulga is with the Department of Computer Science, Michigan State University, East Lansing, MI 48824 USA.

M. Vikram is with the GE Global Research Center, Bengaluru 560066, India.

M. Lu is with the Department of Electrical and Computer Engineering, West Virginia University Institute of Technology, Beckley, WV 25801 USA.

E. Michielssen is with the Department of Electrical Engineering and Computer Science, University of Michigan, Ann Arbor, MI 48109 USA.

Color versions of one or more of the figures in this paper are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/TAP.2018.2882621

on building algorithms based on hierarchical partitioning (HiP-MLFMA) [17] and more recently, its block-wise variant B-HiP-MLFMA [18]. This approach works by transitioning from partitioning in space at lower levels to directions at higher levels; B-HiP-MLFMA optimizes the manner, in which the directions are partitioned. In these methods, one uses a number of processors that is either a power of 2 (for HiP-MLFMA) or a power of 4 (for B-HiP-MLFMA).

Another algorithm [19] that has also shown excellent scalability and performance takes a different approach; using postorder traversal, the resulting self-similarity of the tree as well as direction partitioning yields an algorithm that scales very well [19]. The salient features of this algorithm are: 1) wideband MLFMA for analysis from low to high frequencies; 2) global (exact) interpolation and anterpolation; 3) demonstrated error control (L_2 error norm against analytical data); and 4) adaptive direction partitioning that is dictated by the number of duplicate nodes; for the purposes of this discussion, it will be referred to as AP-MLFMA, for "adaptive partitioning." It was shown that the matrix vector product (matvec) times of AP-MLFMA are highly competitive with others. To a large extent, each of these methods assume a uniform tree.

Despite this progress, several open problems remain, particularly for wideband and multiscale MLFMA systems. These can be summarized as accuracy in deep nonuniform trees and their relative tradeoffs in developing efficient parallel algorithms. To wit, both HiP- and AP-MLFMA offer a solution to a portion of the problem. The former relies on local inter/anterpolation, which has linear asymptotic complexity but requires significant oversampling, resulting in an increase in cost. It is also inaccurate, as one needs a filter at the anterpolation stage [9], [20], [21]. On the positive side, HiP-MLFMA is easier to parallelize. The AP-MLFMA uses global (exact) inter/anterpolations so the accuracy can be well controlled, and the sampling is optimal and predetermined. However, developing parallel algorithms for this approach is a challenge; in addition, memory costs required to store matrices for inter/anterpolation on each node limit the height of the tree. Both HiP- and AP-MLFMA do not include nonuniform trees within their parallel frameworks. Algorithms that address both wideband and multiscale nature of the problems implies that one needs both the capability of handling low frequencies as well as adaptive nonuniform trees [9], [15], [22]; while such algorithms exist in serial, the literature on parallelization is extremely sparse with each [19], [23] addressing a portion of the problem.

The principal contribution of this paper is to address these extant deficiencies; our objective is to provide a parallel computational framework for computing fields arising from realistic distributions with the following properties.

1) It is a methodology whose accuracy can be controlled to desired precision.
2) It includes transitions to low frequencies and nonuniform distributions.
3) The sampling of the spectrum is optimal and overcomes memory bottlenecks associated with global inter/anterpolations.
4) It is efficient in terms of scalability and execution times.

We present a number of results to demonstrate the accuracy of the numerical methods and the performance of the parallel algorithms presented here. The methodology extensively extends the work [15], [19] on ACE-AP-MLFMA.

The rest of this paper is organized as follows. Section II lays out the EM scattering problem to be solved. In Section III, we review the wideband MLFMA, establish the need for filters for anterpolation, develop transitions between spherical harmonics and Fourier series for interpolation/anterpolation, nonuniform trees for multiscale features, and an interpolation method for saving memory. Section IV outlines the parallel algorithm, and finally, Section V demonstrates the error control and performance of the serial and parallel algorithms.

II. Problem Statement

Consider an object residing in a region $\mathcal{D} \in \mathbb{R}^3$ bounded by the surface $\partial\mathcal{D}$, the perfectly conducting portion of which is denoted by S. Let $\{\mathbf{E}^i(\mathbf{r}), \mathbf{H}^i(\mathbf{r})\}$ denote electric and magnetic fields that are incident on this object, μ and ε denote the magnetic permeability and electric permittivity of the medium, respectively, and $k = \omega\sqrt{\mu\varepsilon}$ denote the wavenumber for frequency ω. The intrinsic impedance of the medium is denoted by $\eta = \sqrt{\mu/\varepsilon}$. The field scattered by the object may be obtained using a combined field integral equation (CFIE) formulated in terms of the unknown electric surface current density $\mathbf{J}(\mathbf{r})$ on S. For $\mathbf{r} \in S$

$$\begin{aligned}\alpha\hat{n} \times \hat{n} \times \mathbf{E}^i(\mathbf{r}) + (1-\alpha)\hat{n} \times \mathbf{H}^i(\mathbf{r}) \\ = -\alpha\mathcal{L}\{\mathbf{J}\}(\mathbf{r}) + (1-\alpha)\mathcal{K}\{\mathbf{J}\}(\mathbf{r})\end{aligned} \tag{1}$$

$$\mathcal{L}\{\mathbf{J}\}(\mathbf{r}) \doteq \hat{n} \times \left(\hat{n} \times \int_S \mathbf{G}(\mathbf{r}-\mathbf{r}') \cdot \mathbf{J}(\mathbf{r}')\, dS'\right) \tag{2}$$

$$\mathcal{K}\{\mathbf{J}\}(\mathbf{r}) \doteq \hat{n} \times \frac{1}{jk\eta}\left(\nabla \times \int_S \mathbf{G}(\mathbf{r}-\mathbf{r}') \cdot \mathbf{J}(\mathbf{r}')\, dS'\right) \tag{3}$$

$$\mathbf{G}(\mathbf{r}-\mathbf{r}') \doteq -jk\eta\left[\mathbb{I} + \frac{\nabla\nabla}{k^2}\right] g(\mathbf{r}-\mathbf{r}') \tag{4}$$

where

$$g(\mathbf{r}-\mathbf{r}') \doteq \frac{e^{-jk|\mathbf{r}-\mathbf{r}'|}}{4\pi|\mathbf{r}-\mathbf{r}'|} \tag{5}$$

is the scalar Helmholtz Green's function and $\hat{n} \doteq \hat{n}(\mathbf{r})$ denotes the unit normal to S at the point $\mathbf{r}$. Solving for $\mathbf{J}$ via the method of moments (MoM) involves the discretization of (1), typically by expansion and testing with the classic RWG basis functions [24] defined on a triangular mesh. To do so, we expand the unknown surface current as

$$\mathbf{J}(\mathbf{r}) = \sum_{n=1}^{N_s} I_n \mathbf{f}_n(\mathbf{r}) \tag{6}$$

where $\mathbf{f}_n(\mathbf{r})$ denotes the RWG function associated with the nth edge, and N_s is the total number of edges in the mesh. The $N_s \times N_s$ matrix equation resulting from the MoM procedure may be expressed as

$$ZI = V \tag{7}$$

where $I = \{I_1, \ldots, I_{N_s}\}^T$ is the vector of unknown coefficients, and the elements of the matrix Z and vector V are given by Galerkin testing (1).

Solutions to (7) are typically effected using an iterative solver, for instance, the generalized minimal residual (GMRES) method. Coupling this equation with an accelerator amortizes the cost of the matvec; the most popular of the accelerators being the MLFMA method [25]. Next, we discuss challenges and their resolution when these methods are applied for multiscale analysis.

III. Challenges and Remedies for Parallel Multiscale Analysis

All acceleration methods for accelerating the solution of (7) rely on the fact that the matrix Z may be partitioned into a matrix Z^{NF} of near interactions and a matrix Z^{FF} of far interactions, or

$$Z = Z^{NF} + Z^{FF}. \tag{8}$$

This partitioning is founded on the notion that matvecs with Z^{FF} can be computed in $\mathcal{O}(N_s \log N_s)$ time, and those with Z^{NF} can be computed in $\mathcal{O}(N_s)$ time. Developing criteria for such a partitioning for multiscale problems poses challenges from both computational and parallelization perspectives. These challenges arise from the nonuniformity of discretization and their distribution. To set the stage, assume that the spatial distribution of unknowns can be mapped onto a uniform octree data structure. The root of the tree contains all unknowns, and the leaves of the tree correspond to clusters of unknowns, serving as the interface between the tree and the unknowns. The matvec with Z^{FF} is effected by using appropriate operators to traverse up, down, and across the tree. Assume that the minimum size of a leaf is around 0.2λ. In a multiscale scenario, a leaf box may contain too many unknowns, thus dramatically increasing computational costs of near-field interactions and overwhelming the cost complexity of the MLFMA. This can be overcome using a judicious choice of representations for the Green's function that enables a low-frequency stable decomposition, as in [7], [9], and [11]; here, we employ the ACE method [14], [15]. Enabling this results in a tree that could potentially have MLFMA leaves that are then roots of subtrees that extend downward. From a computational perspective, one needs to develop operators to efficiently transfer information between nodes (both at the same level and between levels).

Another bottleneck is accurate traversal up and down the tree. As will be shown later, a provably accurate approach is to use a global representation of radiated field or use a local bandlimited representation [26]–[28]. For the latter, the price paid is the oversampling required for the same accuracy. However, the downside of using global methods is the challenges it poses in terms of memory requirements for storing interpolation coefficients at higher levels in the tree as well as parallelization. As a result, an efficient strategy for deep trees would use a spherical filter that uses spherical sampling up to a certain level, switches to a fast Fourier transform (FFT)-based filter using uniform sampling [29] to save on memory, and then a local bandlimited filter with linear complexity to leverage parallelism. This concept is illustrated in Fig. 1. While we leave the discussion of local bandlimited interpolation and modification of parallel algorithms to [30], we shall extensively discuss spherical and uniform sampling, transitions between the two, and parallelization strategies that one can employ. In what follows, we will explore each of these issues in greater detail. Note that, the exposition will focus on evaluation of the scalar Green's function; changes necessary to effect these operations for the dyadic Green's function are well-known and only those subtleties that are not well-known are elaborated upon.

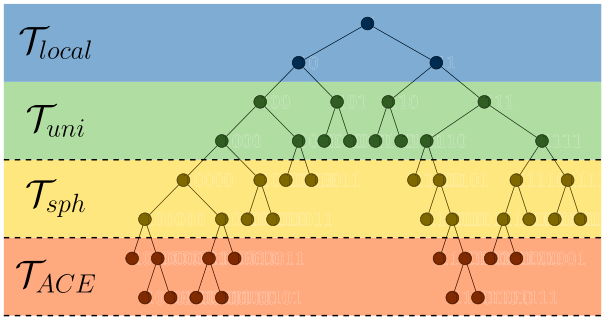

Fig. 1. Illustration of the target computational strategy for representing interactions in a general multiscale geometry. Each $\mathcal{T}$ represents operations performed with the subscripted method, e.g., ACE, spherical and uniform interpolation/anterpolation, and local band-limited interpolation/anterpolation.

A. Nonuniform Trees

MLFMA relies on construction of octrees to partition near and far interactions. Here, we briefly discuss creating such lists for nonuniform distributions. Consider a box b at any level in the tree; we denote its parent by $P(b)$ and its grandparent by $P^2(b)$ and so on up the tree. Boxes that share a spatial location (vertex, edge, or face) with b are identified as being in its near field. If two boxes are of the same size and share a spatial location, they belong to each other's U-lists. Likewise, if two boxes are of the same size, do not share a spatial location, and their parents are in the near field of each other, they are in each other's far-field or V-lists. To handle nonuniform distributions, we develop an adaptive tree [22] starting with a uniform tree representation, where all leaf boxes reside precisely at the same level, merging siblings subject to the rule that their parent does not contain more than some specified s_{max} DoF.

An important consequence of adaptive methods is the introduction of far-field interactions between boxes at different levels of the tree, or cross-level interactions. These interactions, deemed the X- and W-lists, are discovered during the construction of the near and far interaction lists in the following manner. Consider a box b. To construct its interaction lists, we generate all hypothetically interacting boxes with b at the same level according to the scheme described earlier, and we perform a top-down search of the octree to locate each such hypothetical box h. If h itself is found in the tree, it is obviously added to the appropriate U- or V-list. If the search for h dead-ends on an ancestor $P^n(h)$ that is not a leaf,

the interaction pair (b, h) is discarded; however, if $P^n(h)$ is a leaf and it satisfies the conditions of the present interaction list, we retain $(b, P^n(h))$ in the X-list and its reciprocal interaction pair $(P^n(h), b)$ in the W-list.

The computational procedure is as follows.

1) Construct the nonuniform tree and delineate near and far-field interactions.
2) Precompute and store near field interactions.
3) For any matvec,
   a) Compute contributions due to the near field.
   b) Compute contributions due to the far field:
      i) For all leaves, compute charge to multipole (C2M).
      ii) Construct multipole information for all boxes from those of their children [multipole-to-multipole (M2M)]. This operation is level dependent, and one needs the following operators: $\mathrm{M}_{ACE} \longrightarrow \mathrm{M}_{ACE}$, $\mathrm{M}_{ACE} \longrightarrow \mathrm{M}_{sph}$, $\mathrm{M}_{sph} \longrightarrow \mathrm{M}_{sph}$, $\mathrm{M}_{sph} \longrightarrow \mathrm{M}_{uni}$ and $\mathrm{M}_{uni} \longrightarrow \mathrm{M}_{uni}$.
      iii) Translate from multipole to local (M2L) expansion. For nonuniform trees, this would involve constructing operators for both in- and across-level translations.
      iv) Construct local expansions for all boxes from those of its parent [local-to-local (L2L)]. This is a conceptual inverse of the M2M operation and one needs to develop identical operators to traverse down the tree.
      v) From local expansions at leaves, construct far-field information at all observers [local-to-observer (L2O)].
   c) Sum the contributions of near and far fields.

B. Low-Frequency Analysis

Over distances or length scales which are small relative to a wavelength ($< 0.25\lambda$ or so), the oscillatory behavior of g over larger distances gives way to quasi-static behavior which can be accurately modeled using Taylor series expansions. This is the foundation of the ACE method [14], [15] for the low-frequency Helmholtz problem.

Consider the evaluation of the potential $\Psi(\mathbf{r})$ due to a source distribution $u(\mathbf{r})$

$$\Psi(\mathbf{r}) = \sum_{n=1}^{N_s} u_n g(\mathbf{r} - \mathbf{r}'_n). \tag{9}$$

To compute the potentials between regions that are sufficiently separated, ACE [14], [15] dictates that the potential due to s' sources with location and strength $\{\mathbf{r}_m, u_m\}$, $m = 1, \ldots, s'$ within the region Ω_s centered at $\mathbf{r}_s^c$ can be written as

$$\Psi^{FF}(\mathbf{r}) = \sum_{n=0}^{\infty} \mathbf{M}^{(n)}\left(\mathbf{r}_s^c\right) \cdot n \cdot \nabla^{(n)} g\left(\mathbf{r} - \mathbf{r}_s^c\right) \tag{10}$$

where $\cdot n \cdot$ denotes an n-fold tensor contraction, $\mathbf{M}^{(n)}$ is the multipole tensor representing the sources within Ω_s, and the tensor $\nabla^{(n)} g(\mathbf{r} - \mathbf{r}_s^c)$ is an n-fold tensor operand on the Green's function. Details for these expressions, as well as aggregation operators $\mathrm{M}_{ACE} \longrightarrow \mathrm{M}_{ACE}$ and $\mathrm{M}_{ACE} \longrightarrow \mathrm{M}_{sph}$ and their counterparts for disaggregation, are available in [15]. Operators addressing translations across levels are presented later in this section.

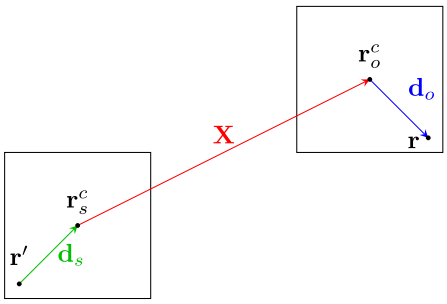

Fig. 2. Graphical illustration of the addition theorem and some notation.

C. MLFMA

The derivation of MLFMA and its multilevel variant are well-known [25], [29], [31], [32]. Given the wealth of papers in this area over the past few decades, it seems somewhat presumptuous to contend that there is significant new information to be added to the literature. The purpose of this section is somewhat different. In keeping with our goal of controllable accuracy, we will describe methods that permit both better understanding and error control. We start with a 2-level description of MLFMA that is based on the integral representation of the Green's function

$$g(\mathbf{X} + \mathbf{d}) \approx -\frac{jk}{(4\pi)^2} \int_{S^2} e^{-j\mathbf{k}\cdot(\mathbf{d}_s + \mathbf{d}_o)} T(\mathbf{k}, \mathbf{X}) d^2\hat{k} \tag{11}$$

where the translation operator T is given by

$$T(\mathbf{k}, \mathbf{X}) \doteq \sum_{n=0}^{\infty} (-j)^n (2n+1) h_n^{(2)}(kX) P_n(\hat{k} \cdot \hat{X}) \tag{12}$$

with $\mathbf{k} = k\hat{k}$. Here, S^2 denotes the unit sphere, parametrized by $(\theta, \phi) \in [0, \pi] \times [0, 2\pi]$. We note that $\mathbf{k} = \mathbf{k}(\theta, \phi)$, and use these notations interchangeably. Fig. 2 illustrates the decomposition of $|\mathbf{r} - \mathbf{r}'|$ into $|\mathbf{X} + \mathbf{d}(1)|$, where $\mathbf{d}(l)$ denotes the sum of particle-to-center vectors at level l with $l = 1$ being the leaf level. We denote the centers of boxes containing $\mathbf{r}'$ and $\mathbf{r}$ as $\mathbf{r}_s^c(1)$ and $\mathbf{r}_o^c(1)$, respectively. It follows that $\mathbf{d} = \mathbf{r} - \mathbf{r}_o^c(1) + \mathbf{r}_s^c(1) - \mathbf{r}'$. The rules to evaluate the spectral integral are well-understood and in place [9], [15], [20], [29], [32]. The far-field part of the potential (9) is evaluated using

$$\Psi^{FF}(\mathbf{r}) \approx -\frac{jk}{(4\pi)^2} \int_{S^2} e^{-j\mathbf{k}\cdot\mathbf{d}_o(1)} U_1(\theta, \phi) d^2\hat{k} \tag{13}$$

where the local expansion U_1 of the observer box is

$$U_1(\theta, \phi) \doteq T(\mathbf{k}, \mathbf{X}) V_1(\mathbf{k}) \tag{14}$$

and V_1 denotes the multipole expansion of the source box

$$V_1(\theta, \phi) = \sum_{i=1}^{s'} u_i e^{-j\mathbf{k}\cdot\left(\mathbf{r}_s^c(1) - \mathbf{r}'_i\right)}. \tag{15}$$

In what follows, we discuss the intimate connection between integration, interpolation ($V_l(\theta,\phi) \longrightarrow V_{l+1}(\theta,\phi)$), and anterpolation ($U_{l+1}(\theta,\phi) \longrightarrow U_l(\theta,\phi)$).

1) Inter/Anterpolation and Spectral Integration: Let us revisit (11). It is well-known that the plane wave expansion gives us

$$e^{-j\mathbf{k}\cdot(\mathbf{d}_s(1)+\mathbf{d}_o(1))} \approx S(kd,\theta,\phi) = \sum_{n=0}^{N_h(1)} \sum_{m=-n}^{n} a_{nm} Y_{nm}(\theta,\phi) \tag{16}$$

where $d = |\mathbf{d}(1)| = |\mathbf{d}_s(1)+\mathbf{d}_o(1)|$, and a_{nm} are known coefficients of the normalized spherical harmonics, $Y_{nm}(\theta,\phi)$, used to represent the plane wave. Given that $T(\mathbf{k},\mathbf{X})$ can be written as $T(\mathbf{k},\mathbf{X}) = \sum_{n=0}^{\infty}\sum_{m=-n}^{n} b_{nm}(k|\mathbf{X}|)Y_{nm}^{*}(\theta,\phi)$, it follows that (16) filters the above representation via (11), i.e.,

$$\begin{aligned}
&\int_{S^2} S(kd,\theta,\phi)T(\mathbf{k},\mathbf{X})d^2\hat{k} \\
&= \int_{S^2} d^2\hat{k} \left[\sum_{n=0}^{N_h(1)} \sum_{m=-n}^{n} a_{nm} Y_{nm}(\theta,\phi)\right] \\
&\quad \times \left[\sum_{n=0}^{\infty} \sum_{m=-n}^{n} b_{nm} Y_{nm}^{*}(\theta,\phi)\right] \\
&= \int_{S^2} d^2\hat{k}\, S(kd,\theta,\phi) \cdot \left[\sum_{n=0}^{\cancel{\infty}N_h(1)} \sum_{m=-n}^{n} b_{nm} Y_{nm}^{*}(\theta,\phi)\right].
\end{aligned} \tag{17}$$

The filter due to $S(kd,\theta,\phi)$ follows from the fact that if higher order terms ($n > N_h(1)$) are present, they integrate to zero analytically. In other words, $S(kd,\theta,\phi)$ effectively band-limits the spectrum of the incoming wave. As a result, the integration rules (and the number of harmonics of the translation operator) are chosen to exactly integrate polynomials of order $2N_h(1)$ present in (11). If polynomials of order higher than $2N_h(1)$ were present, then while they should theoretically integrate to zero, they would not do so numerically if the integration rule is not chosen appropriately.

Given the bandwidth of $S(kd,\theta,\phi)$, the number of quadrature points/samples required to completely represent $S(kd,\theta,\phi)$, i.e, obtain all coefficients a_{nm}, is $(N_h(1)+1)\times(2N_h(1)+1)$. To evaluate the integral (11) numerically, one should be able to recover b_{nm} from samples of $T(\mathbf{k},\mathbf{X})$, and vice-versa. An integration rule to be used can be defined using $N_h(1) = \lceil \chi_s k D(1) \rceil$, $N_\theta(1) = N_h(1)+1$, and $N_\phi(1) = 2\ N_h(1)+1$, where $D(1)$ is the diameter of a sphere that encloses the leaf box, and $\chi_s \geq 1$ is a factor controlling the accuracy of the approximation of g.

In a multilevel scenario, the same logic holds at every level. That is, both the contributions from that level and those from its parents should be representable in terms of spherical harmonics up to a chosen order. This implies that one needs to develop an anterpolation operator to effect a filter of the incoming spectra to enable a transition from parent to child, or in terms of representation, reduce the maximum degree of spherical harmonics used from $N_h(i+1) \longrightarrow N_h(i)$. To control errors arising from this transition, it is necessary to understand the spectral properties of the anterpolation operator used; local algebraic inter-/anterpolation operators do not accomplish the necessary filtering, and errors accrued due to lack of a proper filter are exacerbated in deep trees. For this reason, we will use bandlimited anterpolants. As discussed in Section IV, using such global operations poses bottlenecks in parallelization.

Thus far, we have established the need for filters to perform interpolation on outgoing expansions ($V_l(\theta,\phi)$) and more importantly, anterpolation on incoming expansions $U_l(\theta,\phi)$. In order to treat both expeditiously, we denote both using $\Lambda_l(\theta,\phi)$. The two variations that we have used are spherical harmonics and Fourier transforms. Spherical filters as used in [21] and [33] result in optimal sampling of the far-field data. An alternative approach is to use Fourier transforms in both angular directions [29], requiring uniformly spaced samples. Obviously, the tradeoff is $\mathcal{O}(N_h(l)^3)$ memory required to effect the Legendre transforms versus doubling of the sample count required for the latter. The approach we espouse is to choose a transition level at which the representation switches from spherical to uniform sampling.

a) Spherical filters: Spherical filters are based on projecting an incoming or outgoing field pattern Λ_l onto spherical harmonics as

$$\Lambda_l(\theta,\phi) = \sum_{n=0}^{N_h(l)} \sum_{m=-n}^{n} \lambda_{l,nm} Y_{nm}(\theta,\phi) \tag{18}$$

and then performing the various down- and up-sampling operations. We shall not delve into details, as these are well-known and obtainable from [33].

b) Fourier filters: The concept of using Fourier transforms was first introduced into the FMM literature by Sarvas [29]. The key to this approach was the realization that the angular domain of the associated Legendre function can be extended from $[0,\pi] \rightarrow [0,2\pi]$ [34]. To summarize, $\Lambda_l(\theta,\phi)$ can be extended to the entire sphere as

$$\tilde{\Lambda}_l(\theta,\phi) = \begin{cases} \Lambda(\theta,\phi) & (\theta,\phi) \in [0,\pi]\times[0,2\pi] \\ \Lambda(2\pi-\theta,\phi+\pi) & (\theta,\phi) \in [\pi,2\pi]\times[0,\pi] \\ \Lambda(2\pi-\theta,\phi-\pi) & (\theta,\phi) \in [\pi,2\pi]^2. \end{cases} \tag{19}$$

If the spectrum of $\Lambda_l(\theta,\phi)$ is bandlimited to $N_h(l)$ harmonics, then, it can be written as

$$\tilde{\Lambda}_l(\theta,\phi) = \sum_{n=-N_h(l)}^{N_h(l)} \sum_{m=-N_h(l)}^{N_h(l)} \tilde{\lambda}_{l,nm} e^{-jn\theta} e^{-jm\phi}. \tag{20}$$

It follows that one needs to sample at $(2\ N_h(l)+1)\times(2\ N_h(l)+1)$ uniformly on the sphere to recover the coefficients $\tilde{\lambda}_{l,nm}$. Using the above representation, we can write (13) as

$$\Psi^{FF}(\mathbf{r}) \approx -\frac{jk}{32\pi^2} \int_0^{2\pi} d\theta |\sin\theta| \int_0^{2\pi} d\phi e^{-j\mathbf{k}\cdot\mathbf{d}_o(1)} \tilde{\Lambda}_1(\theta,\phi). \tag{21}$$

In (21), $|\sin\theta|$ is not bandlimited. But as discussed earlier, $\tilde{\Lambda}_1$ filters the rest of the integrand such that one can represent

$|\sin\theta|$ in terms of a finite Fourier series. Specifically,

$$|\sin\theta| \approx \tilde{s}(\theta) = \sum_{n=-N_h(1)}^{N_h(1)} s_n e^{-jn\theta}. \tag{22}$$

As a result, it follows that the highest order of harmonics in θ is $2\,N_h(1)$, and the integration rules must be designed such that one can evaluate (21) using $(4\,N_h(1)+1)\times(2\,N_h(1)+1)$. This is done using uniform sampling in both θ and ϕ directions; reduction of this number by about a factor of two can be accomplished by leveraging symmetries [29]. Furthermore, we note that as is the usual practice, we can embed $|\sin\theta|$ within the translation operator; we can also filter this expression [35]. As a result, one needs to evaluate an integral of the form

$$\Psi^{FF}(\mathbf{r}) \approx -\frac{jk}{32\pi^2}\int_0^{2\pi} d\theta \int_0^{2\pi} d\phi e^{-j\mathbf{k}\cdot\mathbf{d}_o(1)}\tilde{\Upsilon}_1(\theta,\phi) \tag{23}$$

where $\tilde{\Upsilon}_1(\theta,\phi) = \tilde{s}(\theta)T(\mathbf{k},\mathbf{X})$. As before, these ideas can be trivially extended to a multilevel setting.

c) Spherical to uniform and vice-versa: Both spherical and uniform interpolation have their drawbacks in terms of costs. Transitioning from one method to the other yields a computationally efficient scheme. To facilitate this, in what follows, we present the means to do so. Consider the transition from spherical to uniform. The procedure is straightforward, i.e.,

$$\Lambda_l(\theta,\phi) \xrightarrow{sph} \lambda_{l,nm} \xrightarrow{uni} \tilde{\Lambda}_{l+1}(\theta,\phi) \tag{24}$$

where the stacked symbols indicate the sampling regimes and embed a shift operation for upward traversal. Equation (24) indicates that one obtains from the spherical samples of $\Lambda_l(\theta,\phi)$ coefficients $\lambda_{l,nm}$ and then uses these to construct the requisite samples at uniform points in the $[0,2\pi]\times[0,2\pi]$ grid, and then shift these from the center of the child box to that of its parent. Note that, one can exploit symmetry to store field samples only in the $[0,\pi]\times[0,2\pi]$ grid by using an even number of samples in ϕ.

Next, to effect the transition from uniform to nonuniform sampling, we consider (23)

$$\begin{aligned}\Psi^{FF}(\mathbf{r}) &\approx -\frac{jk}{32\pi^2}\int_0^{2\pi} d\theta \int_0^{2\pi} d\phi e^{-j\mathbf{k}\cdot\mathbf{d}_o(1)}\tilde{\Upsilon}_1(\theta,\phi)\\ &= -\frac{jk}{(4\pi)^2}\int_0^{\pi} \sin\theta d\theta \int_0^{2\pi} d\phi e^{-j\mathbf{k}\cdot\mathbf{d}_o(1)}\frac{\tilde{\Upsilon}_1(\theta,\phi)}{\sin\theta}.\end{aligned} \tag{25}$$

As is evident from (25), the integration has been reduced back to the unit sphere. However, the integration is challenging as $\tilde{\Upsilon}_1(\theta,\phi)/\sin\theta$ is not bandlimited in terms of spherical harmonics. To effect the integral, we have to resort to the notion that the the effective bandwidth of $e^{-j\mathbf{k}\cdot\mathbf{d}_o}$ filters the rest of the integrand. That is,

$$\frac{1}{\sin\theta}\tilde{\Upsilon}_1(\theta,\phi) = \sum_{n=0}^{N_h(1)}\sum_{m=-n}^{n} v_{1,nm}Y_{nm}(\theta,\phi). \tag{26}$$

To obtain $v_{1,nm}$ we evaluate

$$v_{1,nm} = \int_0^{\pi} d\theta \int_0^{2\pi} d\phi \tilde{\Upsilon}_1(\theta,\phi)Y^*_{nm}(\theta,\phi). \tag{27}$$

Equation (27) can be evaluated exactly by defining an extended spherical harmonic as

$$\hat{Y}_{nm}(\theta,\phi) = \begin{cases} Y_{nm}(\theta,\phi) & (\theta,\phi)\in[0,\pi]\times[0,2\pi]\\ Y_{nm}(2\pi-\theta,\phi+\pi) & (\theta,\phi)\in[\pi,2\pi]\times[0,\pi]\\ Y_{nm}(2\pi-\theta,\phi-\pi) & (\theta,\phi)\in[\pi,2\pi]^2. \end{cases} \tag{28}$$

The periodic extension of the normalized spherical harmonics can be represented in terms of a Fourier series, and as a result, the integral

$$v_{1,nm} = \frac{1}{2}\int_0^{2\pi} d\theta \int_0^{2\pi} d\phi \tilde{\Upsilon}_1(\theta,\phi)\hat{Y}^*_{nm}(\theta,\phi) \tag{29}$$

can evaluated exactly using uniform samples of $\tilde{\Upsilon}_1(\theta,\phi)$ that are available to us with a trapezoid rule in both dimensions.

2) Subtleties for Vector Fields: The algorithm of this paper uses four trees: one each for component of the magnetic vector potential and one for the scalar potential. This is the point form of MLFMA that addresses integration issues [36]. Alternatively, one can develop MLFMA with either two or three trees. The two tree version is the usual dyadic MLFMA and requires either vector spherical harmonics [21] or the Fourier method modified to account for the antisymmetry of the transverse fields [37], whereas the three tree version requires mapping onto the three components of electric fields, E_x, E_y, E_z, and employing the scalar filters espoused here. Multiplication of fields of the form (18) by the dyad $\hat{\theta}\hat{\theta}+\hat{\phi}\hat{\phi}$ increases the polynomial order by two, requiring a commensurate increase in the size of the integration rule.

D. Cross-Level Interactions

Cross-level interactions, e.g., interactions between boxes of different sizes, occur as a consequence of a nonuniform tree. The addition theorems used to derive either MLFMA or ACE are generally not valid for such interactions. To circumvent this issue, we subsume either $\mathbf{d}_o$ or $\mathbf{d}_s$, depending on the interaction, into the translation heading $\mathbf{X}$. If the source box Ω_s is smaller than the observer box Ω_o, i.e., $|\Omega_s| < |\Omega_o|$, the source multipole expansion is valid everywhere within Ω_o, implying $\mathbf{X} \leftarrow \mathbf{X}+\mathbf{d}_o$. The source multipole expansion is consequently translated from its expansion center $\mathbf{r}_s^c$ directly to each particle within Ω_o. Such interactions are tabulated in the X-list. Conversely, if $|\Omega_s| > |\Omega_o|$, we take $\mathbf{X} \leftarrow \mathbf{X}+\mathbf{d}_s$, translating each particle within Ω_s to the center $\mathbf{r}_o^c$ of Ω_o, adding to the local expansion. These interactions are tabulated in the W-list. Both of these forms of the addition theorem are always valid for well-separated interactions. The essential criterion is that the two boxes are separated by at least one box of the same size as the smaller interacting box.

The interaction "type" is determined by the field representation (ACE or MLFMA) of the smaller of the two boxes.

Consider evaluation of (9) between source and observation points in Ω_s and Ω_o, respectively. Let $\mathbf{M}^{(n)}$ be the ACE multipole tensor representing the sources within Ω_s centered at $\mathbf{r}_s^c$. Then, we may write

$$\Psi(\mathbf{r}) = \begin{cases} \displaystyle\sum_{n=0}^{\infty} \mathbf{M}^{(n)} \cdot n \cdot \nabla^n g\left(\mathbf{r} - \mathbf{r}_s^c\right), & |\Omega_o| > |\Omega_s| \\ \displaystyle\sum_{n=0}^{\infty} \left(\mathbf{r} - \mathbf{r}_o^c\right)^n \cdot n \cdot \mathbf{L}^{(n)}, & |\Omega_o| < |\Omega_s| \end{cases} \tag{30}$$

where

$$\mathbf{L}^{(n)} = \sum_{i=1}^{N} \frac{u_i}{n!} \nabla^n g\left(\mathbf{r}_o^c - \mathbf{r}_i\right). \tag{31}$$

In the MLFMA, we have

$$\Psi(\mathbf{r}) = \begin{cases} \displaystyle\int_{S^2} T\left(\mathbf{k}, \mathbf{r} - \mathbf{r}_s^c\right) V\left(\mathbf{k}, \mathbf{r}_s^c\right) d^2\hat{k}, & |\Omega_o| > |\Omega_s| \\ \displaystyle\int_{S^2} e^{-j\mathbf{k}\cdot\left(\mathbf{r} - \mathbf{r}_o^c\right)} U\left(\mathbf{k}, \mathbf{r}_o^c\right) d^2\hat{k}, & |\Omega_o| < |\Omega_s| \end{cases} \tag{32}$$

where $V(\mathbf{k}, \mathbf{r}_s^c)$ is the source multipole expansion, and

$$U\left(\mathbf{k}, \mathbf{r}_o^c\right) = \sum_{i=1}^{N} u_i T\left(\mathbf{k}, \mathbf{r}_o^c - \mathbf{r}_i\right) \tag{33}$$

represents the local expansion within Ω_o of the sources. As with the interaction type, the band limit for the translation operator T is determined by the smaller of the two boxes.

It should be noted that the ACE translation operator is valid only over electrically small distances [15], so care should be taken to ensure that translations from a box at the coarsest ACE level to points within MLFMA boxes fall within this radius. For this purpose, 2:1 balancing of the octree [38] at this level suffices.

E. Interpolation of MLFMA Translation Operators

To reduce the storage overhead associated with translation operators, interpolation in $\hat{k} \cdot \hat{X}$ may be used [39]. However, this scheme is untenable for large numbers of X- and W-list interactions, as it requires the number of unique translation distances to be small. To ameliorate this cost, we can also interpolate in the argument of the spherical Hankel function to further reduce the storage without significantly reducing the accuracy of the MLFMA. It is well-known that we may pull out a phase–magnitude term independent of n, obtaining $h_n^{(2)}(z) = (e^{-jz}/z)\tilde{h}_n^{(2)}(z)$. It is then common to all elements of the series T, and we may implicitly define a smoother, phase- and amplitude-compensated translation operator $\tilde{T}$

$$T(\mathbf{k}, \mathbf{X}) = \frac{e^{-jkX}}{kX} \tilde{T}(\mathbf{k}, \mathbf{X}). \tag{34}$$

We employ the Lagrange interpolating polynomials $\ell_n(x)$ of order p. For a given level, let kX_{min}, kX_{max} denote the minimum and maximum translation distances, respectively. The interpolation nodes are chosen to give a high density near kX_{min} and a lower density near kX_{max} because the function to be interpolated goes as $\mathcal{O}((kX)^{-m}), m = 0, \ldots, N_h(l) + 1$. The nodes are chosen as follows. Let $a = kX_{min} - \Delta_1$ and $b = kX_{max} + \Delta_2$ be the beginning and end of the sampling interval, respectively, where Δ_1, Δ_2 are small extensions of the interval, and J be the number of kX values to interpolate over. Then, for $j = 1, \ldots, J$, we let

$$x_j = a + (b - a)\left[1 - \cos\left(\frac{(j-1)\pi}{2(J-1)}\right)\right] \tag{35}$$

denote the jth interpolation node. Let p denote the order of interpolation polynomials used. The interval extensions are chosen as

$$\begin{aligned} \Delta_1 &= p(kX_{max} - kX_{min})\left[1 - \cos\left(\frac{\pi}{2(J-1)}\right)\right] \\ \Delta_2 &= p(kX_{max} - kX_{min}) \cos\left(\frac{(J-2)\pi}{2(J-1)}\right) \end{aligned} \tag{36}$$

so that the extensions are the size of only p subintervals at the beginning and end. This prevents extension of the interpolation interval into the explosion zone of the spherical Hankel functions.

We store samples of the compensated translator $\tilde{T}$ for level l in a matrix $\mathbf{T}_l$ defined as

$$[\mathbf{T}_l]_{mn} \doteq \tilde{T}\left(\cos\beta_m, \frac{x_n}{k}\right) \tag{37}$$

where $\{\beta_m\}$ is a set of uniformly spaced samples on $[-q\,\Delta_\beta, \pi + q\,\Delta_\beta]$, with interpolation order q and spacing Δ_β. Then, for each translation distance kX in the interaction list at level l, we precompute the interpolation weights $\{w_j = \ell_j(kX)\}, j = 0, \ldots, p$ for reconstructing the compensated translator $\tilde{T}$ from the columns of $\mathbf{T}_l$ as

$$\tilde{T}(\cos\beta_m, kX) \approx \sum_{j=0}^{p} w_j [\mathbf{T}_l]_{m(Q+j)} \tag{38}$$

where Q is the index of the first-occurring sample x_Q used in the interpolation, determined by binary search. During the evaluation phase, (38) is evaluated to obtain the samples for the interpolation in $\hat{k} \cdot \hat{X}$. The result of this interpolation is then multiplied by the phase–magnitude factor e^{-jkX}/kX to complete the process. The complexity is, therefore, increased only by a constant factor of $(p + 1) \sim \mathcal{O}(1)$. Table I summarizes the reconstruction error over a range of translation distances normalized by box size using $2N_h(l)$ samples and $p = 3$. The average error is reduced by a factor of two when switching from uniform sampling to cosine sampling, and another factor of two is gained by phase–amplitude compensation. The maximum error at any point within the interval is reduced significantly by the approach described here.

TABLE I

COMPARISON OF L_2 RELATIVE ERROR DATA FOR TRANSLATION OPERATOR INTERPOLATION WITH AND WITHOUT COMPENSATION

Sampling	ϵ_{max}	ϵ_{min}	ϵ_{avg}	$kX/D(l)$ at ϵ_{max}
Uniform	7.38e-3	0	1.43e-3	1.81411
Cosine	2.35e-3	2.39e-7	4.16e-4	3.52377
Uniform+comp.	3.28e-3	0	5.59e-4	1.81411
Cosine+comp.	2.05e-4	1.10e-7	7.50e-5	2.3004

IV. Summary of Parallel Wideband MLFMA Algorithm

In this section, we provide a brief outline of the parallel ACE-AP-MLFMA algorithm used to obtain the numerical results presented in Section V (more details will be provided in [30]). The parallel algorithm we use is based on that presented in [19] with a number of important improvements. We would like to point out that this algorithm is effectively a bottom-up tree construction/partitioning algorithm, and compared to the top-down HiP schemes used in [17] and [40], it gives us important advantages in terms of data partitioning and load balancing (LB) for nonuniform trees as discussed in the following.

A. Tree Construction

Our bottom-up tree construction starts by assigning points to leaf boxes according to a prescribed leaf level box size (essentially creating a uniform tree to begin with). Using a parallel bucket-sort algorithm, these leaf boxes are then partitioned into contiguous chunks with the objective of balancing the number of points assigned to different processes. After this initial partitioning, each process constructs the upper levels of its own tree based on its leaf boxes. This scheme will obviously result in sparsely populated leaf nodes given the nonuniform particle distributions in our target problems. Therefore, starting at the leaf level and working up the tree, each process merges its sparsely populated tree nodes. To be precise, if the total number of particles owned by the children of an internal tree node is below a predetermined threshold s_{max}, then those leaf nodes are merged at the parent node, making the parent a new leaf node. Since the octree is distributed across processes, such merge operations are performed recursively at the higher levels in coordination with neighboring processes. Interaction lists are subsequently computed on the merged tree.

B. Load Balancing

Even distribution of the computational work is crucial to achieve scaling to a large number of processes. Two major issues make LB for the MLFMA algorithm challenging: 1) electrically large objects result in deep trees that contain a significant amount of work across a small number of high-level tree nodes and 2) the nonuniform octree structure has a work profile with high variability.

We address the first issue by inheriting the adaptive direction partitioning idea, first presented in [19]. Note that, the bottom-up partitioning approach described above results in overlapped tree regions between processes, i.e., internal tree nodes at process boundaries may appear in the partial trees of multiple processes. We call such nodes as duplicate nodes, and to balance the load in deep trees, we distribute the multipole expansion data for these nodes as well as the computations (inter/anterpolations and translations) associated with them evenly across all duplicating processes. As discussed in [30], this paper significantly improves the implementation of the basic direction partitioning idea by leveraging better parallelism, nonblocking collective communication primitives, and an efficient distributed execution schedule.

The LB issue associated with nonuniform trees is addressed using an empirical cost evaluation technique. Specifically, each process computes the interaction lists for nonduplicate nodes within its own partial tree and evaluates how much time it takes to perform a sample set of kernel operations (inter/anterpolations and translations) on the given hardware. It then assigns a computational load to each tree node, starting from the highest level by taking into account the empirical kernel costs and the number of interactions each tree node has. These costs are then percolated all the way down to the leaf nodes. Finally, the nonuniform MLFMA tree is repartitioned such that the empirical costs of the contiguous chunks of leaf nodes assigned to each process are evenly distributed.

Algorithm 1 Parallel MLFMA matvec

1) **Near-field *matvec***
2) **Charge/Multipole-to-multipole (C2M/M2M)**
   - Compute multipole expansions for unique interior nodes in post-order sequence.
   - Initiate asynchronous reduce-scatters for direction partitioned duplicate nodes
3) **Parallel update and interpolation**
   - Complete reduce-scatters
   - Interpolate/shift duplicate node expansions using coarse-grained parallelism at ACE & spherically sampled levels and fine-grained parallelism at uniformly sampled levels
4) **Multipole-to-local (M2L)**:
   - Initiate asynchronous communications for multipole expansions of source/observer node pairs located at different processes
   - Apply M2L operations for local source/observer nodes located on same process
   - Complete communications for non-local source multipole expansion data & perform M2L
5) **Parallel update and anterpolation**
   - Gather complete local expansion for each duplicate node from all its users
   - Shift/anterpolate duplicate node expansions using fine-grained parallelism at uniformly sampled levels and coarse-grained parallelism at ACE & spherically sampled levels
6) **Local-to-local/observer (L2L/L2O)**
   - Shift and anterpolate local expansions for unique interior nodes in pre-order sequence
   - Compute fields from local expansions

C. Parallel Evaluation

Algorithm 1 provides a broad overview of the parallel matvec. Parallelism is applied in three different forms in the M2M phase. Each process first performs the serial computations for their own unique interior nodes, which are nodes that strictly belong to a single process. Second, duplicate nodes in the ACE and spherically sampled levels are handled using coarse-grained parallelism, where all processes sharing a duplicate node compute the interpolations for the children they own sequentially and reduce or reduce-scatter, for ACE and MLFMA, respectively, the results among all sharing processes

using nonblocking message passing interface (MPI) collectives. As the third form of parallelism, duplicate nodes above the spherical-to-uniform transition level are interpolated using a fine-grained parallel version of the FFT-based interpolation algorithm. The partial multipole data for each duplicate node in this category is shifted, summed, and partitioned over the duplicating processes in parallel.

Translations in the M2L phase follow a similar strategy as above, where all processes begin with sequentially computing translations within their own trees, then exchange multipole data with neighboring processes and process them as they arrive—effectively overlapping computation and communication. Due to the presence of cross-level interactions in nonuniform trees, these computations are organized into three subphases corresponding to X, V, and W lists. We note that one important shortcoming of the present parallel M2L implementation is that all translation data for duplicate nodes need to be communicated to the resident process, i.e., the designated receiving process for each duplicate node, which then applies them one by one—this can potentially create performance bottlenecks for high-level nodes in deep trees.

Finally, the L2L phase is performed by essentially mimicking M2M in the reverse order. As an artifact of the limited parallelization of duplicate nodes mentioned above, L2L phase starts with broadcasting the local expansions of duplicate nodes from the resident process to the duplicating processes. After this initial communication, anterpolation and disaggregation of the duplicate nodes above the spherical-to-uniform transition level are performed in distributed fashion. Once all parallel anterpolations are completed, all processes can independently perform the sequential L2L operations for their internal nodes.

V. Numerical Results

A. Accuracy Analysis

In this section, we demonstrate the controllable accuracy of each interpolation/anterpolation method discussed in this paper, namely, spherical, uniform, and hybrid. We also discuss tradeoffs between these methods in terms of memory and runtime. The examples in this section were run using a single core on a desktop computer with an Intel Xeon E5-2630 CPU with a clock speed of 2.3 GHz and 64 GB RAM. All computations use double precision arithmetic, and FFTW 3.3.5 is used for performing FFTs in the interpolation and anterpolation stages.

1) Hybrid Sampling MLFMA Error Control: To illustrate the accuracy of the proposed hybrid sampling and interpolation/anterpolation method, we consider a uniformly random distribution of 256 000 randomly oriented dipoles of unit strength within an $8\lambda \times 8\lambda \times 8\lambda$ cubical domain. The box size is chosen to be 0.25λ, resulting in a six-level uniform octree. Neither ACE nor interpolation of the MLFMA translation operator is used in this experiment. For the upper levels of the tree, a cap was enforced on the bandwidth to prevent numerical breakdown of the translation operator. Using three buffer boxes, we study the error convergence of the far-field contributions for each of the spherical, uniform, and hybrid methods outlined in Section IV. For the hybrid scheme, the transition was chosen so that only the bottom two levels use the spherical scheme. These data are shown in Fig. 3(a). As the oversampling parameter χ is increased, each method converges in a nearly identical manner, demonstrating that the hybrid method introduces no error. Fig. 3(b) shows the time taken per matvec for each method. The hybrid scheme only adds a few seconds to the baseline time taken by the purely spherical scheme, while the uniform-only scheme takes roughly twice as long on average, as expected. For contrast, direct evaluation of the far field matvec takes 5128.029 s on the same machine.

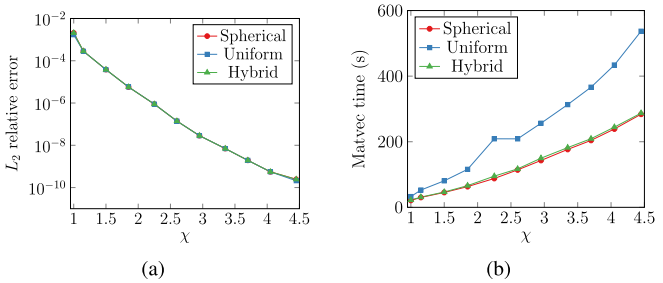

Fig. 3. (a) Error convergence and (b) matvec timings versus oversampling parameter χ for each sampling/interpolation method in the 8λ cube geometry.

TABLE II
Memory Consumption and Timings for Different Hybrid Sampling Transitions

L_{su}	3	4	5	6	7
Uni. levels	0	1	2	3	4
Sph. levels	9	8	7	6	5
P_n^m (GB)	6.0	2.0	0.3	0.04	0.005
Multipole (GB)	13.1	14.5	15.9	17.4	18.8
M2L (GB)	0.8	1.4	1.5	1.6	1.6
Matvec time (s)	730	957	1364	1377	1589

2) Tradeoffs for Hybrid Sampling: We have already seen the effect of purely uniform sampling on runtime due to the doubled sampling rate throughout the tree. The results for the hybrid sampling scheme suggest that the extra cost of fully uniform sampling comes principally from performing M2L operations while oversampling at the finest levels of the tree where most boxes exist, as expected. Conversely, the optimal sampling rate for the spherical scheme gives the best runtime, but at the expense of precomputing and storing samples of associated Legendre functions P_n^m to perform interpolation and anterpolation via the spherical harmonics transform. To study the tradeoffs, we examine the case of 256 000 dipoles distributed on a $256\lambda \times 256\lambda$ square in the xy plane. Using leaf-level boxes of diameter $\lambda/4$, the uniform tree is 11 levels deep, and a one-buffer box rule is used. The oversampling rate is fixed at $\chi = 1.0$.

We examine memory consumption and matvec time for different values of the finest spherical sampling level L_{su}. Table II shows the memory required for storing P_n^m samples, multipole and local expansions, and M2L operators, and the serial matvec time versus L_{su}.

While these data are collected using a serial kernel, we must emphasize that in a parallel setting, the multipole and local data, which consume the majority of the overall memory, is distributed across processes in a largely nonredundant

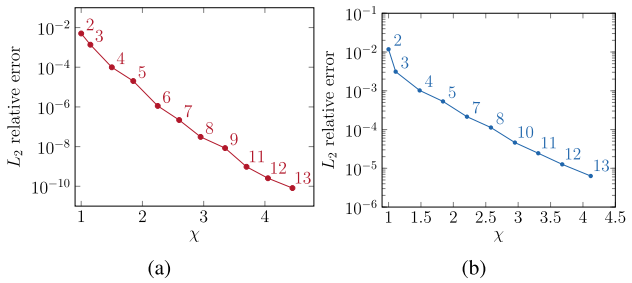

Fig. 4. Error convergence versus oversampling parameter χ for (a) uniform tree ACE-MLFMA for 64 000 dipoles in a flattened box, and (b) nonuniform tree cone-sphere. The number next to each data point indicates the order P of ACE expansions used. Note that, the example in (a) uses three buffer boxes for the far-field, while that in (b) uses only one.

TABLE III
COMPARISON OF UNIFORM AND NONUNIFORM TREES FOR THE CONESPHERE GEOMETRY

	Uniform	Non-uniform
Nodes	1058881	29562
Leaf level	14	8-14
Dipoles/box	1	12
Inter. lists	1149.16 s	25.05 s

fashion, while the P_n^m storage is duplicated on every process. A set of M2L operators specific to each process' partial tree must also be stored. In contemporary high-performance computing environments, the maximum amount of RAM per core is typically around 4 GB, underscoring the importance of finding an acceptable balance between memory consumption and runtime to make most efficient use of computational resources.

3) Uniform-Tree Wideband MLFMA Error Control: We now examine the error in the hybrid-sampling MLFMA with ACE, using a uniform tree. We consider the case of 64 000 dipoles distributed nonuniformly within an $8\lambda \times 8\lambda \times 1/16\lambda$ box. The point locations were generated using $(x, y, z) = (8r_1^{1.5}, 8r_2^{1.5}, r_3/16)$, where r_1, r_2, r_3 are random numbers on $[0, 1]$. The smallest box was chosen as $\lambda/16$, yielding an eight-level tree with two levels of ACE, two levels of spherical MLFMA, and two levels of uniform MLFMA. Fig. 4(a) shows the error convergence of the far field as both χ, P are increased, demonstrating fine-grained control over accuracy over both the low- and mid-frequency regimes.

4) Nonuniform-Tree Wideband MLFMA Error Control: We now examine the error convergence for the nonuniform wideband ACE-MLFMA applied to a multiscale object and compare with the uniform-tree algorithm. Consider a "cone-sphere" surface geometry comprising a long cone terminated on its wide end by a hemispherical surface. The cone portion is 36.33λ long, and the spherical portion has a radius of 4.49λ. The mesh contains 47 367 triangles, and we placed six randomly oriented dipoles within each triangle, totaling 284 202 dipoles. The smallest box size is chosen to be $\lambda/128$, and the tree is merged using $s_{max} = 30$, yielding a 14-level tree with five levels of ACE, three levels of spherical MLFMA, and four levels of uniform MLFMA.

Table III summarizes some of the stark differences between the uniform and nonuniform trees. Notably, the precomputation time required for computing interaction lists decreases by a factor of over 45 as a consequence of the 97% reduction in overall tree nodes. Fig. 4(b) shows convergence in the relative L_2 error of the far-field using the nonuniform tree with a single buffer box rule.

B. Parallel Kernel Evaluation

Having demonstrated the accuracy of the proposed numerical methods, we turn now to the demonstration of the parallel algorithm on electrically large, multiscale objects. In this section, we focus on the parallel N-body evaluation of the Helmholtz potential Ψ, given by

$$\Psi(\mathbf{r}_m) = \sum_{n=1}^{N} g(\mathbf{r}_m - \mathbf{r}_n)u_n, \quad m = 1, \ldots, N \tag{39}$$

where $\mathbf{r}_n$ is the location of the nth source/observer point, and $u_n \equiv 1$ is the source strength. Results in this section and in Section V-C were run on the Cori supercomputer at the National Energy Research Scientific Computing Center, Berkeley, CA, USA. Each computational node contains two sockets, each populated with a 16-core Intel Xeon E5-2698 v3 ("Haswell") with a clock speed of 2.3 GHz, and 128 GB DDR4 RAM. FFTW 3.3.4.11 was used for all FFTs. Again, all computations employ double precision arithmetic. Both the N-body and solver codes are parallelized using MPI with no shared memory parallelism. We now consider an aircraft-shaped distribution comprising 175 764 666 points on the surface, which fits into a bounding box of dimensions $693.5\lambda \times 200.2\lambda \times 754.8\lambda$. The minimum leaf box diameter was set to $\lambda/32$ resulting in a 16-level tree, and the densest box contains 56 points. The four coarsest levels of computation use uniform sampling, and the remaining seven MLFMA levels use spherical sampling. ACE is used for boxes of diameter smaller than $\lambda/4$. The nonuniform tree is merged with $s_{max} = 40$, resulting in leaves containing 20 points on average distributed over the bottom five levels of the tree, and an almost 20-fold reduction in leaf boxes. For accuracy on the order of 10^{-3}, we set $\chi = 1.0$, $P = 3$ [15]. This supposition is backed up by randomly selecting a point on each process to ensure good spatial distribution, computing the exact observed field, and comparing with the computed value; the average relative error at these observers is 3.98×10^{-3}.

With this setup, we examine the tradeoffs for nonuniform trees versus uniform trees and the beneficial effects of LB for parallel matvecs. Fig. 5(a) shows the time taken per process for the far-field matvec using 1024 processes using a uniform tree and a nonuniform tree. The time drops from 61 s on average for the uniform tree to 38 s for the nonuniform tree. Load balancing is employed in both instances. Much of the speedup is achieved in the translation stages. The flatness of portions of the M2L timings for the uniform tree are caused by synchronization barriers. Fig. 5(b) details the effectiveness of LB, showing per-process timings for the tree traversal (M2M+L2L) and translation (V, W,and$X-$lists) both with and without LB with exactly the same parameters. Without

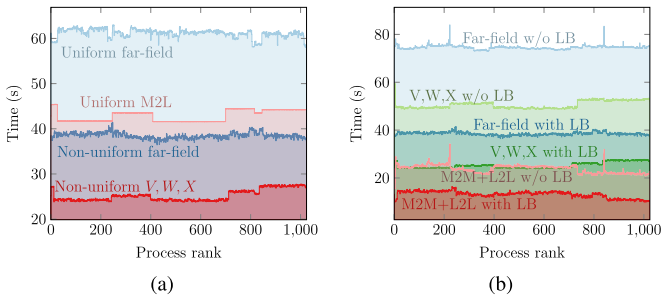

Fig. 5. Per-process timings for far-field matvec stages for the airplane geometry (a) with uniform versus nonuniform tree and (b) with and without LB for nonuniform tree on 1024 processes.

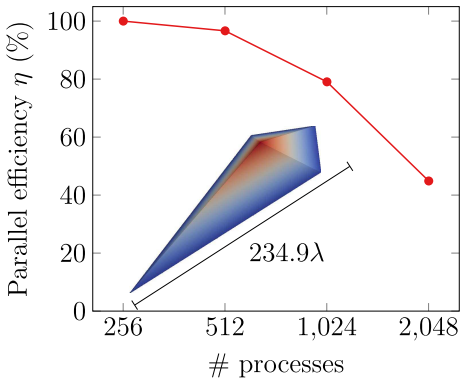

Fig. 6. Parallel efficiency of the complete matvec for the 286M-point arrow geometry with reference to 256 processes. Colorization of geometry solely to illustrate depth.

LB, the translation stage alone takes 10 s longer than the entire matvec with LB.

Next, we consider a collection of 286 312 650 points distributed uniformly on a surface geometry with an arrow-like shape, yielding high variation in point density over the surface. The bounding box for the arrow measures $234.9\lambda \times 77\lambda \times 39.32\lambda$, and we choose the minimum box size to be $\lambda/64$, resulting in a 15-level tree. As always, ACE is used for boxes smaller than $\lambda/4$. The three coarsest MLFMA levels of computation employ uniform sampling, while the rest employ spherical sampling. The tree is merged using $s_{max} = 25$, resulting in nine points per leaf box on average.

A common metric for measuring the scalability of a parallel algorithm is the parallel efficiency η, defined by

$$\eta(t, N_p)(\%) \doteq \frac{t_{ref} N_{p,ref}}{t N_p} \times 100\% \tag{40}$$

where t is the time taken using N_p processes and t_{ref} is the time taken for the reference run using $N_{p,ref}$ processes. The parallel efficiency of the complete matvec (both near- and far-field together) is given in Fig. 6 with reference to 256 processes. Memory limitations prevented us from running the code under the same memory-per-core conditions beneath 256 processes. Up to 1024 processes, the matvec exhibits good scaling, but drops precipitously for 2048 processes. This appears to be associated with (as Amdahl's law suggests) communication steps, particularly within the parallel interpolation/anterpolations, and decreasing work to processor ratio.

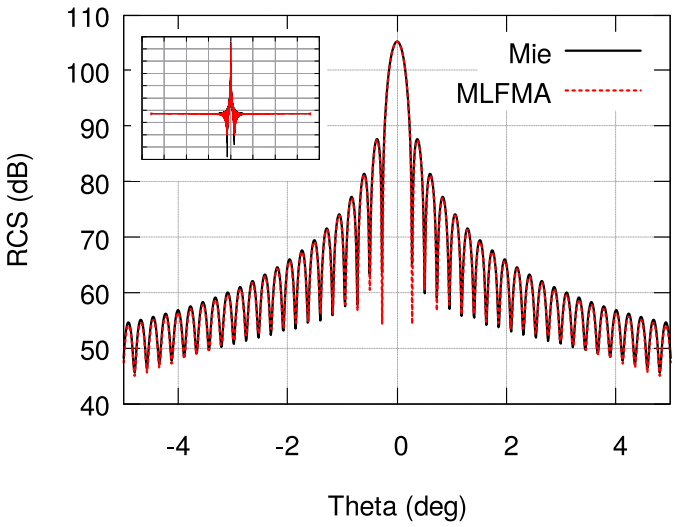

Fig. 7. Comparison of bistatic RCS calculated by parallel MLFMA versus analytical for a sphere of diameter 256λ within $\pm 5°$ of the main lobe.

C. Parallel CFIE Solver Evaluation

In this section, the kernel evaluation is wrapped in a parallel MoM solver for the CFIE (1), and we present results demonstrating its capabilities, all with $\alpha = 0.5$. Parallel GMRES is used as the iterative solver. We note that our use of a four-tree mixed-potential MLFMA as opposed to a two- or three-tree dyadic MLFMA directly results in an increase in runtime.

1) Uniform Tree: First, we consider the evaluation of the radar cross section (RCS) of an electrically large conducting sphere of diameter 256λ, discretized with 84 934 656 unknowns. A seven-point integration rule within each patch was used for both source and testing far-field integrals. The leaf-level box is chosen as 0.25λ, resulting in an 11-level tree. The bottom six levels employ spherical sampling before transitioning to uniform sampling, and $\chi = 1.0$. The incident plane wave is polarized along the $\hat{x}$ direction and traveling in the $+z$ direction. Parallel GMRES with tolerance 10^{-3} and a restart value of 30 is used to solve the resulting matrix system, converging within two outer iterations. The solve took 31 min using 2048 processes, requiring 27.8 s per matvec. The calculated RCS is shown in Fig. 7 along with the analytical Mie series solution, with which the agreement is very good.

2) Nonuniform Trees: Finally, we consider the evaluation of RCS from geometries with nonuniform distributions using nonuniform trees.

First, we consider electrically large objects with sharp corners or other features with regions of high density of unknowns. Our first example is an arrowhead-shaped geometry measuring $470\lambda \times 154\lambda \times 79\lambda$, discretized using 77 257 728 unknowns. A 14-level nonuniform tree is used with minimum leaf box size of 0.0574λ and $s_{max} = 35$, resulting in 10 unknowns per leaf, on average. The RCS is shown in Fig. 8, along with an illustration of the geometry. The

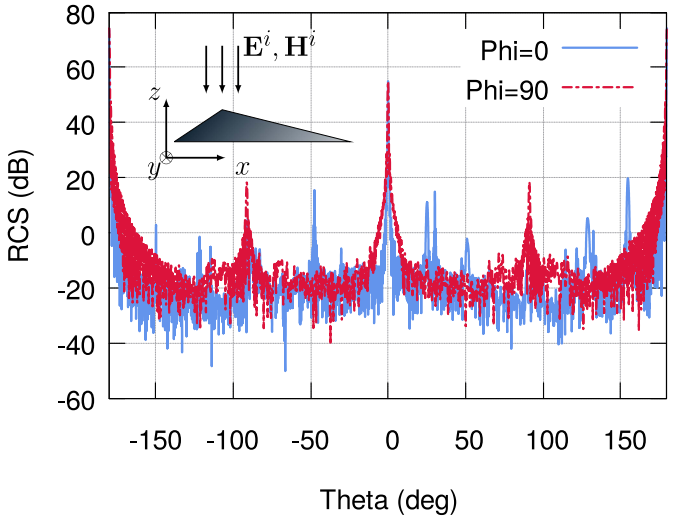

Fig. 8. RCS calculated with parallel MLFMA for 470λ arrowhead geometry.

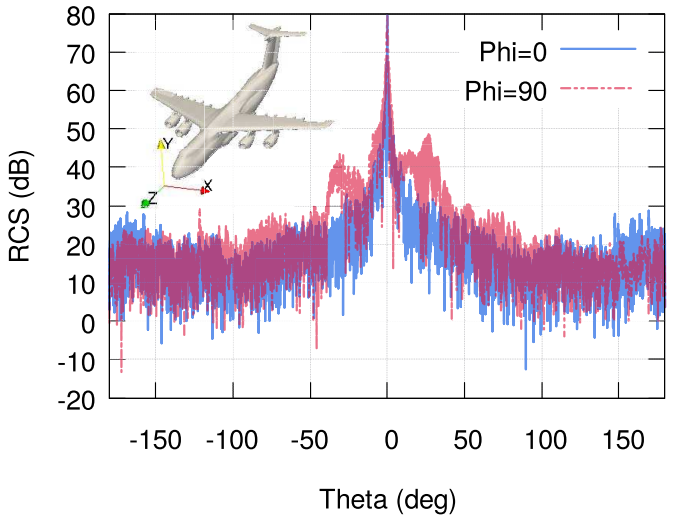

Fig. 9. RCS calculated with parallel MLFMA for 755λ airplane geometry.

incident electric field is given by $\mathbf{E}^i = \hat{x}e^{jkz}$. Solution with tolerance 5×10^{-3} took 19.3 min on 2048 processes, at 23.6 s per matvec. Our last example concerns scattering from the 755λ airplane geometry depicted in the inset of Fig. 9. A $\lambda/10$ discretization of the surface yields 285 425 664 unknowns. The incident field is given as $\mathbf{E}^i = \hat{x}e^{-jkz}$. Using a minimum leaf size of 0.0231λ, the nonuniform tree spans 16 levels and is merged with $s_{max} = 35$, increasing the average number of unknowns per box from 1 to 18 and distributing leaves across the bottom four tree levels. ACE ($P = 3$) is used for boxes smaller than $\lambda/4$, and uniform MLFMA is used at the three uppermost levels. The densest box at the finest level contains 40 unknowns. Using GMRES with a tolerance of 5×10^{-3}, the solution on 2048 processes took 3.17 h, using 1 min 8 s per matvec. The RCS in the $\phi = 0°, 90°$ cuts is shown in Fig. 9.

VI. Conclusion

In this paper, we have presented an extremely wide-band parallel MLFMA with fine-grained control over the error in field evaluation and methods for reducing both storage and computational cost while sacrificing nothing in terms of accuracy. We introduced rigorous operators for computing interactions at any level on a nonuniform octree which is adapted in parallel to fit to the geometry. We also introduced a rigorous method for transitioning from spherical harmonics-based to Fourier-based inter/anterpolation of multipole expansions to optimize storage and solution time. An array of numerical examples demonstrates the accuracy and efficiency of the algorithm, and several scattering examples demonstrate the ability of the solver to accurately solve problems on large, complicated objects with nonuniform spatial distributions of unknowns.

Acknowledgment

This research used resources of the National Energy Research Scientific Computing Center, Berkeley, CA, USA, a DOE Office of Science User Facility supported by the Office of Science of the U.S. Department of Energy under Contract DE-AC02-05CH11231.

References

[1] A. F. Peterson, S. L. Ray, and R. Mittra, *Computational Methods for Electromagnetics*. New York, NY, USA: IEEE Press, 1998.

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

[3] J.-M. Jin, *Theory and Computation of Electromagnetic Fields*. Hoboken, NJ, USA: Wiley, 2011.

[4] M. Vikram and B. Shanker, "An incomplete review of fast multipole methods-from static to wideband-as applied to problems in computational electromagnetics," *Appl. Comput. Electromagn. Soc. J.*, vol. 24, no. 2, pp. 79–108, 2009.

[5] N. Nishimura, "Fast multipole accelerated boundary integral equation methods," *Appl. Mech. Rev.*, vol. 55, no. 4, pp. 299–324, 2002.

[6] L. Greengard and V. Rokhlin, "A new version of the fast multipole method for the Laplace equation in three dimensions," *Acta Numer.*, vol. 6, pp. 229–269, Jan. 1997.

[7] L. Greengard, J. Huang, V. Rokhlin, and S. Wandzura, "Accelerating fast multipole methods for the Helmholtz equation at low frequencies," *IEEE Comput. Sci. Eng.*, vol. 5, no. 3, pp. 32–38, Jul./Sep. 1998.

[8] J.-S. Zhao and W. C. Chew, "Three-dimensional multilevel fast multipole algorithm from static to electrodynamic," *Microw. Opt. Technol. Lett.*, vol. 26, no. 1, pp. 43–48, Jul. 2000.

[9] H. Cheng *et al.*, "A wideband fast multipole method for the Helmholtz equation in three dimensions," *J. Comput. Phys.*, vol. 216, no. 1, pp. 300–325, Jul. 2006.

[10] T. Dufva and J. Sarvas, "Broadband MLFMA with plane wave expansions and optimal memory demand," *IEEE Trans. Antennas Propag.*, vol. 57, no. 3, pp. 742–753, Mar. 2009.

[11] I. Bogaert, J. Peeters, J. Fostier, and F. Olyslager, "NSPWMLFMA: A low frequency stable formulation of the MLFMA in three dimensions," in *Proc. IEEE Antennas Propag. Soc. Int. Symp.*, Jul. 2008, pp. 1–4.

[12] J. Aronsson, I. Jeffrey, and V. Okhmatovski, "Generalization of the Barnes–Hut algorithm for the Helmholtz equation in three dimensions," *IEEE Antennas Wireless Propag. Lett.*, vol. 8, pp. 425–428, 2009.

[13] M. Takrimi, Ö. Ergül, and V. B. Ertürk, "A novel broadband multilevel fast multipole algorithm with incomplete-leaf tree structures for multiscale electromagnetic problems," *IEEE Trans. Antennas Propag.*, vol. 64, no. 6, pp. 2445–2456, Jun. 2016.

[14] B. Shanker and H. Huang, "Accelerated Cartesian expansions—A fast method for computing of potentials of the form R^{-v} for all real v," *J. Comput. Phys.*, vol. 226, no. 1, pp. 732–753, 2007.

[15] M. Vikram, H. Huang, B. Shanker, and T. Van, "A novel wideband FMM for fast integral equation solution of multiscale problems in electromagnetics," *IEEE Trans. Antennas Propag.*, vol. 57, no. 7, pp. 2094–2104, Jul. 2009.

[16] S. Velamparambil, W. C. Chew, and J. Song, "10 million unknowns: Is it that big? [computational electromagnetics]," *IEEE Antennas Propag. Mag.*, vol. 45, no. 2, pp. 43–58, Apr. 2003.

[17] Ö. Ergül and L. Gurel, "Efficient parallelization of the multilevel fast multipole algorithm for the solution of large-scale scattering problems," *IEEE Trans. Antennas Propag.*, vol. 56, no. 8, pp. 2335–2345, Aug. 2008.
[18] B. Michiels, J. Fostier, I. Bogaert, and D. De Zutter, "Weak scalability analysis of the distributed-memory parallel MLFMA," *IEEE Trans. Antennas Propag.*, vol. 61, no. 11, pp. 5567–5574, Nov. 2013.
[19] V. Melapudi, B. Shanker, S. Seal, and S. Aluru, "A scalable parallel wideband MLFMA for efficient electromagnetic simulations on large scale clusters," *IEEE Trans. Antennas Propag.*, vol. 59, no. 7, pp. 2565–2577, Jul. 2011.
[20] M. A. Epton and B. Dembart, "Multipole translation theory for the three-dimensional Laplace and Helmholtz equations," *SIAM J. Sci. Comput.*, vol. 16, no. 4, pp. 865–897, 1995.
[21] B. Shanker, A. A. Ergin, M. Lu, and E. Michielssen, "Fast analysis of transient electromagnetic scattering phenomena using the multilevel plane wave time domain algorithm," *IEEE Trans. Antennas Propag.*, vol. 51, no. 3, pp. 628–641, Mar. 2003.
[22] J. Carrier, L. Greengard, and V. Rokhlin, "A fast adaptive multipole algorithm for particle simulations," *SIAM J. Sci. Stat. Comput.*, vol. 9, no. 4, pp. 669–686, 1988.
[23] A. R. Benson, J. Poulson, K. Tran, B. Engquist, and L. Ying, "A parallel directional fast multipole method," *SIAM J. Sci. Comput.*, vol. 36, no. 4, pp. C335–C352, 2014.
[24] S. M. Rao, D. R. Wilton, and A. W. Glisson, "Electromagnetic scattering by surfaces of arbitrary shape," *IEEE Trans. Antennas Propag.*, vol. AP-30, no. 3, pp. 409–418, May 1982.
[25] J. M. Song and W. C. Chew, "Multilevel fast-multipole algorithm for solving combined field integral equations of electromagnetic scattering," *Microw. Opt. Technol. Lett.*, vol. 10, no. 1, pp. 14–19, Sep. 1995.
[26] M. Lu and E. Michielssen, "A local filtering scheme for FMM/PWTD algorithms," in *Proc. IEEE Antennas Propag. Soc. Symp.*, vol. 4, Jun. 2004, pp. 4523–4526.
[27] B. Michiels, "Parallel fast multipole methods for the simulation of extremely large electromagnetic scattering problems," Ph.D. dissertation, Dept. Inf. Technol., Fac. Eng. Archit., Ghent Univ., Ghent, Belgium, 2013.
[28] A. Capozzoli, C. Curcio, A. Liseno, and A. Riccardi, "Parameter selection and accuracy in type-3 non-uniform FFTs based on Gaussian gridding," *Prog. Electromagn. Res.*, vol. 142, no. 6, pp. 743–770, 2013.
[29] J. Sarvas, "Performing interpolation and anterpolation entirely by fast Fourier transform in the 3-D multilevel fast multipole algorithm," *SIAM J. Numer. Anal.*, vol. 41, no. 6, pp. 2180–2196, 2003.
[30] S. M. Hughey, "Efficient parallelization of non-uniform fast multipole algorithms," Ph.D. dissertation, Dept. Elect. Comput. Eng., Michigan State Univ., East Lansing, MI, USA, 2018.
[31] R. Coifman, V. Rokhlin, and S. Wandzura, "The fast multipole method for the wave equation: A pedestrian prescription," *IEEE Antennas Propag. Mag.*, vol. 35, no. 3, pp. 7–12, Jun. 1993.
[32] J. M. Song, C.-C. Lu, and W. C. Chew, "Multilevel fast multipole algorithm for electromagnetic scattering by large complex objects," *IEEE Trans. Antennas Propag.*, vol. 45, no. 10, pp. 1488–1493, Oct. 1997.
[33] R. Jakob-Chien and B. K. Alpert, "A fast spherical filter with uniform resolution," *J. Comput. Phys.*, vol. 136, no. 2, pp. 580–584, 1997.
[34] S. Y. K. Yee, "Studies on Fourier series on spheres," *Monthly Weather Rev.*, vol. 108, no. 5, pp. 676–678, 1980.
[35] C. Cecka and E. Darve, "Fourier-based fast multipole method for the Helmholtz equation," *SIAM J. Sci. Comput.*, vol. 35, no. 1, pp. A79–A103, 2013.
[36] D. Dault and B. Shanker, "A mixed potential MLFMA for higher order moment methods with application to the generalized method of moments," *IEEE Trans. Antennas Propag.*, vol. 64, no. 2, pp. 650–662, Feb. 2016.
[37] J. Peeters, "Efficient simulation of 3D electromagnetic scattering problems using boundary integral equations," Ph.D. dissertation, Dept. Inf. Technol., Fac. Eng. Archit., Ghent Univ., Ghent, Belgium, 2010.
[38] H. Sundar, R. S. Sampath, and G. Biros, "Bottom-up construction and 2:1 balance refinement of linear octrees in parallel," *SIAM J. Sci. Comput.*, vol. 30, no. 5, pp. 2675–2708, 2008.
[39] J. Song and W. C. Chew, "Interpolation of translation matrix in MLFMA," *Microw. Opt. Technol. Lett.*, vol. 30, no. 2, pp. 109–114, 2001.
[40] B. Michiels, J. Fostier, I. Bogaert, and D. De Zutter, "Full-wave simulations of electromagnetic scattering problems with billions of unknowns," *IEEE Trans. Antennas Propag.*, vol. 63, no. 2, pp. 796–799, Feb. 2015.

Stephen Hughey (S'12) received the B.S. degree in electrical engineering and the Ph.D. degree in electrical engineering and computational science from Michigan State University, East Lansing, MI, USA, in 2012 and 2018, respectively.

Since 2018, he has been with the BerrieHill Research Division, Applied Research Associates, Lansing, MI, USA. His current research interests include integral equation methods and fast algorithms in electromagnetics, finite-element methods, and parallel computing/algorithm design.

H. M. Aktulga received the B.S. degree in computer science from Bilkent University, Ankara, Turkey, in 2004, and the M.S. and Ph.D. degrees in computer science from Purdue University, West Lafayette, IN, USA, in 2009 and 2010, respectively.

He was a Post-Doctoral Researcher with the Scientific Computing Group, Lawrence Berkeley National Laboratory, Berkeley, CA, USA. He joined Michigan State University, East Lansing, MI, USA, in 2014. His current research interests include high-performance computing (HPC), HPC applications, big data analytics and (sparse) numerical linear algebra, the design and development of parallel algorithms, numerical methods, performance models, and software systems that can harness the full potential of state-of-the-art computing platforms to address challenging problems in large-scale scientific computing and big-data analytics. A distinguishing aspect of his research is the close collaborations that he has built with domain experts in a wide range of fields such as molecular modeling and simulation, materials science, computational electromagnetics, and nuclear physics.

Melapudi Vikram received the B.Tech. degree in ME from IIT Madras, Chennai, India, in 2003, and the Ph.D. degree in electrical and computer engineering from Michigan State University, East Lansing, MI, USA, in 2009.

He was with the Electrical and Computer Engineering Department, Michigan State University, from 2009 to 2010, and with Ansys (ANSOFT), Pittsburgh, PA, USA, from 2010 to 2012. He is currently a Lead Research Engineer with the GE Global Research Center, Bengaluru, India. He has authored over 40 referred publications in international journals and conferences. He holds five patents. His current research interests include the development of sensing and imaging methods for industrial and medical applications that includes device design, signal/image processing algorithms, the use of advanced machine learning methods, multiphysics modeling, and fast and parallel algorithms.

Dr. Vikram was a recipient of the Best Student Paper Award (3rd prize) at the 2008 IEEE AP-S International Symposium and USNC/URSI National Radio Science Meeting, San Diego, CA, USA.

Mingyu Lu (M'03–SM'08) received the B.S. and M.S. degrees in electrical engineering from Tsinghua University, Beijing, China, in 1995 and 1997, respectively, and the Ph.D. degree in electrical engineering from the University of Illinois at Urbana–Champaign, Champaign, IL, USA, in 2002.

From 2002 to 2005, he was a Post-Doctoral Research Associate with the Electromagnetics Laboratory, University of Illinois at Urbana–Champaign. From 2005 to 2012, he was an Assistant Professor with the Department of Electrical Engineering, University of Texas at Arlington, Arlington, TX, USA. In 2012, he joined the Department of Electrical and Computer Engineering, West Virginia University Institute of Technology, Beckley, WV, USA, where he is currently an Associate Professor. His current research interests include wireless power transmission, radar systems, microwave remote sensing, antenna design, and computational electromagnetics.

Dr. Lu was a recipient of the First Prize Award in the student paper competition of the IEEE AP-S International Symposium, Boston, MA, USA, in 2001, and the Outstanding Service Award from the IEEE Fort Worth Chapter in 2008. He served as the Chair of the Antennas and Propagation Society of the IEEE Fort Worth Chapter from 2006 to 2011.

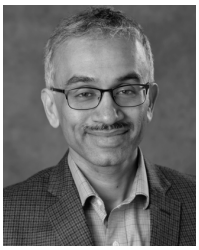

Balasubramaniam Shanker (SM'02–F'10) received the B.Tech. degree from IIT Madras, Chennai, India, in 1989, and the M.S. and Ph.D. degrees from Pennsylvania State University, State College, PA, USA, in 1992 and 1993, respectively.

From 1993 to 1996, he was a Research Associate with the Department of Biochemistry and Biophysics, Iowa State University, Ames, IA, USA, where he focused on the molecular theory of optical activity. From 1996 to 1999, he was a Visiting Assistant Professor with the Center for Computational Electromagnetics, University of Illinois at Urbana–Champaign, Champaign, IL, USA. From 1999 to 2002, he was an Assistant Professor with the Department of Electrical and Computer Engineering, Iowa State University. From 2012 to 2015, he was the Associate Chair of graduate studies with the Department of Electrical and Computer Engineering, Michigan State University, East Lansing, MI, USA, where he was the Associate Chair of the Department of Computational Mathematics, Science and Engineering, from 2015 to 2018, and is currently a University Distinguished Professor with the Department of Electrical and Computer Engineering and also with the Department of Physics and Astronomy. He has authored or co-authored around 400 journal and conference papers and presented a number of invited talks. His current research interests include all aspects of computational electromagnetics (frequency- and time-domain integral equation-based methods, multiscale fast multipole methods, fast transient methods, and higher order finite-element and integral equation methods), propagation in complex media, mesoscale electromagnetics, and particle and molecular dynamics as applied to multiphysics and multiscale problems.

Dr. Shanker was a recipient of the Withrow Distinguished Junior Scholar in 2003, the Withrow Distinguished Senior Scholar in 2010, the Withrow Teaching Award in 2007, and the Beal Outstanding Faculty Award in 2014. He was an Associate Editor of IEEE ANTENNAS AND WIRELESS PROPAGATION LETTERS and the IEEE TRANSACTIONS ON ANTENNAS AND PROPAGATION. He is the Topical Editor of the *Journal of Optical Society of America: A*, and is a Full Member of the USNC-URSI Commission B.

Eric Michielssen (SM'99–F'02) received the M.S. degree *(summa cum laude)* in electrical engineering from Katholieke Universiteit Leuven, Leuven, Belgium, in 1987, and the Ph.D. degree in electrical engineering from the University of Illinois at Urbana–Champaign (UIUC), Champaign, IL, USA, in 1992.

From 1992 to 2005, he was a Faculty Member with UIUC. In 2005, he joined the University of Michigan, Ann Arbor, MI, USA, where he is currently the Louise Ganiard Johnson Professor of engineering and a Professor of electrical engineering and computer science. He also serves as the Associate Vice President of advanced research computing and the Co-Director of precision health initiative with the University of Michigan. He has authored or co-authored over 180 journal papers and book chapters and over 350 papers in conference proceedings. His current research interests include all aspects of theoretical and applied computational electromagnetics, the development of fast frequency and time domain integral-equation-based techniques for analyzing electromagnetic phenomena, and the development of robust optimizers for the synthesis of electromagnetic/optical devices.

Dr. Michielssen is a member of the URSI Commission B. He was a recipient of the Belgian American Educational Foundation Fellowship in 1988, the Schlumberger Fellowship in 1990, the 1994 International Union of Radio Scientists (URSI) Young Scientist Fellowship, the 1995 National Science Foundation CAREER Award, and the 1998 Applied Computational Electromagnetics Society Valued Service Award, the UIUC's 2001 Xerox Award for Faculty Research, the UM College of Engineering David E. Liddle Research Excellence Award in 2011, the IEEE AP-S Chen-To-Tai Distinguished Educator Award in 2014, and the 2018 IEEE AP-S S.A. Schelkunoff Transactions Prize Paper Award. In addition, he was named the 1999 URSI United States National Committee Henry G. Booker Fellow and selected as the recipient of the 1999 URSI Koga Gold Medal. He was appointed the 2002 Beckman Fellow in the UIUC Center for Advanced Studies, named the 2003 Scholar in the Tel Aviv University Sackler Center for Advanced Studies, and selected as the UIUC 2003 University and Sony Scholar.