A FREE-SPACE ADAPTIVE FMM-BASED PDE SOLVER IN THREE DIMENSIONS

M. HARPER LANGSTON,
LESLIE GREENGARD AND DENIS ZORIN

msp

A FREE-SPACE ADAPTIVE FMM-BASED PDE SOLVER IN THREE DIMENSIONS

M. HARPER LANGSTON, LESLIE GREENGARD AND DENIS ZORIN

We present a kernel-independent, adaptive fast multipole method (FMM) of arbitrary order accuracy for solving elliptic PDEs in three dimensions with radiation and periodic boundary conditions. The algorithm requires only the ability to evaluate the Green's function for the governing equation and a representation of the source distribution (the right-hand side) that can be evaluated at arbitrary points. The performance is accelerated in three ways. First, we construct a piecewise polynomial approximation of the right-hand side and compute far-field expansions in the FMM from the coefficients of this approximation. Second, we precompute tables of quadratures to handle the near-field interactions on adaptive octree data structures, keeping the total storage requirements in check through the exploitation of symmetries. Third, we employ shared-memory parallelization methods and load-balancing techniques to accelerate the major algorithmic loops of the FMM. We present numerical examples for the Laplace, modified Helmholtz and Stokes equations.

1. Introduction

Many problems in scientific computing call for the efficient solution to linear partial differential equations with constant coefficients. On regular grids with separable Dirichlet, Neumann or periodic boundary conditions, such equations can be solved using fast, direct methods. For free-space boundary conditions and highly nonuniform source distributions defined on adaptive and/or unstructured grids, alternative approaches are necessary. We describe a *direct* high-order adaptive solver for inhomogeneous linear constant-coefficient PDEs in three dimensions with decay conditions at infinity. A typical case is the Poisson equation

$$-\Delta u = g, \quad \text{supp}(g) \subset \Omega, \tag{1}$$

where Ω is a bounded domain in $\mathbb{R}^3$, and $u(\boldsymbol{x}) = O(1/|\boldsymbol{x}|)$ as $|\boldsymbol{x}|$ goes to infinity. Our solver uses a kernel-independent fast multipole method (FMM) [60; 61] which can be applied to any PDE for which a free-space Green's function evaluation routine is provided. It handles highly nonuniform sources in an efficient manner, using an adaptive approximation of the right-hand side in (1). The structure of the solver allows for natural integration with FMM-based boundary integral equation techniques, leading to the construction of an adaptive kernel-independent solver for inhomogeneous PDEs in complex geometries, to be described elsewhere.

Related work. For regular grids in separable coordinate systems (rectangles, disks, spheres, etc.), fast methods for constant-coefficient second order PDEs are well established [15; 16]. These methods generally rely on cyclic reduction and/or fast Fourier transforms (FFTs) to achieve nearly linear scaling. For many problems, however, adaptive meshes resulting from adaptive mesh refinement (AMR) strategies are essential [2; 7; 49], and existing solvers typically rely on domain decomposition strategies [23] or multigrid acceleration [18; 37; 40; 44]. For complex geometries, unstructured grid generation techniques are often used [45]. In such cases, both the grid generation process and the solution of the resulting linear systems can be computationally expensive. The lack of regularity in the data structures adds complexities in parallelization as well [1; 18].

A more recent class of methods combines ideas from potential theory with finite difference methods. In [39], fast direct solvers were used on a sequence of refined grids with boundary conditions inherited from the coarser levels. This results in discontinuities at coarse-fine interfaces, which are corrected using a second pass through the grid hierarchy. In [4], the method of local corrections (MLC) [3] was combined with multigrid methods to solve the Poisson equation on a hierarchy of nested grids. The fastest free-space Poisson solver for three-dimensional problems of which we are aware is described in [47]. It first solves local Poisson problems on fine grids using FFT-based techniques and then couples together the solutions on coarser grids using MLC. This approach was shown to be very effective in parallel, with good scaling up to 1024 processors. (A similar two-dimensional scheme is described in [29].) For unstructured meshes, the preceding methods do not apply without significant modification and most fast solvers are based on iterative methods using multigrid or domain decomposition acceleration [13; 14; 17].

Here we concentrate on the integral equation (or, more precisely, the integral transform) viewpoint. Rather than solving (1), for example, we simply compute

$$u(\boldsymbol{x}) = \frac{1}{4\pi} \int_{\mathbb{R}^3} \frac{1}{|\boldsymbol{x} - \boldsymbol{y}|} g(\boldsymbol{y}) \, d\boldsymbol{y}. \tag{2}$$

Among the advantages of this approach is the increase in precision in computing derivatives. In PDE-based methods, if first or second derivatives of the solution are

needed, accuracy tends to degrade due to the need for numerical differentiation. Instead, we can differentiate the kernel in (2) and compute derivatives from their integral representation as well. Other advantages are that free-space radiation conditions are automatically satisfied, we can obtain simple *a priori* error estimates, and high order accuracy is straightforward to achieve. However, the computational complexity of a naïve implementation is high: computing the solution u at N points x given N discretization points y requires $O(N^2)$ work. There have been a number of methods proposed to overcome this barrier. These include panel-clustering techniques [11; 36], hierarchical matrices ($\mathcal{H}, \mathcal{H}^2$-matrices) [10; 34; 35], the Barnes–Hut method [5], and the fast multipole method (FMM) [20; 32; 25; 51] originally designed for gravitational/Coulomb interactions. These schemes all achieve linear $O(N)$ or nearly linear $O(N \log N)$ scaling. Most of these methods fall into the class of what are often called *tree codes* because they separate near- and far-field interactions on a hierarchy of spatial scales using quadtree (2D) or octree (3D) data structures. Because it can achieve arbitrary precision at modest cost with straightforward error estimates, we concentrate on the FMM in the present setting. The classical FMM is kernel-specific and relies on detailed separation of variables solutions of the governing PDE. While the FMM references above considered the Laplace equation, the Helmholtz equation was subsequently treated in [52]. A three-dimensional version effective for all frequencies (and additional references) can be found in [21]. The modified Helmholtz equation was discussed in [12; 27], and the biharmonic equation in [28; 33; 58]. The Stokes equations are somewhat exceptional, since they can be handled by a sequence of calls to the original (Coulomb) FMM [56; 50]. An attractive alternative that avoids much of the detailed analytic work of these methods is the kernel-independent approach of [60; 61]. In this approach, expansions in special functions are replaced with equivalent source densities. The result is that the same numerical apparatus can be used for a variety of PDEs, and the user need only supply a subroutine for the evaluation of the relevant Green's function.

While the bulk of the work on FMMs over the last two decades has concentrated on particle interactions or the acceleration of boundary integral equation methods, there has been some work on solving inhomogeneous PDEs. One option is to couple the FMM with finite difference methods to allow for fast solvers in complex geometries [46; 48; 9]. While this is a significant improvement in terms of range of applicability over classical fast solvers, these methods require a regular volume mesh on which is superimposed an irregular boundary. Adaptive FMMs for volume source distributions in two dimensions were described in [22; 24; 29]. The present paper extends these two-dimensional schemes to three dimensions, incorporates them into kernel-independent FMMs, and introduces several new performance optimizations. The result is an efficient, adaptive method that is capable of computing volume

integrals in three dimensions for a broad variety of PDE kernels.

Before turning to the method itself, we should also note that there has been a significant body of work in the quantum chemistry community on accelerating volume integral calculations using the FMM, where collections of Gaussians are typically used to describe the charge distribution [53; 59]. These are Poisson problems in free-space but with a different approach to defining the right-hand side.

2. Equations and kernels

Given a linear, constant-coefficient PDE

$$\mathcal{L}(u)(\boldsymbol{x}) = g(\boldsymbol{x}), \tag{3}$$

classical mathematical methods can be used to compute the corresponding Green's function $K(\boldsymbol{x}, \boldsymbol{y})$ in free space, such that

$$u(\boldsymbol{x}) = \int_{\Omega} K(\boldsymbol{x}, \boldsymbol{y}) g(\boldsymbol{y}) \, d\boldsymbol{y}, \tag{4}$$

where Ω is the support of g. $K(\boldsymbol{x}, \boldsymbol{y})$ is, in general, weakly singular; assuming $g(\boldsymbol{x})$ is given at N points and $u(\boldsymbol{x})$ is desired at N points, the nonlocal character of the integral representation, as indicated above, would lead to an $O(N^2)$ solution procedure. Thus, we need both a suitable quadrature approach and a fast algorithm for (4) to yield a useful numerical technique. Assuming this is achieved, a number of advantages follow. First, no linear system needs to be solved. Second, adaptivity is achieved through the approximation of the right-hand side. Third, as mentioned in the previous section, derivatives can be computed without loss of precision. (There is some loss in accuracy for derivatives of order greater than two, since at that point the integral operator becomes hypersingular and some catastrophic cancellation cannot be avoided.) Finally, we have simple a priori error estimates. To see this, let $\hat{g}(\boldsymbol{x})$ be the approximation to $g(\boldsymbol{x})$, and let $\hat{Q}[f](\boldsymbol{x})$ denote the quadrature approximation of $\int_{\Omega} K(\boldsymbol{x}, \boldsymbol{y}) f(\boldsymbol{y}) \, d\boldsymbol{y}$. Assuming the near field is computed exactly, the quadrature error satisfies an estimate of the form

$$\left| \hat{Q}[f](\boldsymbol{x}) - \int_{\Omega} K(\boldsymbol{x}, \boldsymbol{y}) f(\boldsymbol{y}) \, d\boldsymbol{y} \right| \leq \epsilon \|f\|_1,$$

where ϵ is the approximation error in the FMM. In turn, ϵ is controlled by the parameter p that determines the number of discretization points used for equivalent densities, as described in Section 4.1 and [60].

To estimate the total error, let us assume $\hat{g}(\boldsymbol{x})$ is a k-th order polynomial approximation of the right-hand side,

$$\hat{g}(\boldsymbol{x}) - g(\boldsymbol{x}) \leq \delta = O(h^k),$$

and that

$$\hat{u}(\boldsymbol{x}) = \hat{\boldsymbol{Q}}[\hat{g}](\boldsymbol{x}).$$

(5)

Then

$$
\begin{aligned}
e(\boldsymbol{x}) = u(\boldsymbol{x}) - \hat{u}(\boldsymbol{x}) &= \int_\Omega K(\boldsymbol{x}, \boldsymbol{y}) g(\boldsymbol{y}) d\boldsymbol{y} - \int_\Omega \hat{\boldsymbol{Q}}[\hat{g}](\boldsymbol{x}) \\
&\leq \int_\Omega K(\boldsymbol{x}, \boldsymbol{y})[g(\boldsymbol{y}) - \hat{g}(\boldsymbol{y})] d\boldsymbol{y} + \left| \int_\Omega K(\boldsymbol{x}, \boldsymbol{y}) \hat{g}(\boldsymbol{y}) d\boldsymbol{y} - \hat{\boldsymbol{Q}}[\hat{g}](\boldsymbol{x}) \right| \\
&\leq C_1 \| g(\boldsymbol{y}) - \hat{g}(\boldsymbol{y}) \|_\infty + \| \hat{g}(\boldsymbol{y}) \|_1 \, \epsilon,
\end{aligned}
$$

(6)

where

$$C_1 = \max_{\boldsymbol{x}} \int_\Omega |K(\boldsymbol{x}, \boldsymbol{y})| \, d\boldsymbol{y} \leq C_1 \delta + \| \hat{g}(\boldsymbol{y}) \|_1 \, \epsilon.$$

This estimate is much sharper than one typically obtained when discretizing the PDE itself, where the order of accuracy is determined by high derivatives of the solution. Here, it depends only on the quality of the approximation of the right-hand side ($\delta = O(h^k)$) and the FMM tolerance (ϵ). Note that the constant C_1 is a bounded quantity determined by the volume of Ω with no dependence on the data. If ϵ is chosen to be of the same order as δ, the scheme is formally k-th order accurate. In practice, it is convenient to decouple the right-hand side approximation error from the FMM tolerance, as above, permitting the user to control them independently.

The principal drawback with the integral formulation is that, when implemented naïvely, the complexity of the approach is quadratic in the number of sample points. FMM algorithms overcome this computational barrier by making systematic use of the smoothness of distant interactions on a hierarchy of spatial scales [6; 24; 30]. The kernel-independent versions of the FMM [60; 61] are particularly useful because of their generality; they make it possible to compute solutions of the form (4) for any (nonoscillatory) elliptic PDE, provided only a module which evaluates the kernel.

After describing the details of the approach, we demonstrate its performance for the Poisson equation

$$-\Delta u(\boldsymbol{x}) = g(\boldsymbol{x}),$$

(7)

the modified Helmholtz equation

$$\alpha u(\boldsymbol{x}) - \Delta u(\boldsymbol{x}) = g(\boldsymbol{x}), \quad \alpha > 0,$$

(8)

and the Stokes equations

$$
\begin{aligned}
\nabla p(\boldsymbol{x}) - \mu \Delta \boldsymbol{u}(\boldsymbol{x}) &= \boldsymbol{g}(\boldsymbol{x}), \quad \mu > 0, \\
\nabla \cdot \boldsymbol{u}(\boldsymbol{x}) &= 0.
\end{aligned}
$$

(9)

Defining $r = x - y$ and $r = \|r\|$, the corresponding kernels in three dimensions are given, respectively, by

$$K(x, y) = \frac{1}{4\pi r}, \tag{10}$$

$$K(x, y) = \frac{1}{4\pi r}e^{-\sqrt{\alpha}r}, \tag{11}$$

$$K(x, y) = \frac{1}{8\pi\mu}\left(\frac{1}{r}I + \frac{r \otimes r}{r^3}\right). \tag{12}$$

The classical FMM is reviewed briefly in Section 3, the kernel-independent method is described in Section 4, and symmetries for optimizing storage are discussed in Section 5. Numerical experiments are presented in Section 6 as well as a brief discussion on extending our method to periodic boundary conditions or including a singular source component along with a smooth background force. We provide additional error analysis in the Appendix as well as a brief summary of how the method is optimized using OpenMP and load-balancing techniques to achieve near-linear strong scaling.

3. Analytic fast multipole method

We briefly review the structure of the original two-dimensional FMM for the case of particle interactions [30]. Given a set of N_{src} charges of strength $g(y_i)$ at locations $(y)_i$, the FMM was designed to compute the induced potentials u_j at N_{trg} target locations, x_j,

$$u_j = u(x_j) = \sum_{i=1}^{N_{\text{src}}} K(x_j, y_i)g(y_i), \quad j = 1, \ldots, N_{\text{trg}}, \tag{13}$$

where $K(x, y) = -\log|\{x - y\}|/2\pi$. For $N_{\text{src}} \approx N_{\text{trg}} = N$, the FMM decreases the computational cost from $O(N^2)$ to $O(N)$ for fixed user-prescribed accuracy by introducing a hierarchical partition (represented by a tree data structure T) of a regular bounding domain D and two series expansions for each box at each level of the hierarchy. More precisely, the root of T is associated with the entire box D and defined to be at *level* $\ell = 0$. Level $\ell + 1$ is obtained from level ℓ recursively, dividing each subdomain at level ℓ into four equal-sized children. For a regular box B of width H, B's *near field*, $\mathcal{N}^B$, is defined as the set of all boxes in D that lie within a box centered at B of width $3H$. The *neighbor list* L_N^B is defined as the set of boxes in $\mathcal{N}^B$ which share a vertex with B. In the nonadaptive case, $L_N^B = \mathcal{N}^B$. The *far field*, $\mathcal{F}^B$, is the complement of the near field: $\mathcal{F}^B = D \setminus \mathcal{N}^B$. Finally, the *interaction list* L_I^B is the set of children of B's parent's neighbors that are not neighbors themselves. Thus, $L_I^B \subseteq \mathcal{F}^B$. The depth of T is chosen so that the smallest boxes (leaves in T) contain no more than some fixed number of points—

say, s. We first consider uniformly refined trees, where all leaves T are at the same level. Note that the total number of boxes in a 2D quadtree is bounded by $4N/3s$ (and $8N/3s$ in a 3D octree). Thus, if the workload per box is constant, the net algorithm has $O(N)$ complexity.

A *local expansion* is used to represent within each box B the influence of all sources in the far field of B. A *multipole expansion* about the center of B is used to represent the influence of sources *inside B* on boxes in the far field $\mathcal{F}^B$ [30].

The FMM computes the total field at a target point in leaf box B as the sum of (a) the field due to the source points contained in the boxes of the neighbor list L_N^B and (b) the contribution from sources in the far field $\mathcal{F}^B$. The contributions from source points inside the boxes of L_N^B are computed directly using (13), while the contributions from $\mathcal{F}^B$ are obtained by evaluating the local expansion of box B at the target. The essential task of the FMM is the *construction* of the local expansions in a hierarchical manner. This takes place in two steps.

The upward pass. This pass begins at the finest level of the tree data structure, converting charge strengths at source points into multipole expansions for each leaf box; this computation is carried out by the *source-to-multipole* (S2M) operator. Multipole expansions for each nonleaf box B at each coarser level are obtained recursively. More precisely, the multipole expansions for the four children of B are merged into a single expansion about B's center using the *multipole-to-multipole* (M2M) operator.

The downward pass. For each box B, starting at the coarsest level, the local expansion of $\mathcal{F}^B$ is obtained by shifting the local expansion of B's parent to the center of B using the *local-to-local* (L2L) operator and by mapping multipole expansions centered at each box in L_I^B to B's local expansion using the *multipole-to-local* (M2L) operator. For leaf box B, local expansions are then evaluated at each target point using the *local to target* (L2T) operator.


Figure 1. Boxes used by M2M, L2L and M2L operators. For box B at level ℓ, P in the L2L operator represents the parent of B at level $\ell - 1$, and in the M2M operator, C represents the children of B at level $\ell + 1$. Boxes labeled V in the M2L operator reside in L_I^B.

Summary. FMM uses S2M, M2M, M2L, L2L, and L2T linear operators: M2M and L2L operators are determined uniquely by the relative position of a box and its parent; each M2L operator is determined by the relative position of a box in the interaction list; S2M and L2T operators depend on source and target point locations and can be different for each (finest level) box. Figure 1 on the previous page illustrates the data flow involved in the M2M, M2L and L2L operators.

For the Laplace kernel in three dimensions, far-field expansions are represented using a mixture of spherical harmonics [31] and plane-wave representations [32].

We turn now to the kernel-independent approach [60; 61] in order to design a volume integral FMM in three dimensions that can handle a broad class of PDEs.

4. 3D kernel-independent FMM volume integral solver

Given an octree T for our 3D bounding domain D, let $D = \sum\{B_i\}$, $i = 1 \dots M$ be the set of leaf boxes resulting from hierarchical subdivision. For a single-layer kernel K, we compute the integral (4) at some point $\boldsymbol{x}$ as

$$u(\boldsymbol{x}) = \sum_{i=1}^{M} K[B_i, g^{B_i}](\boldsymbol{x}), \tag{14}$$

where $K[B, g^B](\boldsymbol{x}) = \int_B K(\boldsymbol{x}, \boldsymbol{y}) g(\boldsymbol{y}) \, d\boldsymbol{y}$, and g^B represents the restriction of the source distribution to the box B.

The principal difference between the approach of this paper and the analytic FMM for point sources is that we use *sampled equivalent densities* instead of classical special functions and series expansions to account for far-field interactions, as in [60; 61]. This requires only a *black-box* kernel evaluation routine and allows for a kernel-independent implementation. A second difference between the approach of this paper and prior kernel-independent FMM schemes is that we are dealing with a continuous source distribution rather than a collection of point-like particles. To extend the method of [60; 61] to this setting, we use polynomial basis functions to approximate the source distribution g on each leaf box, following the two-dimensional approach of [22; 29; 24]. More precisely, we assume that the input source is given on each leaf box B by a polynomial g^B of degree $k + 1$ with coefficients γ^B:

$$g^B = \sum_{j=1}^{N_k} \gamma_j^B \beta_j (2^\ell (\boldsymbol{x} - \boldsymbol{c}_B)), \tag{15}$$

where β_j are polynomial basis functions, ℓ is the depth of the box B ($\ell = 0$ at the root of T), and $\boldsymbol{c}_B$ is its center. We use monomials for low-order accuracy and tensor-product Chebyshev polynomials for higher-order accuracy. The number of coefficients is $N_k = k(k + 1)(k + 2)/6$ for each scalar source function g. We describe an interpolation scheme to convert a set of source values defined on a

grid of sample points to a polynomial representation in Section 4.6. As output, our algorithm can return either point values of the potential at each target point or a polynomial approximation of the potential on each leaf box (which can then be evaluated at arbitrary locations).

To simplify the exposition, we present our algorithm first for a uniformly refined octree of depth ℓ and then discuss the changes necessary for the adaptive octree case separately. The final algorithmic steps are outlined in Section 4.8, and we briefly discuss how the major loops are optimized for shared-memory parallelization in the Appendix (Section A.4).

4.1. *Equivalent densities.* The kernel-independent approach to translation operators is based on the following idea. For kernel K, suppose we have an arbitrary (smooth or nonsmooth) source distribution g_s in a volume Ω_s with surface Γ_s. Let Γ_t denote an auxiliary surface in the exterior of Γ_s, and let Γ_{check} denote yet another auxiliary surface in the exterior of Γ_t. Finally, let E denote the exterior of Γ_{check}. We will compute a charge density ϕ_t on Γ_t such that the potentials $K[\Omega_s, g_s]$ and $K[\Gamma_t, \phi_t]$ coincide in E. This is always possible if the exterior Dirichlet problem on Γ_t has a unique solution and the exterior field can be represented in terms of a single layer potential.[1]

Remark. For some problems, such as the Helmholtz equation, a combination of single and double layer sources may be required because of nonphysical resonances in the single layer representation, but it is generally sufficient for nonoscillatory kernels ([41] for the Poisson equation, [42] for the Stokes equations).

Our goal is to use $K[\Gamma_t, \phi_t]$ to represent the far-field instead of a multipole expansion. For this, we let Γ_{check} approximate the outer boundary of the neighbor list L_N^B and solve a Fredholm integral equation of the first kind for ϕ_t,

$$K[\Gamma_t, \phi_t](\boldsymbol{x}) = K[\Omega_s, g_s](\boldsymbol{x}) \quad \text{for all } \boldsymbol{x} \in \Gamma_{\text{check}}. \tag{16}$$

Having matched the field on Γ_{check}, the fields will match in the exterior E (with precise estimates depending on the specific kernel). We refer to Γ_t as an *equivalent surface* with *equivalent density* ϕ_t, and Γ_{check} as a *check surface*. In the case when the original density is concentrated on the surface Γ_s, then (16) can be written as

$$K[\Gamma_t, \phi_t](\boldsymbol{x}) = K[\Gamma_s, \phi_s](\boldsymbol{x}) \quad \text{for all } \boldsymbol{x} \in \boldsymbol{x}_t. \tag{17}$$

We match the field created by charges *outside* the near neighbors of a box by a discretized layer potential defined on a surface enclosing the box, and a different equivalent density will be used to replace the local expansion. The number of samples used to represent the equivalent density is the analog of the number of

---

[1]For some kernels (Stokes), a low-dimensional nullspace may need to be eliminated.

expansion terms in a classical FMM. Γ_t and Γ_{check} are cubic surfaces, uniformly sampled at p locations. In discretized form, (17) can be written as

$$\boldsymbol{K}^{\Gamma_t, \boldsymbol{x}_t} \phi_t = \boldsymbol{K}^{\Gamma_s, \boldsymbol{x}_t} \phi_s, \tag{18}$$

where ϕ_s and ϕ_t are vectors of point-sampled densities, and $\boldsymbol{K}^{a,b}$ are matrices with entries given by $\boldsymbol{K}_{ij}^{a,b} = K(a_i, b_j)$ for sample points a_i and a_j on surfaces a and b. For known ϕ_s and solving for ϕ_t, (18) is a discretization of a Fredholm equation of the first kind. For large p, linear systems may be poorly conditioned; in such cases, we choose to utilize Tikhonov regularization methods [41] to invert $\boldsymbol{K}^{\Gamma_t, \boldsymbol{x}_t}$. We discuss this approach and its accuracy in the Appendix (sections A.1 and A.3).

Kernel invariance and matrix precomputation. For all equations we consider, the kernels are invariant with respect to a rigid transformation $\mathcal{T}$: for scalar kernels, $K(\mathcal{T}\boldsymbol{x}, \mathcal{T}\boldsymbol{y}) = \mathcal{T}K(\boldsymbol{x}, \boldsymbol{y})$, and for matrix kernels, $K(\mathcal{T}\boldsymbol{x}, \mathcal{T}\boldsymbol{y}) = \mathcal{T}K(\boldsymbol{x}, \boldsymbol{y})\mathcal{T}^T$. Hence, all matrices $\boldsymbol{K}$ need to be computed only once for each class of pairs of equivalent surfaces, closed with respect to a specific $\mathcal{T}$. Furthermore, many kernels are *homogeneous*: for all $c > 0$, there exists a *scaling exponent* $r \neq 0$ such that $K(c\boldsymbol{x}, c\boldsymbol{y}) = c^r K(\boldsymbol{x}, \boldsymbol{y})$, further reducing the number of classes of surface pairs requiring separate matrices. We consider optimizations due to invariance for each translation operator in the next sections, assuming scalar kernels for simplicity, although our implementation can handle matrix kernels.

4.2. Upward pass. For the upward pass, recall now that for each leaf box, the *sources* are polynomials approximating the source distribution. For consistency with the FMM summary above, we use S, M, etc. in describing translation operator names.

Source to multipole S2M *translations.* For each leaf box B, we choose $\boldsymbol{y}^{B,u}$, the *upward equivalent surface*, and $\boldsymbol{x}^{B,u}$, the *upward check surface*, as in [60]. Equation (16) for upward equivalent density $\phi^{B,u}$ in this case becomes

$$K[\boldsymbol{y}^{B,u}, \phi^{B,u}](\boldsymbol{x}) = K[B, g^B](\boldsymbol{x}), \tag{19}$$

$$K[B, g^B](\boldsymbol{x}) \approx \sum_{j=1}^{N_k} \gamma_j^B F_j^B(\boldsymbol{x}), \tag{20}$$

$$\text{where } F_j^B(\boldsymbol{x}) = \int_B \beta_j(2^\ell(\boldsymbol{y} - \boldsymbol{c}_B))K(\boldsymbol{x}, \boldsymbol{y})d\boldsymbol{y} \quad \text{for } \boldsymbol{x} \in \boldsymbol{x}^{B,u}. \tag{21}$$

By translation invariance, $F_j^B(\boldsymbol{x})$ depends only on the choice of β_j and level ℓ of B. To evaluate the integrals in (21), we use adaptive Gaussian quadrature [8]. In matrix form, the Nyström discretization of (19)–(21) at p sample points $\boldsymbol{\phi}^{B,u}$ on $\boldsymbol{y}^{B,u}$ yields

$$\boldsymbol{K}_{\text{S2M}}^B \boldsymbol{\phi}^{B,u} = \boldsymbol{F}_{\text{S2M}}^B \gamma^B, \tag{22}$$

where $\boldsymbol{F}_{\mathrm{S2M}}^{B}$ is the matrix of precomputed weights (21) and $\boldsymbol{K}_{\mathrm{S2M}}^{B}$ is the matrix with entries $K(\boldsymbol{x}_i, \boldsymbol{y}_j)$, $i = 1 \dots p$, $j = 1 \dots p$. Solving for $\boldsymbol{\phi}^{B,u}$,

$$\boldsymbol{\phi}^{B,u} = (\boldsymbol{K}_{\mathrm{S2M}}^{B})^{-1} \boldsymbol{F}_{\mathrm{S2M}}^{B} \gamma^{B} = \boldsymbol{T}_{\mathrm{S2M}}^{B} \gamma^{B}. \tag{23}$$

For a uniformly refined tree, $\boldsymbol{T}_{\mathrm{S2M}}^{B}$ depends only on ℓ, so one matrix is computed. Figure 2, left, illustrates the computation of $\phi^{B,u}$ from γ^{B}.

*Multipole to multipole (*M2M*) translations.* M2M translation operators translate $\phi^{C,u}$ at a child box C to $\phi^{B,u}$ for the parent box B, shown in Figure 2, right: for all $\boldsymbol{x} \in \boldsymbol{x}^{B,u}$,

$$K[\boldsymbol{y}^{B,u}, \phi^{B,u}](\boldsymbol{x}) = \sum_C K[\boldsymbol{y}^{C,u}, \phi^{C,u}](\boldsymbol{x}),$$

or, in matrix form,

$$\boldsymbol{K}_{\mathrm{M2M}}^{B,B} \boldsymbol{\phi}^{B,u} = \sum_C \boldsymbol{K}_{\mathrm{M2M}}^{C,B} \phi^{C,u}. \tag{24}$$

Similar to the S2M computations, these systems are solved as

$$\boldsymbol{\phi}^{B,u} = \sum_C (\boldsymbol{K}_{\mathrm{M2M}}^{B,B})^{-1} \boldsymbol{K}_{\mathrm{M2M}}^{C,B} \phi^{C,u} = \sum_C \boldsymbol{T}_{\mathrm{M2M}}^{C,B} \phi^{C,u}. \tag{25}$$

For any two children C_1 and C_2, rotation R maps C_1 to C_2; therefore, only *one* $\boldsymbol{T}_{\mathrm{M2M}}^{C,B}$ is computed per level, with the contribution to $\boldsymbol{\phi}^{B,u}$ from any other child obtained by composing this matrix with an appropriate permutation of $\phi^{C,u}$. Further, for *homogeneous* kernels, only one matrix is stored at a single level ℓ and scaled as necessary.

4.3. *Downward pass.* In the downward pass *downward equivalent densities* are computed through the M2L, L2L, and L2T operators.


**Figure 2.** Kernel-independent FMM translation operators for S2M (left) and M2M (right).

Figure 3. Kernel-independent FMM translation operators for M2L (left) and L2L (right).

Multipole to local (M2L) translations. For any box B, M2L operators (Figure 3, left) translate $\phi^{V,u}$, approximating the field of sources inside of $V \in L_I^B$, to a *downward equivalent density* $\phi^{B,d}$. In this case, we seek to induce identical potentials *inside* of B, effectively swapping upward equivalent and check surfaces to obtain downward equivalent and check surfaces: $\mathbf{y}^{B,d} = \mathbf{x}^{B,u}$ and $\mathbf{x}^{B,d} = \mathbf{y}^{B,u}$. Equation (17) takes the form

$$K[\mathbf{y}^{B,d}, \phi^{B,d}](\mathbf{x}) = \sum_V K[\mathbf{y}^{V,u}, \phi^{V,u}](\mathbf{x}) \quad \text{for all } \mathbf{x} \in \mathbf{x}^{B,d}, \qquad (26)$$

where $\phi^{B,d}$ is discretized at p uniformly spaced samples on $\mathbf{y}^{B,d}$. The right-hand side of (26) is computed and stored as a *downward check potential*, $u^{B,d}$ at $\mathbf{x}^{B,d}$, and $\boldsymbol{\phi}^{B,d}$ is recovered after the L2L contribution is added.

$$u_{\text{M2L}}^{B,d} = \sum_V \boldsymbol{K}_{\text{M2L}}^{V,B} \boldsymbol{\phi}^{V,u}. \qquad (27)$$

We efficiently evaluate $u^{B,d}$ with FFTs by treating densities as being defined on extensions of $\mathbf{y}^{V,u}$ and $\mathbf{x}^{B,d}$ to 3D Cartesian grids with zero values in the interior. This results in $O(p^3/2)$ sample locations, and the computational cost of $O(p^3/2 \log p)$ for evaluation.

There are at most 189 possible locations for $V \in L_I^B$ relative to any particular B; however, using translation and rotation invariance of the kernel as discussed in Section 5, we store at most 16 total $\boldsymbol{K}_{\text{M2L}}^{V,B}$ matrices for a homogeneous kernel.

Local to local (L2L) translations. Contributions from $\mathscr{F}^B \setminus L_I^B$ are captured through the local field computed for B's parent box, P, using L2L operators (Figure 3,

right). We translate $\phi^{P,d}$ at $\boldsymbol{y}^{P,d}$ to $\phi^{B,d}$ at $\boldsymbol{y}^{B,d}$ using the equation

$$K[\boldsymbol{y}^{B,d}, \phi^{B,d}](\boldsymbol{x}) = K[\boldsymbol{y}^{P,d}, \phi^{P,d}](\boldsymbol{x}) \quad \text{for all } \boldsymbol{x} \in \boldsymbol{x}^{B,d}. \tag{28}$$

The right-hand side is computed as a contribution to $u^{B,d}$, so (28) for B at depth ℓ becomes

$$u_{\text{L2L}}^{B,d} = \boldsymbol{K}_{\text{L2L}}^{P,B} \boldsymbol{\phi}^{P,d}, \tag{29}$$

such that

$$\boldsymbol{\phi}^{B,d} = (\boldsymbol{K}_{\text{L2L}}^{B,B})^{-1}\big(u_{\text{M2L}}^{B,d} + u_{\text{L2L}}^{B,d}\big). \tag{30}$$

The precomputation of matrix $\boldsymbol{K}_{\text{L2L}}^{P,B}$ is completely analogous to $\boldsymbol{K}_{\text{M2M}}^{C,B}$, with parent and child swapped.

*Local to grid target (*L2T*) translations.* For each leaf box B, we evaluate $u^{B,g}$ at grid locations, $\boldsymbol{x}^{B,g}$. At depth ℓ, $\phi^{B,d}$ accounts for all contributions from $\mathscr{F}^B$ while direct near-field calculations (discussed in detail in Section 4.4) account for the contributions from $\mathscr{N}^B$. The far-field potential is computed using L2T operators (Figure 4).

$$u(\boldsymbol{x}) = K[\boldsymbol{y}^{B,d}, \boldsymbol{\phi}^{B,d}](\boldsymbol{x}), \quad \boldsymbol{x} \in \boldsymbol{x}^{B,g},$$

or, in matrix form

$$u^{B,g} = \boldsymbol{K}_{\text{L2T}}^{B,B} \boldsymbol{\phi}^{B,d}. \tag{31}$$

For a uniformly refined tree, all leaves are at the same level, so we precompute and store one $\boldsymbol{K}_{\text{L2T}}^{B,B}$ matrix.

4.4. *Near-field interactions.*

After the far-field contributions are computed, the final step is to compute near-field interactions for leaf boxes. This is the most expensive step in the computation, if carried out naïvely, and it is essential to


Figure 4. Kernel-independent FMM translation operators for L2T.

optimize this part of the algorithm. For each leaf box B, we need to compute the influence of the volume density g^U for every box $U \in L_B^N$ (the near field boxes). Given a polynomial approximation γ^U to g^U, we evaluate the potential on an $n \times n \times n$ grid of samples $x^{B,g}$ on B, which we then add to the far field contribution computed in (31).

The principal mechanism to accelerate this step is based on the observation that we may use a regular grid pattern of points in B, permitting the use of precomputation. More precisely,

$$u^{B,g}(x) = \sum_U K[U, g](x) = \sum_U \sum_{j=1}^{N_k} \gamma_j^U F_j^{U,B}(x), \tag{32}$$

$$F_j^{U,B}(x) = \int_U \beta_j \left(2^\ell(y - c_U)\right) K(x, y)\, dy \quad \text{for } x \in x^{B,g}, \tag{33}$$

where c_U is the center of box U. We evaluate $u^{B,g}$ on a uniform grid $x_i^{B,g}$, $i = 1 \ldots n^3$ for $n < 6$ and on a tensor product Chebyshev grid for $n > 6$ to avoid condition problems, as discussed in Section A.2 of the Appendix. In matrix form (32) becomes

$$u^{B,g} = \sum_U F^{U,B} \gamma^U. \tag{34}$$

For a uniform octree, there are at most 27 possible locations for $U \in L_N^B$ with respect to B itself; using symmetries, however, only 4 are unique up to translation and rotation (Section 5). As in the S2M computations, adaptive Gaussian quadrature [8] is used to precompute and store the weights for these matrices. This can be done to machine precision for the function u and its first or second derivatives.

As each leaf box, B is not dependent on the near-field computations of any leaf box in T, it can quickly be seen how even the simplest approaches can take advantage of parallel architectures in the near-field computations. In Section A.4 of the Appendix, we discuss how we use OpenMP [19] and load-balancing approaches to parallelize the near-field and other computational steps of the FMM for shared-memory, multiprocessor architectures.

4.5. *Polynomial approximation of the solution.* In order to compute the value of u^B at an arbitrary point in the box, it is convenient to approximate it as a polynomial υ^B using a least-squares fit: minimizing

$$\sum_{i=1}^{n^3} \left\| u^B(x_i) - \sum_{j=1}^{N_n} \upsilon_j^B \beta_j(x_i - c_B) \right\|^2 \quad \text{for } x_i \in x^B, \tag{35}$$

where $\beta_j \in \{P_a(x)P_b(y)P_c(z), 0 \le a + b + c \le n - 1\}$ using either a monomial or Chebyshev polynomial basis, depending on the desired order n. For $n \le 6$ it is more convenient to use regular grids, while for $n > 6$ Chebyshev grid points provide

greater stability. In Section A.2 we demonstrate the accuracy of equispaced points and Chebyshev points for $n = 4, 6, 8$. For box B at depth ℓ, if Γ is the matrix with entries $\Gamma_{ij} = \beta_j(2^\ell(\boldsymbol{x} - \boldsymbol{c}_B))$, (35) leads to the equation $v^B = \Gamma^{(+)} u^{B,g}$, where the pseudoinverse $\Gamma^{(+)}$ needs to be precomputed only once as it does not depend on the kernel and is scale-invariant in *all* cases; that is, $\Gamma_{ij} = \beta_j(\boldsymbol{x}_i^*)$ where $\boldsymbol{x}_i^*$ are grid points in $B^* = [-1, 1]^3$. Once the v_j^B are known, we can evaluate the solution at an arbitrary point $\boldsymbol{x}_t \in B$ as

$$u(\boldsymbol{x}_t) = \sum_{j=1}^{N_n} v_j^B \beta_j(\boldsymbol{x}_t - \boldsymbol{c}_B). \tag{36}$$

In general, we assume that k, the order of the approximation γ^B of the force g^B, is equal to n, the order of approximation of v^B; however, as source and target locations need not be the same, k and n can be different.

4.6. *Polynomial force approximation from grid samples.* We have assumed the right-hand side is already given as a polynomial. However, if the force is available in another form (e.g., as samples on an AMR grid or polynomials on an unstructured finite element grid), we need simply to build a k-th order approximation of the right-hand side to the desired tolerance at regular grid points on each leaf node, followed by conversion to a polynomial representation, as in the preceding section. We view this step as outside the scope of the present paper.

4.7. *Nonuniform source distributions and adaptive FMM.* For nonuniform source distributions, leaf boxes may appear at different levels, leading to several additional types of interactions between boxes that need to be taken into account. For adaptive octrees, the number of relative positions of boxes one needs to consider can become very large. To avoid storing large number of precomputed matrices, we consider *level-restricted refinement*: we require adjacent leaf boxes be within one level of each other, a common restriction in tree codes and structured grids. Many fast approaches exist to convert arbitrary octrees to ones satisfying this constraint [55]; we currently use a straightforward sequential algorithm similar to [24].

Lists for adaptive FMM. Our definitions and notation follow [26; 30; 31]. For leaf box B, we define U and W lists:

- The *U-list*, L_U^B, consists of leaves adjacent to B, including itself; $L_U^B = L_N^B$ for uniform trees.
- The *W-list*, L_W^B, is the set of descendants of B's neighbors, not adjacent to B, but whose parents are adjacent to B. For any $W \in L_W^B$, W is at a finer level than B and $W \in \mathcal{N}^B$ (conversely, $B \in \mathcal{F}^W$).

For leaf and nonleaf boxes B, we define V and X lists.

- The *V-list*, L_V^B, is the set of B's parent's neighbor's children, not adjacent to B. $L_V^B = L_I^B$ for uniform trees.

- The *X-list*, L_X^B, is the set of boxes A such that $B \in L_W^A$.

We note that

$$B \in L_U^A \iff A \in L_U^B, \quad B \in L_V^A \iff A \in L_V^B, \quad B \in L_W^A \iff A \in L_X^B.$$

An example domain with labeled lists is shown in [60]; possible positions of boxes in the L_U^B, L_V^B, L_W^B and L_X^B are shown in Figure 5.

By the following lemma, for level-restricted trees, boxes in W and X lists have finite possible positions.

Lemma 4.1. *For a level-restricted tree T in which all neighboring leaf boxes are within one level of each other in the octree, for a box, B, all boxes in L_W^B and L_X^B must also be within one level of B.*

Proof. For box B assume there exists $W \in L_W^B$ such that $\ell_W - \ell_B \geq 2$. Then W's parent, P_W, satisfies $\ell_{P_W} - \ell_B \geq 1$, so there exist descendants D of P_W with $D \in L_U^B$ and $\ell_D - \ell_B \geq 2$, violating our tree-level restriction. Thus, $\ell_W - \ell_B \leq 1$. Since $W \in L_W^B$ implies $B \in L_X^W$, we have $\ell_B - \ell_X \leq 1$ for all $X \in L_X^B$. $\qquad\square$

Boxes in L_U^B and L_V^B are handled as boxes in L_N^B and L_I^B, respectively, are in the uniform case. For leaf box B, if $W \in L_W^B$, then $W \notin \mathscr{F}^B$; therefore, W's contribution to B is not accounted for through its parent, P_B, but since $B \in \mathscr{F}^W$, we can evaluate $\phi^{W,u}$ at $\boldsymbol{x}^{B,g}$. Hence, using notation analogous to other operators, M2T operators need to be defined. Further, for $X \in L_X^B$, $B \in \mathcal{N}^X$, but $X \in \mathscr{F}^B$. Thus, we need to


same level U   finer level U   coarser level U

W   V   X

Figure 5. Possible box positions for different lists in a level-restricted tree in 2D. The configurations in 3D are analogous.

evaluate contributions from X directly but can apply them to $\phi^{B,d}$; that is, we need to define an S2L operator.

To summarize, for adaptive FMM, in addition to M2M, M2L, L2L, and L2T, two additional operators, M2T and S2L need to be defined, and S2M and near-field (S2T) operators need to handle leaf boxes at arbitrary levels. We begin by describing changes to S2M and S2T operators and follow with a discussion of M2T and S2L operators.

S2M *operators for the adaptive case.* For *homogeneous* kernels, we store a *single* matrix $T_{\mathrm{S2M}}^{B^*}$, scaling for level ℓ as was done for the M2M and L2L operators. Let $B^* = [-1, 1]^3$ at $\ell = 0$. Then, for $x \in B$ at level ℓ, let $x^* = 2^\ell(x - c^B)$ for $x^* \in B^*$. For scaling exponent r, we have $K(x_i, y_j) = 2^{-r\ell} K(x_i^*, y_j^*)$, and (21) becomes

$$F_j^B(x) = 2^{-(r+2)\ell} K[B^*, \beta_j](x^*) = 2^{-(r+2)\ell} F_j^{B^*}(x^*) \quad \text{for all } x \in x^{B,g}, \ x^* \in x^{B^*,g}.$$

In matrix form, $F_{\mathrm{S2M}}^B = 2^{-(r+2)\ell} F_{\mathrm{S2M}}^{B^*}$ and $K_{\mathrm{S2M}}^B = 2^{-r\ell} K_{\mathrm{S2M}}^{B^*}$. Solving for $\phi^{B,u}$, (23) becomes

$$\phi^{B,u} = T_{\mathrm{S2M}}^{B^*} \gamma^B, \tag{37}$$

where $T_{\mathrm{S2M}}^{B^*}$ is precomputed and stored. (For inhomogeneous kernels, we store one matrix per leaf level).

Neighbor list interactions for adaptive trees. For adaptive level-restricted trees, leaves may exist at any level and $U \in L_U^B$ may exist at one level finer or coarser than B. As above for the S2M operators, for homogeneous kernels with scaling exponent r, we only compute matrices for pairs (B, U) with B scaled to B^* (U is appropriately scaled as well to U^*) such that (33) and (34) become

$$F_j^{U,B}(x) = 2^{-(r+2)\ell} K[U^*, \beta_j](x^*) = 2^{-(r+2)\ell} F_j^{(U,B)^*}(x^*),$$
$$\text{for all } x \in x^{B,g}, x^* \in x^{B^*,g},$$

$$u^{B,g} = \sum_U F_{\mathrm{S2T}}^{U,B} \gamma^U = 2^{-(r+2)\ell} \sum_U F_{\mathrm{S2T}}^{(U,B)^*} \gamma^U.$$

Along with 27 possible same-level neighbors, there are 56 fine-level neighbors (one level deeper) and 7 coarse-level neighbors (one level higher), all constituting the 90 possible locations for boxes in L_U^B in a level-restricted octree. Using symmetries (Section 5), we only precompute and store 10 matrices. For inhomogeneous kernels, this set of matrices is precomputed for each level for which leaf boxes exist.

M2T *and* S2L *operators.* For leaf box B and $W \in L_W^B$, we need an operator that evaluates $\phi^{W,u}$ at $x^{B,g}$:

$$u^{B,g}(x) = \sum_W K[y^{W,u}, \phi^{W,u}](x) \text{for } x \in x^{B,g},$$

or, in matrix form,

$$u^{B,g} = \sum_W \boldsymbol{K}^{W,B}_{\text{M2T}} \phi^{W,u}, \tag{38}$$

for precomputed $\boldsymbol{K}^{W,B}_{\text{M2T}}$. For all boxes B, L^B_X contains leaves X, for which contributions to B are computed at $\boldsymbol{x}^{B,d}$:

$$u^{B,d}(\boldsymbol{x}) = \sum_X K[X, g^X](\boldsymbol{x}) \approx \sum_X \sum_{j=1}^{N_k} \gamma^X_j F_j(\boldsymbol{x}) \quad \text{for } \boldsymbol{x} \in \boldsymbol{x}^{B,d},$$

or, in matrix form,

$$u^{B,d} = \sum_X \boldsymbol{F}^{X,B}_{\text{S2L}} \gamma^X. \tag{39}$$

There are 152 possible locations for $W \in L^B_W$; however, only six locations are distinct up to translation and rotation; also, due to the inverse relationship between L^B_X and L^B_W, the number of symmetry classes is the same (Section 5). For homogeneous kernels, only six $\boldsymbol{K}^{W,B}_{\text{M2T}}$ and six $\boldsymbol{F}^{X,B}_{\text{S2L}}$ matrices are precomputed for level $\ell = 0$ and scaled as necessary. For inhomogeneous kernels we compute and store these sets for each leaf level.

Remark. In cases where the order of γ is low compared to the order of $\phi^{B,u}$ and $\phi^{B,d}$, the size of M2T and S2L operators may actually be *larger* than those needed for direct computation of contributions from $W \in L^B_W$ or $X \in L^B_X$ to $\boldsymbol{x}^{B,g}$. Assuming we have a homogeneous kernel, if $W \in L^B_W$ is a leaf box, we can replace $\boldsymbol{K}^{W,B}_{\text{M2T}}$ by $\boldsymbol{F}^{W,B}_{\text{S2T}}$, constructed exactly in the same way as for boxes in the neighbor list L^B_U. Similarly, for leaf box B and a box $X \in L^B_X$, we can replace $\boldsymbol{F}^{X,B}_{\text{S2L}}$ with $\boldsymbol{F}^{X,B}_{\text{S2T}}$. For homogeneous kernels, these operators are computed for B^* only and scaled as necessary as in Section 4.4.

4.8. *Pseudocode and complexity for kernel-independent FMM volume solver.*

Pseudocode. The algorithm is summarized on the next page. We assume that a tree-level restricted octree T already exists [24] and that for each box, B, we are given the approximation, γ^B, to the force g^B (we discuss how to construct γ from g in Section 4.6). For clarity, we do not include the optimization of replacing M2T and S2L with S2T operators when more efficient as discussed above.

Computational complexity and storage requirements. We analyze the complexity for a uniformly refined octree. The analysis for the adaptive FMM is similar but slightly more complicated. We assume a homogeneous scalar kernel such as the Laplace kernel in (10) for analyzing the storage and computational complexities. Further, we assume that there are ℓ levels in the octree T. For a uniform tree, this implies we have $M_\ell = 8^\ell$ leaves and $M_t = (8^{\ell+1} - 1)/7$ total boxes in T. If we are using a k-th order polynomial approximation to the force at each leaf, we further assume there are approximately $N = M_\ell n^3$ total target points and $C = M_\ell N_k$ total

STEP 1 - BUILD LISTS
for each box B in *preorder* traversal of T **do**
   build L_U^B, L_W^B, L_X^B, and L_V^B (Section 4.7)
end for
STEP 2 - UPWARD PASS (Section 4.2)
for each box B in *postorder* traversal of T **do**
  if B is a leaf box **then**
    Convert local force approximations to upward densities:
$$\phi^{B,u} := \boldsymbol{T}_{\text{S2M}}^{B} \gamma^{B} \quad (23)$$
  else
    Translate children's upward densities to parent's upward density:
$$\phi^{B,u} := \sum_C \boldsymbol{T}_{\text{M2M}}^{C,B} \phi^{C_i,u} \quad (25)$$
  end if
end for
STEP 3 - DOWNWARD PASS (Section 4.3)
for each nonroot box B in *preorder* traversal of T **do**
  Add potentials due to parent downward density, U and X boxes to get the downward check potential:
$$u^{B,d} := \boldsymbol{K}_{\text{L2L}}^{P,B} \phi^{P,d} + \sum_{V \in L_V^B} \boldsymbol{K}_{\text{M2L}}^{V,B} \phi^{V,u} + \sum_{X \in L_X^B} \boldsymbol{F}_{\text{S2L}}^{X,B} \gamma^X \quad (29), (27), (39)$$
  Translate the check potential to the downward density:
$$\phi^{B,d} := (\boldsymbol{K}_{\text{L2L}}^{B,B})^{-1} u^{B,d} \quad (30)$$
  if B is a leaf box **then**
    Compute potentials from adjacent and W boxes to the potential at grid locations:
$$u^{B,g} := \sum_{U \in L_U^B} \boldsymbol{F}_{\text{S2T}}^{U,B} \gamma^U + \sum_{W \in L_W^B} \boldsymbol{K}_{\text{M2T}}^{W,B} \phi^{W,u} \quad (34), (38)$$
    Add the potential from the far field:
$$u^{B,g} := u^{B,g} + \boldsymbol{F}_{\text{L2T}} \phi^{B,d} \quad (31)$$
  end if
end for

Algorithm 1. Kernel-independent volume FMM.

coefficients. Let p be the number of coefficients sought in the multipole expansion, affecting the size of the equivalent densities and surfaces. For a desired level of precision in the expansion, say $\epsilon_{\text{fmm}} = 10^{-n_p}$, we have $p = n_p^3 - (n_p - 2)^3$. In Table 1, we indicate the computational complexity of each step of the nonadaptive FMM algorithm as well as the amount of precomputation and storage used for operators at each step. For nonuniform source distributions, we store additional

Operator	Complexity	Storage
S2M: $\boldsymbol{T}_{\text{S2M}}^{B}$	$O(Cp)$	pN_k
M2M: $\boldsymbol{T}_{\text{M2M}}^{C,B}$	$O((M_t - M_\ell)p^2)$	p^2
M2L: $\boldsymbol{K}_{\text{M2L}}^{V,B}$	$O(M_t p^{3/2} \log p + 189 M_t p^{3/2})$	$16p^{3/2}$
L2L: $\boldsymbol{K}_{\text{L2L}}^{P,B}, (\boldsymbol{K}_{\text{L2L}}^{B,B})^{-1}$	$O(M_t p^2)$	$2p^2$
L2T: $\boldsymbol{K}_{\text{L2T}}^{B,B}$	$O(Np)$	pn^3
Near Interaction: $\boldsymbol{F}_{\text{S2T}}^{U,B}$	$O(27NN_k)$	$4N_k n^3$
U-list (adaptive): $\boldsymbol{F}_{\text{S2T}}^{U,B}$		$10N_k n^3$
W-list: $\boldsymbol{F}_{\text{S2T}}^{W,B}, \boldsymbol{K}_{\text{M2T}}^{W,B}$		$6n^3(N_k + p)$
X-list: $\boldsymbol{F}_{\text{S2T}}^{X,B}, \boldsymbol{F}_{\text{S2L}}^{X,B}$		$6N_k(n^3 + p)$

Table 1. Computational complexity and storage requirements for a scalar homogeneous kernel. These values scale linearly for matrix and inhomogeneous kernels.

operators for the near-field interactions in the U, W, and X operators; the complexity of these operators are based on the degree of adaptivity.

Finally, we note that the computational and storage complexities will scale linearly for matrix or inhomogeneous kernels. For example, for the Stokes kernel in (12), the number of coefficients, p, scales as a results of the matrix kernel size to $p = 9(n_p^3 - (n_p - 2)^3)$. For the modified Helmholtz kernel in (11), the inhomogeneous nature of the kernel results in an increased storage complexity, which varies depending on the number of different levels in the tree.

5. Symmetries for precomputed interaction operators

For a box B and all boxes in L_U^B, L_V^B, L_W^B and L_X^B, the number of different relative positions can be large, so precomputing all possible interaction matrices may require significant time and storage. Performance can also be affected by the need for random access of large amounts of precomputed data. The number of precomputed matrices can be substantially reduced via symmetries; that is, many box positions are equivalent in the sense that there is a rigid transformation $\mathcal{T}$, mapping box Z_1 to Z_2 and box B to itself. We store a single matrix for a representative box for each symmetry class, obtaining matrices for all elements of the class by applying $\mathcal{T}$ to the matrix for the representative box.

For every list type $Z \in \{U, V, W, X\}$, we define a set of possible positions $\text{Pos}(Z)$ and a set of symmetry classes which form a partition of $\text{Pos}(Z)$. For each class, we

define a *reference box*, and for each box position in $\text{Pos}(Z)$, we need an efficient way to determine its class and a transformation $\mathcal{T}(B) : \mathbb{R}^3 \to \mathbb{R}^3$ mapping it to the reference box.

For all lists, the symmetries are related to the transformations of space which map a grid of cubes to itself. We consider N^3 grids of sizes 1^3 to 7^3 (we discuss which lists correspond to which cubes in more detail below) and begin by classifying all symmetries of such grids.

Grid symmetries. The cubes on the N^3 grid are indexed by (i, j, k) with values $-M \ldots -1, 0, 1 \ldots M$ for odd $N = 2M + 1$ and $-M \ldots -1, 1 \ldots M$ for even $N = 2M$. We skip index 0 for even grids to ensure cube centers and indices are transformed by symmetries in the same way. If the cube size is 1, cube centers are exactly the indices (i, j, k) for odd N and differ by $\pm\frac{1}{2}$ for even N, depending on the index sign. Each N^3 grid can be partitioned into M (for even N) or $M + 1$ (for odd N) layers. Layer 0 consists of one cube and exists for odd N, and layer M consists of cubes on the surface of the N^3 grid. For odd N, layer l has size $(2l + 1)^3$ and for even N, layer l has size $(2l)^3$.

The group of symmetries G_{cube} of a cube has order 48. For a cube centered at zero, transformations in G_{cube} are compositions of rotations and reflections, mapping each axis direction to another, possibly with orientation reversed. Any permutation of directions is possible, so we identify the group with $S_3 \times J^3$, where S_3 is the group of permutations of length 3, and J is the two-element group of reflections. The rotational part of any element of G_{cube} can be specified as a permutation of length 3 on the set of axes $\{x, y, z\}$, with an orientation 1 or -1 specified for each axis. Transformations from G_{cube} encoded in this way can be applied to points very efficiently: for a point $x \in \mathbb{R}^3$, the permutation is applied to its coordinates, which are then scaled by 1 or -1.

For the N^3 grid, the equivalence classes under the action of G_{cube} can be enumerated combinatorially. If two indices (i, j, k) and (i', j', k') differ only by signs of components, corresponding cubes are in the same class, mapped by reflections. To enumerate all classes, we consider cubes with nonnegative indices. Two cubes with nonnegative indices (i, j, k) and (i', j', k') are in the same class if and only if there is a permutation mapping (i, j, k) to (i', j', k'). For $i \neq j \neq k$ and $i, j, k \in [1, M]$, seven series of equivalence classes are easily enumerated, corresponding to signatures (i, j, k), (i, i, j), (i, i, i), $(0, i, i)$, $(0, 0, i)$ and $(0, 0, 0)$. A reference box in every class is uniquely defined by requiring that its three indices are all nonnegative and are in nondecreasing order. As an example, the figure shows the representative box for the $(1, 2, 3)$ class in the 7^3 grid for $M = 3$.

representative

$(-M,-M,-M)$

The properties of classes in each series are summarized in Table 2. For a box Z, with grid index (i, j, k) relative to B, the reference box is obtained by taking absolute values and sorting the indices; sign changes and a permutation mapping (i, j, k) to the reference box index also encode the transformation.

series signature	7^3 classes	reference cube	classes per grid	classes per layer	class size
(i, j, k)		$(|i|, |j|, |k|)$, $|i| < |j| < |k|$	$\binom{M}{3}$	$\binom{M-1}{2}$	48
(i, i, j)		$(|i|, |i|, |j|)$, $|i| < |j|$ or $(|i|, |j|, |j|)$	$M(M-1)$	$2(M-1)$	24
(i, i, i)		$(|i|, |i|, |i|)$	M	1	8
$(0, i, j)$		$(0, |i|, |j|)$	$\binom{M}{2}$	$M-1$	24
$(0, i, i)$		$(0, |i|, |i|)$	M	1	12
$(0, i, i)$		$(0, 0, |i|)$	M	1	6
$(0, 0, 0)$	—	$(0, 0, 0)$	1	—	1

Table 2. Series of equivalence classes of cubes in an N^3 grid, with $M = \lfloor N/2 \rfloor$. For even N only the first 3 series of classes may be nonempty. For odd N all classes are present. For $M \leq 2$, (i, j, k) classes are empty, and for $M = 1$, (i, i, j) and $(0, i, j)$ classes are also empty. Class $(0, 0, 0)$, corresponding to the center of the grid, exists only in layer 0. Boxes in different classes in one series are marked with circles of different colors; representative boxes are marked with circles with black border. The view is from the top, with first index direction to the right, second direction up and third towards the viewer, as in the figure on page 99. The total number of classes for $(2M)^3$ layers is $(M+1)M/2$ (classes (i, j, M) with $i, j = 1 \ldots M$, $i \leq j$), and for $(2M+1)^3$ layers, it is $(M+2)(M+1)/2$ (classes (i, j, M) with $i, j = 0 \ldots M$, $i \leq j$).

Symmetries of L_U^B. Due to the tree-level restriction, boxes $U \in L_U^B$ are either neighbors of B, neighbors of B's parent and adjacent to B, or adjacent children of neighbors of B. We denote these three sublists of L_U^B by $L_{U,n}^B$, $L_{U,p}^B$ and $L_{U,c}^B$ respectively. Note that $A \in L_{U,p}^B$ is equivalent to $B \in L_{U,c}^A$; hence, it is sufficient to consider $L_{U,n}^B$ and $L_{U,c}^B$. The neighbors of B on the same level as B form a 3^3 grid centered at B, so from Table 2, the number of classes is 4: $(1, 1, 1)$, $(0, 1, 1)$, $(0, 0, 1)$, and $(0, 0, 0)$. Locations of $U \in L_{U,c}^B$ can be thought of as the outer layer of a 4^3 grid, with $M = 2$ and B as the 2^3 subgrid in the center, so we obtain 3 classes: $(1, 1, 2)$, $(1, 2, 2)$, $(2, 2, 2)$, giving 10 classes for L_U^B.

Symmetries of L_V^B. Boxes in L_V^B are children of neighbors of the parent of B, so they can all be represented by cubes of a 6^3 grid; however, the group of rigid transformations of the grid mapping to itself do not necessarily preserve B. Hence, instead regard L_V^B as a subset of a 7^3 grid centered at B with $M = 3$. All $V \in L_V^B$ are in layers 2 and 3, and there are 10 classes: for layer 3, classes $(i, j, 3)$, $i, j = 1 \ldots 3$, $i \le j$ and for layer 2, classes $(i, j, 2)$, $i, j = 0, 1, 2$, $i \le j$. Because we consider only a subset of the full 7^3 grid, the class sizes are smaller, but it can easily be seen that no class becomes empty, so the number is optimal.

Symmetries of L_W^B, L_X^B. For a level-restricted tree, boxes $W \in L_W^B$ are children of neighbors of B not adjacent to B, that is, they reside in the surface layer of a 6^3 grid with B as the central 2^3 grid. For $M = 3$, we have 6 classes from Table 2: $(i, j, 3)$, $i, j = 1 \ldots 3$, $i \le j$. Due to duality, the number of classes for L_X^B is the same, but the class sizes may *not* be the same.

Summary. For a given pair (B, Z), if $Z \in \{L_{U,n}^B, L_{U,c}^B, L_V^B, L_W^B\}$, determine the translation and scaling which map B to the central box or 2^3 subgrid of a larger grid. Then, apply the same transformation to the center of Z; resulting coordinates yield the index (i, j, k), which is translated into the reference box and rotation as described above.

6. Numerical results

Our algorithm has been implemented in C++, and we have tested several kernels and source and target point distributions. Our tests were run on an Intel Xeon-based X7560 (2.27 GHz, 64 bit) system with 16 CPUs and 128 GB of RAM; the major computation loops are accelerated with OpenMP [19] as discussed in Section A.4.

We first test the free-space Poisson solver on three different types of problems designed to show how our algorithm handles increasing levels of complexity in the force distribution.

We use an adaptive-refinement strategy similar to [24]. For this, we compute a k-th order polynomial approximation, γ^B, to the force $g^B(x)$ sampled on a

Figure 6. Sample force distributions based on adaptive refinement. Each point, colored by its tree level, indicates the center of a leaf box B. Top left: A single sharply peaked Gaussian function (Example 1). Top right: A discontinuous force distribution, equal to 1 inside a sphere and 0 outside (Example 2). Bottom: A discontinuous force distribution involving oscillatory functions restricted to the interiors of a set of three spheres (Example 3).

$k \times k \times k$ grid. We let $\tilde{g}^B$ be the force evaluated on a refined $2k \times 2k \times 2k$ grid. If $\|g^B(x) - \tilde{g}^B(x)\|_2 > \epsilon_{\text{rhs}}$, B is subdivided, and the octree is balanced as needed. Three force distributions, used in Examples 1–3 below, are shown in Figure 6.

Example 1. The first experiment tests the accuracy of our method for solving the Poisson equation — Equation (7) with kernel (10) — with a fast-decaying smooth right-hand side:

$$-\Delta u(x) = \sum_{i=0}^{8} -e^{-L\|x-x_i\|^2}(4L\|x-x_i\|^2 - 6L), \quad L = 250,$$

with solution

$$u(x) = \sum_{i=0}^{8} -e^{-L\|x-x_i\|^2},$$

where $\boldsymbol{x}_i = \left(\pm \frac{3}{40}, \pm \frac{3}{40}, \pm \frac{3}{40}\right)$ inside the $[-1, 1]^3$ box. This test requires a high degree of adaptivity to achieve good accuracy with a limited number of points.

In Table 3, ϵ_{fmm} is the precision of the translation operators, ϵ_{rhs} is the refinement criterion for the adaptive refinement of the source distribution, and M_ℓ is the number of leaves in the tree T with L_T levels. The number of points N_{pts} is computed as $M_\ell k^3$ where k is the order of the polynomial. This number of points per leaf is chosen to be sufficiently large to build the polynomial approximation of order k. The computation time T_{FMM} is given in seconds, and the *rate* is in points per second. E_2 and E_∞ are the relative L^2 and L^∞ errors, respectively. Timings include FMM evaluation times only; when the precision ϵ_{fmm} remains constant, the rate of work

ϵ_{fmm}	ϵ_{rhs}	M_ℓ	N_{pts}	L_T	E_2	E_∞	T_{FMM}	rate
				fourth-order force approximation				
10^{-2}	10^{-2}	736	47104	6	$2.3 \cdot 10^{-2}$	$2.4 \cdot 10^{-2}$	$9.1453 \cdot 10^{-3}$	$5.15 \cdot 10^{+6}$
10^{-4}	10^{-2}	736	47104	6	$1.1 \cdot 10^{-3}$	$6.8 \cdot 10^{-4}$	$2.4328 \cdot 10^{-2}$	$1.94 \cdot 10^{+6}$
10^{-4}	10^{-4}	3088	197632	7	$1.1 \cdot 10^{-4}$	$2.7 \cdot 10^{-4}$	$9.6590 \cdot 10^{-2}$	$2.05 \cdot 10^{+6}$
10^{-6}	10^{-4}	3088	197632	7	$3.8 \cdot 10^{-5}$	$2.5 \cdot 10^{-5}$	$3.7906 \cdot 10^{-1}$	$5.21 \cdot 10^{+5}$
10^{-6}	10^{-6}	19328	1236992	8	$3.8 \cdot 10^{-6}$	$3.6 \cdot 10^{-6}$	$2.3889 \cdot 10^{+0}$	$5.18 \cdot 10^{+5}$
10^{-8}	10^{-6}	19328	1236992	8	$3.7 \cdot 10^{-6}$	$1.3 \cdot 10^{-6}$	$6.2057 \cdot 10^{+0}$	$1.99 \cdot 10^{+5}$
10^{-8}	10^{-8}	143088	9157632	9	$1.6 \cdot 10^{-7}$	$8.8 \cdot 10^{-8}$	$4.5280 \cdot 10^{+1}$	$2.02 \cdot 10^{+5}$
				sixth-order force approximation				
10^{-4}	10^{-4}	1408	304128	6	$1.1 \cdot 10^{-4}$	$2.3 \cdot 10^{-4}$	$1.1016 \cdot 10^{-1}$	$2.76 \cdot 10^{+6}$
10^{-6}	10^{-4}	1408	304128	6	$1.4 \cdot 10^{-5}$	$3.5 \cdot 10^{-5}$	$1.8717 \cdot 10^{-1}$	$1.62 \cdot 10^{+6}$
10^{-6}	10^{-6}	4936	1066176	7	$9.0 \cdot 10^{-7}$	$2.2 \cdot 10^{-6}$	$6.6737 \cdot 10^{-1}$	$1.60 \cdot 10^{+6}$
10^{-8}	10^{-6}	4936	1066176	7	$3.3 \cdot 10^{-7}$	$1.6 \cdot 10^{-7}$	$1.6155 \cdot 10^{+0}$	$6.60 \cdot 10^{+5}$
10^{-8}	10^{-8}	20112	4344192	8	$2.4 \cdot 10^{-8}$	$6.2 \cdot 10^{-8}$	$6.8652 \cdot 10^{+0}$	$6.33 \cdot 10^{+5}$
10^{-10}	10^{-8}	20112	4344192	8	$1.8 \cdot 10^{-8}$	$1.0 \cdot 10^{-8}$	$1.5771 \cdot 10^{+1}$	$2.76 \cdot 10^{+5}$
10^{-10}	10^{-10}	92072	19887552	9	$6.3 \cdot 10^{-9}$	$9.7 \cdot 10^{-9}$	$7.5103 \cdot 10^{+1}$	$2.65 \cdot 10^{+5}$
				eighth-order force approximation				
10^{-6}	10^{-6}	2024	1036288	7	$9.2 \cdot 10^{-7}$	$3.5 \cdot 10^{-6}$	$4.6688 \cdot 10^{-1}$	$2.22 \cdot 10^{+6}$
10^{-8}	10^{-6}	2024	1036288	7	$3.7 \cdot 10^{-7}$	$8.2 \cdot 10^{-7}$	$7.4189 \cdot 10^{-1}$	$1.40 \cdot 10^{+6}$
10^{-8}	10^{-8}	5440	2785280	7	$1.9 \cdot 10^{-8}$	$6.6 \cdot 10^{-8}$	$2.1128 \cdot 10^{+0}$	$1.32 \cdot 10^{+6}$
10^{-10}	10^{-8}	5440	2785280	7	$7.7 \cdot 10^{-9}$	$7.6 \cdot 10^{-9}$	$4.3588 \cdot 10^{+0}$	$6.40 \cdot 10^{+5}$
10^{-10}	10^{-10}	22800	11673600	8	$4.1 \cdot 10^{-9}$	$5.6 \cdot 10^{-9}$	$1.9449 \cdot 10^{+1}$	$6.00 \cdot 10^{+5}$
10^{-12}	10^{-10}	22800	11673600	8	$2.6 \cdot 10^{-9}$	$4.7 \cdot 10^{-9}$	$4.1602 \cdot 10^{+1}$	$2.81 \cdot 10^{+5}$
10^{-12}	10^{-12}	50352	25780224	9	$2.1 \cdot 10^{-9}$	$4.6 \cdot 10^{-9}$	$9.2139 \cdot 10^{+1}$	$2.80 \cdot 10^{+5}$

Table 3. Results for free-space Poisson equation (Example 1): Gaussian bump at the origin.

per source and target point remains close to constant, as we would expect since the FMM algorithm scales linearly.

Example 2. In this example, we consider a discontinuous right-hand side, with $g(x) = 1$ inside a sphere of radius $R = 0.75$, and $g(x) = 0$ outside the sphere. Letting $r = \|x\|$, the problem becomes

$$-\Delta u(x) = \begin{cases} 1 & \text{if } r \leq R, \\ 0 & \text{else,} \end{cases}$$

with solution

$$-\Delta u(x) = \begin{cases} (R^2 - r^2)/6 + R^2/3 & \text{if } r \leq R, \\ R^3/3r^2 & \text{else.} \end{cases}$$

While this problem can be handled analytically, it serves as a useful test of performance on adaptive data structures that are refined in the neighborhood of a surface. The number of points indicates the total number of points both inside and outside the sphere. Since the coefficient representation of the force for a leaf node entirely outside of the sphere is zero, these boxes are ignored in all evaluation phases; this increases the computed rate somewhat. A greater speedup is achieved from the observation that leaf nodes entirely in the interior have a constant source distribution, so that only one polynomial coefficient is nonzero. This significantly accelerates both the near-field and S2M calculation stages. Results are shown in Table 4.

ϵ_{fmm}	ϵ_{rhs}	k	M_ℓ	N_{pts}	L_T	E_2	E_∞	T_{FMM}	rate
				fourth-order force approximation					
10^{-2}	10^{-2}	4	232	14848	4	$1.1 \cdot 10^{-2}$	$1.6 \cdot 10^{-2}$	$1.8346 \cdot 10^{-3}$	$8.09 \cdot 10^{+6}$
10^{-3}	10^{-3}	4	1184	75776	5	$4.3 \cdot 10^{-3}$	$4.6 \cdot 10^{-3}$	$1.1007 \cdot 10^{-2}$	$6.88 \cdot 10^{+6}$
10^{-4}	10^{-4}	4	5888	376832	6	$1.4 \cdot 10^{-4}$	$2.5 \cdot 10^{-4}$	$1.1327 \cdot 10^{-1}$	$3.33 \cdot 10^{+6}$
10^{-5}	10^{-5}	6	11432	2469312	7	$1.6 \cdot 10^{-5}$	$3.3 \cdot 10^{-5}$	$6.2147 \cdot 10^{-1}$	$3.97 \cdot 10^{+6}$
10^{-6}	10^{-6}	6	80088	17299008	8	$7.2 \cdot 10^{-6}$	$1.8 \cdot 10^{-5}$	$6.0647 \cdot 10^{+0}$	$2.85 \cdot 10^{+6}$
10^{-7}	10^{-7}	8	127856	65462272	8	$9.4 \cdot 10^{-7}$	$3.3 \cdot 10^{-6}$	$1.8852 \cdot 10^{+1}$	$3.47 \cdot 10^{+6}$
10^{-8}	10^{-8}	8	528984	270839808	10	$3.7 \cdot 10^{-7}$	$9.8 \cdot 10^{-7}$	$1.2335 \cdot 10^{+2}$	$2.20 \cdot 10^{+6}$
10^{-9}	10^{-9}	8	2074360	1062072320	10	$3.2 \cdot 10^{-8}$	$2.6 \cdot 10^{-7}$	$6.4900 \cdot 10^{+2}$	$1.64 \cdot 10^{+6}$

Table 4. Free-space Poisson equation (Example 2): discontinuous force.

Example 3. For our third example, we replicate an experiment from [47] for a highly oscillatory force with discontinuities along multiple surfaces, setting

$$f_m(r) = \begin{cases} ((r - r^2) \sin(2m\pi r))^2 & \text{if } r < 1, \\ 0 & \text{if } r \geq 1, \end{cases}$$

$$-\Delta u(x) = \frac{1}{R^3} \sum_{i=0}^{2} f_m(|x - c_i|/R), \tag{40}$$

where $c_0 = \left(\frac{3}{16}, \frac{7}{16}, \frac{13}{16}\right)$, $c_1 = \left(\frac{7}{16}, \frac{13}{16}, \frac{3}{16}\right)$, $c_2 = \left(\frac{13}{16}, \frac{3}{16}, \frac{7}{16}\right)$, $R = 0.05$, and the wavelength of f_m is $\lambda_m = R/(2m) = (1/40m)$. Defining

$$
\phi_m(r) = \frac{(10r^6 - 28r^5 + 21r^4 - 7)}{840} + \frac{(60r - 120)}{r\lambda_m^6} - \frac{9}{\lambda_m^4}
$$
$$
+ \left[\frac{(300r - 120)}{r\lambda_m^6} - + \frac{(30r^2 - 36r + 9)}{\lambda_m^4} + \frac{(r^4 - 2r^3 + r^2)}{2\lambda_m^2}\right]\cos(r\lambda_m)
$$
$$
+ \left[-\frac{360}{r\lambda_m^7} + \frac{(120r^2 - 96r + 12)}{r\lambda_m^5} + \frac{(5r^3 + 8r^2 - 3r)}{\lambda_m^3}\right]\sin(r\lambda_m),
$$

and

$$
\theta_m(r) = \left(\frac{360}{\lambda_m^6} - \frac{12}{\lambda_m^4} - \frac{1}{120}\right)/r,
$$

we write the solution to (40) as

$$
u^{\text{exact}}(x) = \begin{cases}
\phi(\|x - c_0\|/R) + \sum_{i=1,2}\theta(\|x - c_i\|/R) & \text{if } \|x - c_0\| < R, \\
\phi(\|x - c_1\|/R) + \sum_{i=0,2}\theta(\|x - c_i\|/R) & \text{if } \|x - c_1\| < R, \\
\phi(\|x - c_2\|/R) + \sum_{i=0,1}\theta(\|x - c_i\|/R) & \text{if } \|x - c_2\| < R, \\
\sum_{i=0}^{2}\theta(\|x - c_i\|/R) & \text{else.}
\end{cases}
$$

In order to compare our results to [47], we use the error metric introduced there. Let ϵ^B be the vector of errors calculated as the difference between the calculated and exact solutions on B and calculate the following norm over all leaf boxes.

$$
\|\epsilon_{\text{all}}^B\|_2 = \sum_B \left(\int_B \frac{\epsilon^B}{\|u^{\text{exact}}\|_\infty}\right)^{1/2}.
$$

As indicated in [47],

$$
\|u^{\text{exact}}\|_\infty = \left|\left(-\frac{6}{\lambda_m^4} - \frac{1}{120}\right)/R + \left(\frac{720}{\lambda_m^6} - \frac{24}{\lambda_m^4} - \frac{1}{120}\right)/\|c_i - c_j\|\right|
$$
$$
\text{for all } i, j = 0, 1, 2, i \neq j.
$$

Our automatic refinement strategy refines within or near the sphere surfaces, with refinement taking place in the exterior of the spheres only for the purpose of tree-balancing. We build coefficients only on leaf boxes which contain nonzero source distributions: either interior to or intersecting one of the three spheres. Results are shown in Table 5.

While the performance of our code cannot be compared easily to the optimized and parallelized scheme presented in [47], we have implemented schemes accurate to much higher order. Thus, as expected, we are able to reach comparable accuracies with significantly fewer points. To compare the number of points required, we consider the number of points in the finest level solve of their three-level examples. For

m	ϵ_{fmm}	ϵ_{rhs}	M_ℓ	N_{pts}	L_T	k	$\|\epsilon_{\mathrm{all}}^B\|_2$	$\|\epsilon_{\mathrm{all}}^B\|_\infty$
1	10^{-6}	10^{-6}	3984	254976	8	4	$2.3 \cdot 10^{-7}$	$2.2 \cdot 10^{-5}$
1	10^{-8}	10^{-8}	7296	1575936	9	6	$1.1 \cdot 10^{-8}$	$7.1 \cdot 10^{-7}$
1	10^{-10}	10^{-10}	24144	12361728	9	8	$1.6 \cdot 10^{-10}$	$1.8 \cdot 10^{-8}$
7	10^{-6}	10^{-6}	93816	6004224	10	4	$2.2 \cdot 10^{-6}$	$1.8 \cdot 10^{-4}$
7	10^{-8}	10^{-8}	195984	42332544	10	6	$9.3 \cdot 10^{-9}$	$1.1 \cdot 10^{-6}$
7	10^{-10}	10^{-10}	228312	116895744	10	8	$1.1 \cdot 10^{-10}$	$4.3 \cdot 10^{-8}$
15	10^{-6}	10^{-6}	140568	8996352	10	4	$6.4 \cdot 10^{-7}$	$8.6 \cdot 10^{-5}$
15	10^{-8}	10^{-8}	1092456	235970496	11	6	$7.2 \cdot 10^{-8}$	$4.2 \cdot 10^{-6}$
15	10^{-10}	10^{-10}	1596672	817496064	11	8	$4.8 \cdot 10^{-10}$	$6.1 \cdot 10^{-8}$
30	10^{-6}	10^{-6}	148272	9489408	10	4	$5.8 \cdot 10^{-7}$	$6.2 \cdot 10^{-5}$
30	10^{-8}	10^{-8}	1491216	322102656	11	6	$2.1 \cdot 10^{-8}$	$3.8 \cdot 10^{-6}$
30	10^{-10}	10^{-10}	1720152	880717824	12	8	$4.9 \cdot 10^{-9}$	$2.0 \cdot 10^{-6}$
60	10^{-6}	10^{-6}	150288	9618432	11	4	$6.0 \cdot 10^{-7}$	$9.6 \cdot 10^{-5}$
60	10^{-8}	10^{-8}	1502592	324559872	11	6	$9.2 \cdot 10^{-8}$	$1.4 \cdot 10^{-5}$
60	10^{-10}	10^{-9}	1659312	849567744	11	8	$7.7 \cdot 10^{-8}$	$1.7 \cdot 10^{-5}$

Table 5. Free-space Poisson equation (Example 3): Discontinuities along several spherical surfaces containing oscillating source distributions.

$m = 7$, we achieve accuracy on par with their most accurate tests with approximately $\frac{1}{100}$ as many points. For $m = 15$, we require approximately $\frac{1}{5}$ as many points, and with $\frac{1}{4}$ as many points, we achieve about two orders of magnitude greater accuracy. For $m = 30$, we achieve equivalent results with approximately $\frac{1}{4}$ as many points. Additionally, we extended the examples for an even higher wavenumber component ($m = 60$), decreasing the wavelength to $4.17 \cdot 10^{-4}$, and achieving good results with fewer than 10^9 points.

Example 4. For the modified Helmholtz equation — (8) with kernel (11) — we use a right-hand side similar to that of Example 1, setting the Helmholtz parameter (inverse Debye length) to $\alpha = \pi$:

$$\alpha u(\boldsymbol{x}) - \Delta u(\boldsymbol{x}) = \sum_{i=0}^{8} -e^{-L\|\boldsymbol{x}-\boldsymbol{x}_i\|^2}(4L\|\boldsymbol{x}-\boldsymbol{x}_i\|^2 - 6L - \alpha), \quad L = 250,$$

with solution

$$u(\boldsymbol{x}) = \sum_{i=0}^{8} -e^{-L\|\boldsymbol{x}-\boldsymbol{x}_i\|^2}, \quad \text{for } \boldsymbol{x}_i = \left(\pm\tfrac{3}{40}, \pm\tfrac{3}{40}, \pm\tfrac{3}{40}\right)$$

ϵ_{fmm}	ϵ_{rhs}	M_ℓ	N_{pts}	L_T	E_2	E_∞
fourth-order force approximation						
10^{-2}	10^{-2}	736	47104	6	$2.3 \cdot 10^{-2}$	$2.4 \cdot 10^{-2}$
10^{-4}	10^{-2}	736	47104	6	$6.0 \cdot 10^{-4}$	$5.6 \cdot 10^{-4}$
10^{-4}	10^{-4}	3088	197632	7	$1.1 \cdot 10^{-4}$	$1.8 \cdot 10^{-4}$
10^{-6}	10^{-4}	3088	197632	7	$2.4 \cdot 10^{-5}$	$2.1 \cdot 10^{-5}$
10^{-6}	10^{-6}	19328	1236992	8	$2.3 \cdot 10^{-6}$	$3.4 \cdot 10^{-6}$
sixth-order force approximation						
10^{-4}	10^{-4}	1408	304128	6	$1.1 \cdot 10^{-4}$	$2.4 \cdot 10^{-4}$
10^{-6}	10^{-4}	1408	304128	6	$1.4 \cdot 10^{-5}$	$3.5 \cdot 10^{-5}$
10^{-6}	10^{-6}	4936	1066176	7	$8.5 \cdot 10^{-7}$	$2.5 \cdot 10^{-6}$
10^{-8}	10^{-6}	4936	1066176	7	$1.4 \cdot 10^{-7}$	$1.3 \cdot 10^{-7}$
10^{-8}	10^{-8}	20112	4344192	8	$1.5 \cdot 10^{-8}$	$7.5 \cdot 10^{-8}$
10^{-10}	10^{-8}	20112	4344192	8	$8.9 \cdot 10^{-9}$	$7.4 \cdot 10^{-9}$
10^{-10}	10^{-10}	92072	19887552	9	$2.2 \cdot 10^{-9}$	$5.7 \cdot 10^{-9}$
eighth-order force approximation						
10^{-6}	10^{-6}	2024	1036288	7	$1.9 \cdot 10^{-6}$	$4.7 \cdot 10^{-6}$
10^{-8}	10^{-6}	2024	1036288	7	$1.8 \cdot 10^{-7}$	$4.6 \cdot 10^{-7}$
10^{-8}	10^{-8}	5440	2785280	7	$1.2 \cdot 10^{-8}$	$1.4 \cdot 10^{-8}$
10^{-10}	10^{-8}	5440	2785280	7	$7.7 \cdot 10^{-9}$	$7.6 \cdot 10^{-9}$
10^{-10}	10^{-10}	22800	11673600	8	$3.4 \cdot 10^{-9}$	$5.6 \cdot 10^{-9}$
10^{-12}	10^{-10}	22800	11673600	8	$1.3 \cdot 10^{-9}$	$2.0 \cdot 10^{-9}$
10^{-12}	10^{-12}	50352	25780224	9	$2.0 \cdot 10^{-9}$	$2.6 \cdot 10^{-9}$

Table 6. Free-space modified Helmholtz equation (Example 4): Gaussian bump at the origin.

inside of the $[-1, 1]^3$ box. All translation matrices are computed to a precision of $\epsilon_{\mathrm{fmm}}/10$. These matrices can be computed at run-time in a lazy manner; if α is known before run-time, these tables can be precomputed, stored, and loaded as necessary. Additionally, since the right-hand side is the same as in Example 1, we use the same point distribution; hence, the timings are essentially the same as in that example and are omitted here. Results are shown in Table 6.

Example 5. We test the ability of our code to handle matrix kernels by solving the Stokes equations — (9) with kernel (12) — with the following divergence-free fast-decaying force.

$$-\mu \Delta u(\boldsymbol{x}) + \nabla p(\boldsymbol{x}) = \sum_{i=0}^{8} (8L^3 \|\boldsymbol{x} - \boldsymbol{x}_i\|^2 - 20L^2) e^{-L\|\boldsymbol{x} - \boldsymbol{x}_i\|^2} (\nabla \times (\boldsymbol{x} - \boldsymbol{x}_i)),$$

ϵ_{fmm}	ϵ_{rhs}	M_ℓ	N_{pts}	L_T	E_2	E_∞	T_{FMM}	rate
				fourth-order force approximation				
10^{-2}	10^{-2}	2038	130432	6	$1.3 \cdot 10^{-1}$	$1.5 \cdot 10^{-1}$	$1.3485 \cdot 10^{-1}$	$9.67 \cdot 10^{+5}$
10^{-4}	10^{-2}	2038	130432	6	$1.0 \cdot 10^{-3}$	$8.4 \cdot 10^{-4}$	$9.4899 \cdot 10^{-1}$	$1.37 \cdot 10^{+5}$
10^{-4}	10^{-4}	10606	678784	7	$9.1 \cdot 10^{-4}$	$9.4 \cdot 10^{-4}$	$5.1145 \cdot 10^{+0}$	$1.33 \cdot 10^{+5}$
10^{-6}	10^{-4}	10606	678784	7	$8.4 \cdot 10^{-6}$	$1.1 \cdot 10^{-5}$	$2.2359 \cdot 10^{+1}$	$3.04 \cdot 10^{+4}$
10^{-6}	10^{-6}	69140	4424960	8	$7.5 \cdot 10^{-6}$	$1.4 \cdot 10^{-5}$	$1.4655 \cdot 10^{+2}$	$3.02 \cdot 10^{+4}$
10^{-8}	10^{-6}	69140	4424960	8	$2.2 \cdot 10^{-7}$	$4.4 \cdot 10^{-7}$	$4.8311 \cdot 10^{+2}$	$9.16 \cdot 10^{+3}$
10^{-8}	10^{-8}	484408	31002112	9	$1.4 \cdot 10^{-7}$	$4.4 \cdot 10^{-7}$	$3.2253 \cdot 10^{+3}$	$9.61 \cdot 10^{+3}$
				sixth-order force approximation				
10^{-4}	10^{-4}	2696	582336	7	$4.9 \cdot 10^{-4}$	$8.6 \cdot 10^{-4}$	$1.5720 \cdot 10^{+0}$	$3.70 \cdot 10^{+5}$
10^{-6}	10^{-4}	2696	582336	7	$4.3 \cdot 10^{-6}$	$9.9 \cdot 10^{-6}$	$5.9948 \cdot 10^{+0}$	$9.71 \cdot 10^{+4}$
10^{-6}	10^{-6}	10396	2245536	7	$8.1 \cdot 10^{-6}$	$1.3 \cdot 10^{-6}$	$2.3602 \cdot 10^{+1}$	$9.51 \cdot 10^{+4}$
10^{-8}	10^{-6}	10396	2245536	7	$1.6 \cdot 10^{-7}$	$4.1 \cdot 10^{-7}$	$7.7862 \cdot 10^{+1}$	$2.88 \cdot 10^{+4}$
10^{-8}	10^{-8}	59830	12923280	8	$1.4 \cdot 10^{-7}$	$4.3 \cdot 10^{-7}$	$4.1067 \cdot 10^{+2}$	$3.15 \cdot 10^{+4}$
10^{-10}	10^{-8}	59830	12923280	8	$5.4 \cdot 10^{-9}$	$1.3 \cdot 10^{-8}$	$1.0326 \cdot 10^{+3}$	$1.25 \cdot 10^{+4}$
10^{-10}	10^{-10}	295100	63741600	9	$5.2 \cdot 10^{-9}$	$1.1 \cdot 10^{-8}$	$5.3102 \cdot 10^{+3}$	$1.20 \cdot 10^{+4}$
				eighth-order force approximation				
10^{-6}	10^{-6}	4894	2505728	7	$4.8 \cdot 10^{-6}$	$1.5 \cdot 10^{-5}$	$1.4287 \cdot 10^{+1}$	$1.75 \cdot 10^{+5}$
10^{-8}	10^{-6}	4894	2505728	7	$8.0 \cdot 10^{-8}$	$4.3 \cdot 10^{-7}$	$4.1362 \cdot 10^{+1}$	$6.06 \cdot 10^{+4}$
10^{-8}	10^{-8}	12860	6584320	7	$1.5 \cdot 10^{-7}$	$4.5 \cdot 10^{-7}$	$1.0268 \cdot 10^{+2}$	$6.41 \cdot 10^{+4}$
10^{-10}	10^{-8}	12860	6584320	7	$6.0 \cdot 10^{-9}$	$1.5 \cdot 10^{-8}$	$2.3717 \cdot 10^{+2}$	$2.78 \cdot 10^{+4}$
10^{-10}	10^{-10}	55854	28597248	8	$4.7 \cdot 10^{-9}$	$1.6 \cdot 10^{-8}$	$1.0489 \cdot 10^{+3}$	$2.73 \cdot 10^{+4}$
10^{-12}	10^{-10}	55854	28597248	8	$6.3 \cdot 10^{-9}$	$8.8 \cdot 10^{-9}$	$1.9641 \cdot 10^{+3}$	$1.46 \cdot 10^{+4}$
10^{-12}	10^{-12}	132490	67834880	9	$5.6 \cdot 10^{-9}$	$7.0 \cdot 10^{-9}$	$4.4599 \cdot 10^{+3}$	$1.52 \cdot 10^{+4}$

Table 7. Free-space Stokes equation (Example 5): Gaussian bump at the origin.

with solution

$$u(x) = \frac{2L}{\mu} \sum_{i=0}^{8} e^{-L\|x - x_i\|^2} \big(\nabla \times (x - x_i) \big),$$

for $x_i = \left(\pm \frac{3}{40}, \pm \frac{3}{40}, \pm \frac{3}{40} \right)$, $\mu = 1$, $L = 125$, inside of the $[-1, 1]^3$ box. Errors are again similar to those seen in the fast-decaying experiments from examples 1 and 4; timings are worse, as expected, since we are dealing with nine times as many degrees of freedom per point. Results are shown in Table 7.

Example 6. It is straightforward to extend the solver infrastructure described above to the case of periodic boundary conditions, using the classical method of images of

Lord Rayleigh [54], following the discussion of [24]. The influence of all separated image boxes can be incorporated using either a recursive approach [38] or a scheme based on lattice sums [30]. In either case, the additional work depends only on ϵ_{fmm} and not on the number of degrees of freedom. The main difference is that the unit cell B now has near neighbors, whose influence must be accounted for. This, too, has relatively little impact on performance. A small number of additional boxes are added to both the interaction and near neighbor lists, but no additional data structures are created; instead, everything is handled via careful book-keeping to minimize additional memory consumption.

As an example, we consider the periodic source function

$$f(\boldsymbol{x}) = C M^2 \pi^2 \sin(\pi M x) \sin(\pi M y) \sin(\pi M z),$$

for which the solution is

$$u(\boldsymbol{x}) = C \sin(\pi M x) \sin(\pi M y) \sin(\pi M z).$$

We conduct our experiments for a nontrivial oscillatory force, choosing $C = 5$ and $M = 7$ on the domain $[-1, 1]^3$ with varying degrees of depth and precision. We check the relative L_2 and L_∞ with results shown in Table 8.

ϵ_{fmm}	ϵ_{rhs}	M_ℓ	N_{pts}	L_T	E_2	E_∞	T_{FMM}	rate
\multicolumn			fourth-order force approximation					
10^{-2}	10^{-2}	32768	2097152	6	$2.6 \cdot 10^{-2}$	$3.3 \cdot 10^{-2}$	$3.3250 \cdot 10^{-1}$	$6.31 \cdot 10^{+6}$
10^{-4}	10^{-4}	262144	16777216	7	$2.9 \cdot 10^{-4}$	$7.6 \cdot 10^{-4}$	$1.6296 \cdot 10^{+1}$	$1.03 \cdot 10^{+6}$
10^{-6}	10^{-6}	2097152	134217728	8	$5.6 \cdot 10^{-6}$	$1.9 \cdot 10^{-5}$	$2.9403 \cdot 10^{+2}$	$4.56 \cdot 10^{+5}$
			sixth-order force approximation					
10^{-2}	10^{-2}	32768	7077888	6	$2.7 \cdot 10^{-2}$	$3.4 \cdot 10^{-2}$	$6.0730 \cdot 10^{-1}$	$1.17 \cdot 10^{+7}$
10^{-4}	10^{-4}	37248	8045568	7	$2.7 \cdot 10^{-4}$	$8.2 \cdot 10^{-4}$	$1.7895 \cdot 10^{+0}$	$4.50 \cdot 10^{+6}$
10^{-6}	10^{-6}	262144	56623104	7	$5.5 \cdot 10^{-6}$	$1.8 \cdot 10^{-5}$	$3.6250 \cdot 10^{+1}$	$1.56 \cdot 10^{+6}$
10^{-8}	10^{-8}	2097152	452984832	8	$1.0 \cdot 10^{-7}$	$3.0 \cdot 10^{-7}$	$7.2101 \cdot 10^{+2}$	$6.28 \cdot 10^{+5}$
			eighth-order force approximation					
10^{-2}	10^{-2}	4096	2097152	5	$1.8 \cdot 10^{-2}$	$2.8 \cdot 10^{-2}$	$2.5445 \cdot 10^{-1}$	$8.24 \cdot 10^{+6}$
10^{-4}	10^{-4}	32768	16777216	6	$3.7 \cdot 10^{-4}$	$9.5 \cdot 10^{-4}$	$2.9459 \cdot 10^{+0}$	$5.70 \cdot 10^{+6}$
10^{-6}	10^{-6}	242432	124125184	7	$6.5 \cdot 10^{-6}$	$2.0 \cdot 10^{-5}$	$5.4881 \cdot 10^{+1}$	$2.26 \cdot 10^{+6}$
10^{-8}	10^{-8}	262144	134217728	7	$1.4 \cdot 10^{-7}$	$5.3 \cdot 10^{-7}$	$1.3578 \cdot 10^{+2}$	$9.88 \cdot 10^{+5}$
10^{-10}	10^{-10}	2097152	1073741824	8	$9.0 \cdot 10^{-9}$	$4.3 \cdot 10^{-8}$	$1.8685 \cdot 10^{+3}$	$5.75 \cdot 10^{+5}$

Table 8. Periodic boundary conditions: Example 6.

It is straightforward to extend this approach to a variety of homogeneous Dirichlet, Neumann or mixed boundary conditions by the method of images as well with very little additional effort.

Example 7. A number of applications require the modeling of source distributions that contain both a smooth component and a singular component. In electrostatics, for example, positively charged ions are often approximated as point charges and the neutralizing electrons as an We consider such a case here. The relevant Poisson equation takes the form

$$\Delta u(\boldsymbol{x}) = f_{\text{smooth}}(\boldsymbol{x}) + \sum_{i=1}^{N} q_i \delta(\boldsymbol{x} - \boldsymbol{x}_i),$$

where the q_i are positive and the neutralizing background takes the form of a sum of Gaussian distributions

$$f^i_{\text{smooth}}(\boldsymbol{x}) = 1/(\sqrt{2\pi\sigma^2})e^{-(\boldsymbol{x}-\sigma)^2/2\sigma^2}$$

centered on each δ-function.

The smooth portion can be handled as above, while the particle sources can be handled with the corresponding particle-based kernel-independent FMM [61].

N_{pts}	ϵ_{fmm}	M_ℓ	L_T	S2M/M2M	Near	M2L	L2L/L2T	T_{FMM}
$1.024 \cdot 10^{+6}$	10^{-2}	4096	5	$2.52 \cdot 10^{-2}$	$6.42 \cdot 10^{-1}$	$5.09 \cdot 10^{-2}$	$4.84 \cdot 10^{-3}$	$7.23 \cdot 10^{-1}$
	10^{-4}	4096	5	$6.97 \cdot 10^{-2}$	$6.49 \cdot 10^{-1}$	$1.29 \cdot 10^{-1}$	$2.35 \cdot 10^{-2}$	$8.71 \cdot 10^{-1}$
	10^{-6}	4096	5	$1.41 \cdot 10^{-1}$	$6.50 \cdot 10^{-1}$	$3.77 \cdot 10^{-1}$	$6.66 \cdot 10^{-2}$	$1.24 \cdot 10^{+0}$
	10^{-8}	4096	5	$2.48 \cdot 10^{-1}$	$6.41 \cdot 10^{-1}$	$1.03 \cdot 10^{+0}$	$1.58 \cdot 10^{-1}$	$2.08 \cdot 10^{+0}$
$4.096 \cdot 10^{+6}$	10^{-2}	32768	6	$1.11 \cdot 10^{-1}$	$1.58 \cdot 10^{+0}$	$2.39 \cdot 10^{-1}$	$1.85 \cdot 10^{-2}$	$1.95 \cdot 10^{+0}$
	10^{-4}	32768	6	$3.04 \cdot 10^{-1}$	$1.61 \cdot 10^{+0}$	$1.32 \cdot 10^{+0}$	$1.01 \cdot 10^{-1}$	$3.34 \cdot 10^{+0}$
	10^{-6}	32768	6	$6.29 \cdot 10^{-1}$	$1.61 \cdot 10^{+0}$	$3.55 \cdot 10^{+0}$	$2.92 \cdot 10^{-1}$	$6.08 \cdot 10^{+0}$
	10^{-8}	32768	6	$1.17 \cdot 10^{+0}$	$1.60 \cdot 10^{+0}$	$9.20 \cdot 10^{+0}$	$6.48 \cdot 10^{-1}$	$1.26 \cdot 10^{+1}$
$1.638 \cdot 10^{+7}$	10^{-2}	262144	7	$1.58 \cdot 10^{+0}$	$4.31 \cdot 10^{+1}$	$2.14 \cdot 10^{+0}$	$2.58 \cdot 10^{-1}$	$4.71 \cdot 10^{+1}$
	10^{-4}	262144	7	$4.31 \cdot 10^{+0}$	$4.34 \cdot 10^{+1}$	$1.14 \cdot 10^{+1}$	$1.48 \cdot 10^{+0}$	$6.06 \cdot 10^{+1}$
	10^{-6}	262144	7	$8.79 \cdot 10^{+0}$	$4.38 \cdot 10^{+1}$	$3.05 \cdot 10^{+1}$	$4.20 \cdot 10^{+0}$	$8.72 \cdot 10^{+1}$
	10^{-8}	262144	7	$1.55 \cdot 10^{+1}$	$4.32 \cdot 10^{+1}$	$8.26 \cdot 10^{+1}$	$9.96 \cdot 10^{+0}$	$1.51 \cdot 10^{+2}$
$6.554 \cdot 10^{+7}$	10^{-2}	2097152	8	$7.00 \cdot 10^{+0}$	$1.04 \cdot 10^{+2}$	$1.80 \cdot 10^{+1}$	$1.40 \cdot 10^{+0}$	$1.31 \cdot 10^{+2}$
	10^{-4}	2097152	8	$1.93 \cdot 10^{+1}$	$1.05 \cdot 10^{+2}$	$9.14 \cdot 10^{+1}$	$6.76 \cdot 10^{+0}$	$2.23 \cdot 10^{+2}$
	10^{-6}	2097152	8	$4.01 \cdot 10^{+1}$	$1.07 \cdot 10^{+2}$	$2.61 \cdot 10^{+2}$	$1.93 \cdot 10^{+1}$	$4.27 \cdot 10^{+2}$
	10^{-8}	2097152	8	$7.34 \cdot 10^{+1}$	$1.06 \cdot 10^{+2}$	$7.00 \cdot 10^{+2}$	$4.19 \cdot 10^{+1}$	$9.21 \cdot 10^{+2}$

Table 9. Example 7: Poisson equation with a mixture of smooth and singular sources.

However, it is trivial to modify our solver to incorporate the particle sources into the S2M operator of Section 4.2 by modifying (22):

$$\boldsymbol{K}_{\text{S2M}}^{B}\boldsymbol{\phi}^{B,u} = \boldsymbol{F}_{\text{S2M}}^{B}\gamma^{B} + \sum_{i=1}^{N} q_i G(\boldsymbol{x}, \boldsymbol{x}_i), \tag{41}$$

where G is the kernel used for evaluating the singular component, consisting of the point charges. Once the point charges are incorporated into $\boldsymbol{\phi}^{B,u}$, the rest of the components of the algorithm (i.e., M2M, M2L, and L2L) take care of the far-field interactions. We need only calculate the influence of near-field particle interactions directly, and evaluate both the local expansions (L2T) and the smooth contributions at particle locations. The latter is done by interpolation, as discussed in Section 4.5.

The performance of our scheme is shown in Table 9 on the previous page.

The upward pass timings are minimally larger than for the particle-only case, and the downward pass timings are agnostic about the nature of the sources, dependent only on the tree-structure itself. The Near computation timings are simply the sum of the particle and volume-based cases, since there is no amortization of cost in this step.

7. Conclusions

We have presented a kernel-independent FMM for solving a variety of constant-coefficient elliptic PDEs in free space, allowing for arbitrary levels of adaptivity, highly nonhomogeneous forces and arbitrarily distributed target locations. Results for the Poisson, modified Helmholtz, and Stokes equations show that the performance is similar for each. Applying the method to other equations requires only a kernel evaluation routine.

Compared to the state-of-the-art technique [47], our method is accurate to higher order and therefore solves similar problems with fewer degrees of freedom, and the work per point is approximately the same. Our current implementation uses OpenMP (but not MPI), although we expect the extension to be straightforward, as our solver is built on top of the MPI-based code of [61]. We discuss how the major loops are optimized for OpenMP shared-memory parallelization in Section A.4 for this current implementation.

As in [24], we have extended our solver to handle periodic, Dirichlet and Neumann boundary conditions for problems on cubic domains using the method of images. We are also coupling the present volume integral code with boundary integral methods to allow for the solution of linear, constant-coefficient, inhomogeneous elliptic PDEs in complex geometries, as in [9]. Additional current work involves incorporating this solver into the state-of-the-art in [43]. These extensions will be reported at a later date.

Appendix

In this appendix, we first verify numerically that the equivalent density representation yields the expected accuracy, followed by a discussion of how the choice of grids affects the order of convergence and an overall numerical justification for the use of Tikhonov regularization. We close with a discussion of how we accelerate the major computation loops of our algorithm using OpenMP shared-memory parallelization and load-balancing techniques.

A.1. *Equivalent density accuracy.* As discussed in Section 4.1, we invert several matrices of discretized Fredholm equations of the first kind in order to build out far-field representations,

$$K^{y_d,x_d}\phi_d = K^{y_s,x_d}\phi_s.$$

As in [60], we choose to use Tikhonov regularization [41] when solving these ill-conditioned systems. This solves two problems: in this way, we eliminate the null space in the cases when it is present (Stokes kernel) and we significantly improve accuracy of the inversion for higher numbers of samples (Section A.3). We verify the potential we get from $\phi^{B,u}$, computed using our regularized method in the S2M operation, approximates well $u(x)$, computed directly from a force. We test using $g^B = \sum_{a,b,c}^{(a+b+c)\le(k-1)} x^a y^b z^c$ for box B of width 2 and compute

$$u(x) = \int_B K(x,y)g^B(y)dy, \quad x \in x^{B,u},$$

to within 10^{-16} accuracy using adaptive Gaussian quadrature [8]. We then compute $\phi^{B,u}$ at $y^{B,u}$ using (23) where $(K_{S2M})^{-1}$ is replaced with

$$(\alpha I + (K_{S2M})^* K_{S2M})^{-1} K_{S2M}^*.$$

For FMM precision, n_p, we choose $\alpha = 10^{-(n_p+1)}$. More details on the choice of α are available in [60]. Our algorithm relies on the fact that for surfaces outside the near field of B, $\phi^{B,u}$ is a sufficiently accurate representation of B's volume force. We compute

$$u(x)_{\text{equiv}} = \int_{y^{B,u}} K(x,y)\phi^{B,u}(y)\,dy$$

for $x \in S$, some surface. To evaluate the accuracy of this approximation, we compute

$$u(x)_{\text{exact}} = \int_B K(x,y)g^B(y)\,dy$$

up to an accuracy of 10^{-16} [8]. In Figure 7, we compare the infinity-norm of the resulting error for three different kernels (Laplace, modified Helmholtz, Stokes) and varying levels of the polynomial approximation and multiple degrees of FMM
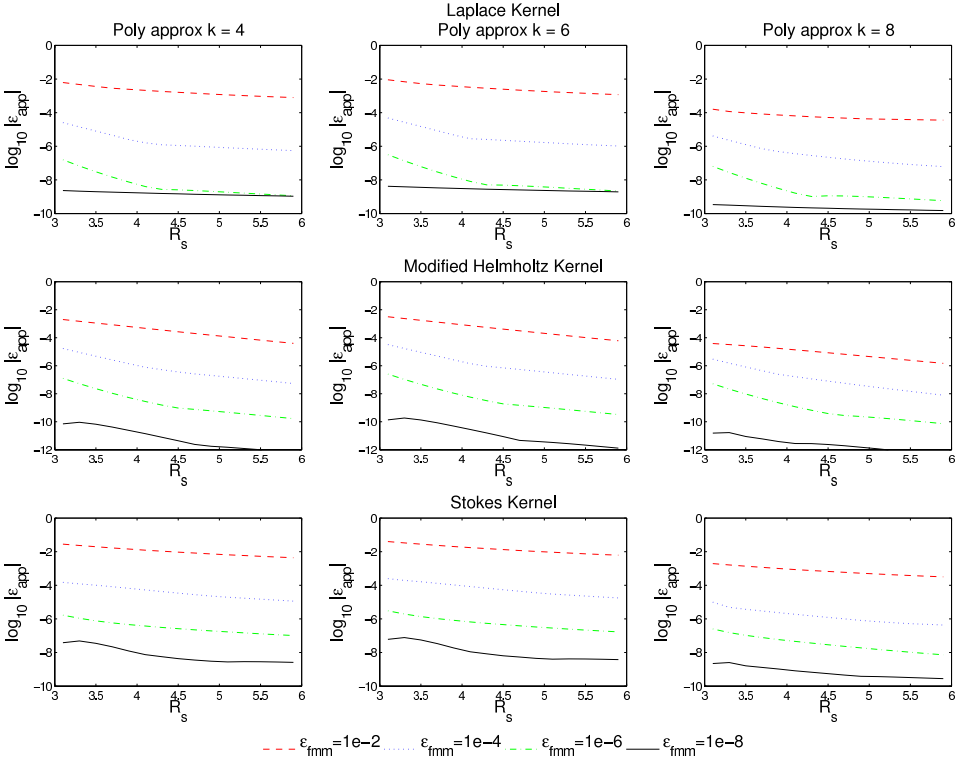
Figure 7. Error due to upward equivalent density approximation of the field. From left to right, three columns show the errors for the polynomial force approximations of degree 4, 6 and 8. Each plot shows four levels of FMM precision, $\epsilon_{\text{fmm}} = 10^{-n_p}$, $p = n_p^3 - (n_p - 2)^3$ points are used on the surfaces $\mathbf{y}^{B,u}$ and $\mathbf{x}^{B,u}$. For the evaluation surfaces S, we vary the radius R_S from 3.1 to 5.9, the region covering $L_I^B \in \mathcal{F}^B$. The y-axis of each plot is the infinity norm $\|u_{\text{equiv}} - u_{\text{exact}}\|_\infty$ computed over 488 samples on S.

evaluation precision. For each of the kernels of interest, $\phi^{B,u}$, computed by inverting our ill-conditioned kernels, is recovered on each surface S to within the requested degree of precision. For evaluating the accuracy of the kernel inversion and regularization in the computation of $\phi^{B,d}$, we note that this computation is equivalent to the particle-based FMM, of which numerical analysis for the M2L and L2L operators is available in [60].

A.2. *Polynomial basis and grid spacing.* As discussed in Section 4.4, we evaluate the solution at a leaf box B on a grid of points $\mathbf{x}^{B,g}$, and construct an approximating polynomial from these points. Additionally, we construct a k-th order polynomial

Figure 8. For each of the test examples, the x-axis indicates the negative log of the requested FMM accuracy, ϵ_{fmm}, and the y-axis indicates the log of E_2. The number of points chosen for each ϵ_{fmm} is equivalent to those in Example 1 (page 102) for $\epsilon_{\text{rhs}} = \epsilon_{\text{fmm}}$. Left: for polynomial approximation of degree 4 and $x^{B,g}$ of size 4^3 on each leaf B, overall relative error is close for equispaced and Chebyshev points. Middle: For $n, k = 6$ differences are visible but insignificant. Right: For $n, k = 8$, solutions based on equispaced grid are less accurate.

approximation to B's distributed force if g^B is given on a grid (see Section 4.6). For consistency with AMR codes and efficiency of implementation, it would have been desirable to use uniform grid samples. This approach works well for $n \leq 6$, but it is well-known for large n that equispaced grids lead to instabilities [57]; as a result, for $n > 6$ we use Chebyshev grid points. To show that regularly spaced grid points perform poorly for $n, k > 6$, we consider the following test case:

$$-\Delta u(\boldsymbol{x}) = e^{-L(\|\boldsymbol{x}\|_2)^2}(4L(\|\boldsymbol{x}\|_2)^2 - 6L), \quad L = 250, \boldsymbol{x} \in [-1, 1]^3.$$

In Figure 8, we compare the overall relative L^2 error, E_2, for solutions using equispaced and Chebyshev grid points in the evaluation of the solution and construction of the polynomial approximations of degree 4, 6 and 8. Errors for discretizations using equispaced or Chebyshev grid points are similar for $k \leq$, but for $k = 8$, Chebyshev points are more accurate.

A.3. *Tikhonov regularization.* As discussed in Section A.1, we use Tikhonov regularization [41] to invert Fredholm equations of the first kind, specifically the S2M, M2M, and L2L operators in Section 4. Further, in Section A.1, we looked specifically at the accuracy resulting from this inversion process. To justify the overall use of Tikhonov regularization, we consider for the Poisson equation the test case

$$-\Delta u(\boldsymbol{x}) = e^{-L(\|\boldsymbol{x}\|_2)^2}(4L(\|\boldsymbol{x}\|_2)^2 - 6L), \quad L = 250, \boldsymbol{x} \in [-1, 1]^3,$$

Figure 9. For each of the test examples, the x-axis indicates the negative log of the requested FMM accuracy, ϵ_{fmm}, and the y-axis indicates the log of E_2. The number of points chosen for each ϵ_{fmm} is equivalent to those in Examples 1 and 5 of Section 6 for $\epsilon_{\text{rhs}} = \epsilon_{\text{fmm}}$. Left: for polynomial approximation of degree 6 for the Laplace kernel with and without regularization. Right: for polynomial approximation of degree 6 for the Stokes kernel with and without regularization.

and for the Stokes equation

$$-\Delta u(\boldsymbol{x}) + \nabla p(\boldsymbol{x}) = \big(8L^3\|\boldsymbol{x} - \boldsymbol{x}_i\|^2 - 20L^2\big)e^{-L\|\boldsymbol{x}-\boldsymbol{x}_i\|^2}\big(\nabla \times (\boldsymbol{x} - \boldsymbol{x}_i)\big),$$
$$L = 125, \ \boldsymbol{x} \in [-1, 1]^3.$$

In Figure 9, we compare the overall relative L^2 error, E_2, solutions, resulting from Tikhonov regularization versus no regularization and the construction of polynomial approximations of degree $k = 6$ for the right-hand sides (errors for $k = 4, 8$ are similar). For decreasing levels of ϵ_{fmm}, we choose $\epsilon_{\text{fmm}} = \epsilon_{\text{rhs}}$.

We notice that for $\epsilon_{\text{fmm}} > 10^{-7}$, the effect of not employing regularization is equivalent to using regularization for both the Laplace and Stokes operators. However, as ϵ_{fmm} decreases, the number of sample points on the equivalent and check surfaces increases, resulting in larger linear systems, which as mentioned earlier, may be poorly conditioned. Indeed, for such larger systems resulting from $\epsilon_{\text{fmm}} \leq 10^{-7}$, it is necessary to regularize the systems to achieve desirable results.

A.4. *Shared-memory parallelization and load-balancing.* We have designed the code to take advantage of shared-memory architectures through the use of OpenMP (see Section 7). In particular, we highlight the steps to accelerate the various major steps in Algorithm 1. For details on the nature of OpenMP and its usage, see [19].

S2M *and* M2M *computations.* In the upward pass (step 2 of Algorithm 1 and Section 4.2), we begin by building a list of all leaf boxes, B in the octree T, which have sources. We then do a simple OpenMP parallelization step over these boxes

for the S2M step. As all components of (23) are of the same size for each leaf box, there is no need to rebalance the load among threads.

In order to ensure proper order of computation, we proceed by sorting all non-leaves in reverse-order by depth. For each nonleaf level in T, beginning at the deepest level, we translate a box B's children's upward equivalent densities to its own through the M2M computation in (25). Again, as each of the components is of the same size, there is no need to rebalance among threads. As we parallelize only among boxes at the same depth in T, level ℓ is not processed until $\ell + 1$ has completed. Further, once we have reached coarse level $\ell = 1$ (which only occurs for periodic or Dirichlet boundary conditions), we discontinue the parallelization.

M2L, L2L, and L2T *computations.* In the downward pass of Algorithm 1 (see Section 4.3), we perform a similar operation as above for the M2M step. First, we sort all boxes B in T from the shallowest to deepest levels in the tree. For each level, ℓ, we parallelize among the boxes being processed at that level for the M2L and L2L computations. The L2L components in (29) are of equivalent size for each box B; however, for each box B, the size of L_V^B vary greatly from other boxes (for example, this list is much smaller for boxes on the edge or corners of our domain). To ensure proper balancing among threads, we further sort all boxes for each level, ℓ by the size of L_V^B and then reorder the boxes such that the sum of all L_V^B for each thread is of roughly the same size.

For the L2T computations in (31), we once again build a list of only leaf boxes, for which the target solution is desired, and we parallelize the computations in this list. The components of the discretized equation are all the same size, as with the S2M computation, so there is no need to rebalance among threads for this step.

Near-field computations. We focus our discussion here on the U-list computations. Parallelizing the near-field computations in (34) is the most straightforward in that no leaf box B is dependent on the completion of computations by any other box. That is, we can simply parallelize the computations among leaf boxes, for which the L_U^B exists. However, even more so than with the M2L computations, the sizes of L_U^B can be very different among leaf boxes (especially in the most adaptively refined octrees). Thus, we sort all leaf boxes B in T by the size of L_U^B and reorder the list of leaves such that the sum of the size of L_U^B among each thread is roughly equivalent, ensuring a relatively well-balanced load among threads. The size of the components and operators are the same for each box B, so balancing by list sizes is optimal. We note that this rebalancing is largely unnecessary for uniformly refined trees.

Additionally, for matrix kernels (e.g., Stokes) and larger orders of polynomial approximation, constantly loading large matrices into memory results in little speedup as we increase the number of processes. To correct this, for each equivalence class as described in Section 5, we perform all of the operations involving a single

class first before performing all computations for other classes of operators. Hence, we only load each matrix operator at most once per processor.

We note that for the M2L step, as we have to process level ℓ before moving to level $\ell - 1$, operators will constantly have to be reloaded. Performing all computations in order for each equivalence class at each level is done, but we have seen little time savings for this in practice as opposed to the near-field computations, where it is essential for good speedup.

Remark. As with the near-field computations, for adaptively refined trees, we rebalance the loads among threads for the X and W lists, which involve additional near-field S2T M2T, and S2L computations in Equations (39) and (38), based on the sizes of L_X^B and L_W^B, respectively. Additionally, we perform all computations in order of equivalence class, again loading each matrix operator at most once.

Timing results versus number of processors. To see the effect of our use of OpenMP and load-balancing strategies, we investigate the strong scaling of two fixed problems. First, in Example 1 (page 102), we set the polynomial order, ϵ_{rhs}, and ϵ_{fmm} to 8. The reasoning behind this is to ensure that for a single processor, neither the near-field nor far-field computations fully dominate the timings. In Table 10

N_{procs}	S2M/M2M	Near	M2L	L2L/L2T	T_{FMM}	scaling rate
Poisson equation (Example 1) $\epsilon_{\text{rhs}} = \epsilon_{\text{fmm}} = 8$ $M_\ell = 5440$ $N_{\text{pts}} = 2785280$						
1	$1.125 \cdot 10^{+0}$	$8.534 \cdot 10^{+0}$	$1.464 \cdot 10^{+1}$	$9.340 \cdot 10^{-1}$	$2.523 \cdot 10^{+1}$	
2	$5.880 \cdot 10^{-1}$	$5.112 \cdot 10^{+0}$	$7.285 \cdot 10^{+0}$	$4.779 \cdot 10^{-1}$	$1.346 \cdot 10^{+1}$	$1.874 \cdot 10^{+0}$
4	$3.750 \cdot 10^{-1}$	$2.377 \cdot 10^{+0}$	$4.559 \cdot 10^{+0}$	$2.927 \cdot 10^{-1}$	$7.604 \cdot 10^{+0}$	$1.770 \cdot 10^{+0}$
8	$1.838 \cdot 10^{-1}$	$1.231 \cdot 10^{+0}$	$2.279 \cdot 10^{+0}$	$1.459 \cdot 10^{-1}$	$3.841 \cdot 10^{+0}$	$1.979 \cdot 10^{+0}$
16	$9.700 \cdot 10^{-2}$	$6.894 \cdot 10^{-1}$	$1.240 \cdot 10^{+0}$	$8.546 \cdot 10^{-2}$	$2.112 \cdot 10^{+0}$	$1.818 \cdot 10^{+0}$
Stokes equations (Example 5) $\epsilon_{\text{rhs}} = \epsilon_{\text{fmm}} = 6$ $M_\ell = 4894$ $N_{\text{pts}} = 2505728$						
1	$8.794 \cdot 10^{+0}$	$5.416 \cdot 10^{+1}$	$9.402 \cdot 10^{+1}$	$6.832 \cdot 10^{+0}$	$1.638 \cdot 10^{+2}$	
2	$5.182 \cdot 10^{+0}$	$2.857 \cdot 10^{+1}$	$5.213 \cdot 10^{+1}$	$3.617 \cdot 10^{+0}$	$8.951 \cdot 10^{+1}$	$1.830 \cdot 10^{+0}$
4	$2.866 \cdot 10^{+0}$	$1.307 \cdot 10^{+1}$	$3.029 \cdot 10^{+1}$	$1.733 \cdot 10^{+0}$	$4.797 \cdot 10^{+1}$	$1.865 \cdot 10^{+0}$
8	$1.569 \cdot 10^{+0}$	$6.781 \cdot 10^{+0}$	$1.613 \cdot 10^{+1}$	$8.367 \cdot 10^{-1}$	$2.532 \cdot 10^{+1}$	$1.894 \cdot 10^{+0}$
16	$8.248 \cdot 10^{-1}$	$3.611 \cdot 10^{+0}$	$9.452 \cdot 10^{+0}$	$3.918 \cdot 10^{-1}$	$1.428 \cdot 10^{+1}$	$1.772 \cdot 10^{+0}$

Table 10. Timings (in wall-time seconds) for the various components of the FMM volume solver for two fixed problem sizes. The tree level, L_T, is 7 and the polynomial order is 8 and in each case. N_{procs}, M_ℓ, and N_{pts} are the number of processors, leaves, and points; we scale N_{procs} linearly. We separate the S2M/M2M, Near (U, W, X-list computations), M2L (V-list computations), and L2L/L2T timings, with the total shown as T_{FMM}.

Figure 10. Log-log plots for timings from Table 10.

(top part) we look at the timings for the different algorithmic steps (note that the near-field computation times include U, W, and X list computation times) and we plot the decreasing times in Figure 10, top.

For our second study of the effect of shared-memory parallelization, we look at the Stokes kernel tests from Example 5 (page 107). We fix the polynomial order at

8 and look at $\epsilon_{\text{rhs}} = \epsilon_{\text{fmm}} = 6$, again in an effort to not have one step fully dominate the computational time, allowing us to look at the effect of scaling the number of processors. Timing results can be seen in the bottom parts of Table 10 and Figure 10.

As can be seen in Table 10, our scheme exhibits the desirable result of nearly linear speedup as we scale the number of processors. As indicated in the conclusion, current work is being done to incorporate this work with [43] in order to achieve parallelization on a significantly larger scale.

Acknowledgement

References

[1]   M. F. Adams and J. Demmel, *Parallel multigrid solver algorithms and implementations for 3D unstructured finite element problem*, Internat. J. Numer. Methods Engrg. **48**, no. 8, 1241–1262.

[2]   M. J. Aftosmis, M. J. Berger, and J. E. Melton, *Adaptive Cartesian mesh generation*, The handbook of grid generation (J. F. Thompson, ed.), CRC Press, Boca Raton, FL, 1998, pp. 22–1–22–26.

[3]   C. R. Anderson, *A method of local corrections for computing the velocity field due to a distribution of vortex blobs*, J. Comput. Phys. **62** (1986), no. 1, 111–123. MR 87d:76050 Zbl 0575. 76031

[4]   G. T. Balls and P. Colella, *A finite difference domain decomposition method using local corrections for the solution of Poisson's equation*, J. Comput. Phys. **180** (2002), no. 1, 25–53. MR 2003c:65102 Zbl 1003.65140

[5]   J. Barnes and P. Hut, *A hierarchical $O(N \log N)$ force calculation algorithm*, Nature **324** (1986), 446–449.

[6]   R. Beatson and L. Greengard, *A short course on fast multipole methods*, Wavelets, multilevel methods and elliptic PDEs (M. Ainsworth et al., eds.), Clarendon Press, Oxford, 1997, pp. 1–37. MR 99a:65142 Zbl 0882.65106

[7]   M. J. Berger, M. Aftosmis, and J. Melton, *Accuracy, adaptive methods and complex geometry*, Proc. 1st AFOSR Conference on Dynamic Motion CFD (L. Sakell and D. Knight, eds.), 1996.

[8]   J. Berntsen, T. O. Espelid, and A. Genz, *Algorithm 698: DCUHRE: an adaptive multidimensional integration routine for a vector of integrals*, ACM Trans. Math. Software **17** (1991), no. 4, 452–456. MR 1140035 Zbl 0900.65053

[9]   G. Biros, L. Ying, and D. Zorin, *A fast solver for the Stokes equations with distributed forces in complex geometries*, J. Comput. Phys. **193** (2004), no. 1, 317–348. MR 2022697 Zbl 1047. 76065

---

[2]http://www.nyu.edu/its/research/hpc

[10] S. Börm, H^2-*matrix arithmetics in linear complexity*, Computing **77** (2006), no. 1, 1–28. MR 2006k:65111 Zbl 1086.65036

[11] S. Börm and W. Hackbusch, *Hierarchical quadrature for singular integrals*, Computing **74** (2005), no. 2, 75–100. MR 2006c:41037 Zbl 1003.65140

[12] A. H. Boschitsch, M. O. Fenley, and W. K. Olson, *A fast adaptive multipole algorithm for calculating screened Coulomb (Yukawa) interactions*, J. Comput. Phys. **151** (1999), no. 1, 212–241. MR 1701576 Zbl 1017.92500

[13] A. Brandt, *Multi-level adaptive solutions to boundary-value problems*, Math. Comp. **31** (1977), no. 138, 333–390. MR 55 #4714 Zbl 0373.65054

[14] W. L. Briggs, V. E. Henson, and S. F. McCormick, *A multigrid tutorial*, 2nd ed., Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, 2000. MR 2001h:65002 Zbl 0958.65128

[15] B. L. Buzbee, G. H. Golub, and C. W. Nielson, *On direct methods for solving Poisson's equations*, SIAM J. Numer. Anal. **7** (1970), 627–656. MR 44 #4920 Zbl 0217.52902

[16] C. Canuto, M. Y. Hussaini, A. Quarteroni, and T. A. Zang, *Spectral methods in fluid dynamics*, Springer, New York, 1988. MR 89m:76004 Zbl 0658.76001

[17] T. F. Chan, R. Glowinski, J. Périaux, and O. B. Widlund, *Domain decomposition methods*, SIAM, Philadelphia, 1989. MR 89j:65010

[18] T. F. Chan and B. F. Smith, *Domain decomposition and multigrid algorithms for elliptic problems on unstructured meshes*, Electron. Trans. Numer. Anal. **2** (1994), no. Dec., 171–182. MR 95i:65173 Zbl 0852.65108

[19] B. Chapman, G. Jost, and R. Pas, *Using openmp: Portable shared memory parallel programming (scientific and engineering computation)*, (2007).

[20] H. Cheng, L. Greengard, and V. Rokhlin, *A fast adaptive multipole algorithm in three dimensions*, J. Comput. Phys. **155** (1999), no. 2, 468–498. MR 2000h:65178 Zbl 0937.65126

[21] H. Cheng, W. Y. Crutchfield, Z. Gimbutas, L. F. Greengard, J. F. Ethridge, J. Huang, V. Rokhlin, N. Yarvin, and J. Zhao, *A wideband fast multipole method for the Helmholtz equation in three dimensions*, J. Comput. Phys. **216** (2006), no. 1, 300–325. MR 2007a:65193 Zbl 1093.65117

[22] H. Cheng, J. Huang, and T. J. Leiterman, *An adaptive fast solver for the modified Helmholtz equation in two dimensions*, J. Comput. Phys. **211** (2006), no. 2, 616–637. MR 2006e:65242 Zbl 1117.65161

[23] G. Chesshire and W. D. Hanshaw, *Composite overlapping meshes for the solution of partial differential equations*, J. Comput. Phys. **90** (1990), no. 1, 1–64. MR 91f:76043 Zbl 0709.65090

[24] F. Ethridge and L. Greengard, *A new fast-multipole accelerated Poisson solver in two dimensions*, SIAM J. Sci. Comput. **23** (2001), no. 3, 741–760. MR 2002i:65146 Zbl 1002.65131

[25] L. Greengard, *The rapid evaluation of potential fields in particle systems*, MIT Press, Cambridge, MA, 1988. MR 89k:31008 Zbl 1001.31500

[26] ———, *Fast algorithms for classical physics*, Science **265** (1994), no. 5174, 909–914. MR 95f:65236

[27] L. Greengard and J. Huang, *A new version of the fast multipole method for screened Coulomb interactions in three dimensions*, J. Comput. Phys. **180** (2002), no. 2, 642–658. MR 2003h:78014 Zbl 1143.78372

[28] L. Greengard, M. C. Kropinski, and A. Mayo, *Integral equation methods for Stokes flow and isotropic elasticity in the plane*, J. Comput. Phys. **125** (1996), no. 2, 403–414. MR 97a:73022 Zbl 0847.76066

[29] L. Greengard and J.-Y. Lee, *A direct adaptive Poisson solver of arbitrary order accuracy*, J. Comput. Phys. **125** (1996), no. 2, 415–424. MR 96m:65090 Zbl 0851.65090

[30] L. Greengard and V. Rokhlin, *A fast algorithm for particle simulations*, J. Comput. Phys. **73** (1987), no. 2, 325–348. MR 88k:82007 Zbl 0629.65005

[31] ———, *The rapid evaluation of potential fields in three dimensions*, Vortex methods: Proceedings of the U.C.L.A. Workshop (C. Anderson and C. Greengard, eds.), Lecture Notes in Math., no. 1360, Springer, Berlin, 1988, pp. 121–141. MR 979565 Zbl 0661.70006

[32] L. Greengard and V. Rokhlin, *A new version of the fast multipole method for the Laplace equation in three dimensions*, Acta Numer., no. 6, Cambridge Univ. Press, 1997, pp. 229–269. MR 99c:65012 Zbl 0889.65115

[33] N. A. Gumerov and R. Duraiswami, *Fast multipole method for the biharmonic equation in three dimensions*, J. Comput. Phys. **215** (2006), no. 1, 363–383. MR 2006j:65373 Zbl 1103.65122

[34] W. Hackbusch, *A sparse matrix arithmetic based on H-matrices. I. Introduction to H-matrices*, Computing **62** (1999), no. 2, 89–108. MR 2000c:65039 Zbl 0927.65063

[35] W. Hackbusch and S. Börm, H^2-*matrix approximation of integral operators by interpolation*, Appl. Numer. Math. **43** (2002), no. 1-2, 129–143. MR 1936106 Zbl 1019.65103

[36] W. Hackbusch and Z. P. Nowak, *On the fast matrix multiplication in the boundary element method by panel clustering*, Numer. Math. **54** (1989), no. 4, 463–491. MR 89k:65162 Zbl 0641.65038

[37] W. Hackbusch and U. Trottenberg (eds.), *Multigrid methods*, Lecture Notes in Mathematics, no. 960, Springer, Berlin, 1982. MR 84b:65007 Zbl 0497.00015

[38] J. Helsing, *Fast and accurate calculations of structural parameters for suspensions*, Proc. Roy. Soc. Lond. A **445** (1994), 127–140.

[39] J. Huang and L. Greengard, *A fast direct solver for elliptic partial differential equations on adaptively refined meshes*, SIAM J. Sci. Comput. **21** (1999/00), no. 4, 1551–1566. MR 2001c:65132 Zbl 0957.65091

[40] H. Johansen and P. Colella, *A Cartesian grid embedded boundary method for Poisson's equation on irregular domains*, J. Comput. Phys. **147** (1998), no. 1, 60–85. MR 99m:65231 Zbl 0923.65079

[41] R. Kress, *Linear integral equations*, 2nd ed., Applied Mathematical Sciences, no. 82, Springer, New York, 1999. MR 2000h:45001 Zbl 0920.45001

[42] O. A. Ladyzhenskaya, *The mathematical theory of viscous incompressible flow*, 2nd ed., Mathematics and its Applications, no. 2, Gordon and Breach, New York, 1969. MR 40 #7610 Zbl 0184.52603

[43] I. Lashuk, A. Chandramowlishwaran, M. Langston, T. Nguyen, R. Sampath, A. Shringarpure, R. Vuduc, L. Ying, D. Zorin, and G. Biros, *A massively parallel adaptive fast-multipole method on heterogeneous architectures*, SC'2009 Conference, IEEE/ACM SIGARCH, 2009.

[44] D. Martin and K. Cartwright, *Solving Poisson's equations using adaptive mesh refinement*, technical report M96/66, Electronic Research Laboratory, University of California, Berkeley, 1996.

[45] D. J. Mavriplis, *Unstructured grid techniques*, Annu. Rev. Fluid Mech. (1997), no. 29, 473–514. MR 97j:76044

[46] A. Mayo, *Fast high order accurate solution of Laplace's equation on irregular regions*, SIAM J. Sci. Statist. Comput. **6** (1985), no. 1, 144–157. MR 86i:65066 Zbl 0559.65082

[47] P. McCorquodale, P. Colella, G. T. Balls, and S. B. Baden, *A local corrections algorithm for solving Poisson's equation in three dimensions*, Commun. Appl. Math. Comput. Sci. **2** (2007), 57–81. MR 2008i:65291 Zbl 1133.65106

[48] A. McKenney, L. Greengard, and A. Mayo, *A fast Poisson solver for complex geometries*, J. Comput. Phys. **118** (1995), no. 2, 348–355. MR 96a:65179 Zbl 0823.65115

[49] M. L. Minion, *A projection method for locally refined grids*, J. Comput. Phys. **127** (1996), no. 1, 158–178. MR 97g:76072 Zbl 0859.76047

[50] G. J. Rodin and Y. Fu, *Fast solution method for three-dimensional Stokesian many-particle problems*, Comm. Numer. Methods Engrg. **16** (2000), no. 2, 145–149.

[51] V. Rokhlin, *Rapid solution of integral equations of classical potential theory*, J. Comput. Phys. **60** (1985), no. 2, 187–207. MR 86k:65120 Zbl 0629.65122

[52] ———, *Rapid solution of integral equations of scattering theory in two dimensions*, J. Comput. Phys. **86** (1990), no. 2, 414–439. MR 90k:76081 Zbl 0686.65079

[53] M. Strain, G. Scuseria, and M. Frisch, *Achieving linear scaling for the electronic quantum Coulomb problem*, Science **271** (1996), 51–53.

[54] J. W. Strutt (Lord Rayleigh), *On the influence of obstacles arranged in rectangular order upon the properties of a medium*, Phil. Mag. **34** (1892), 481–502.

[55] 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. **30** (2008), no. 5, 2675–2708. MR 2010d:68192 Zbl 1186.68554

[56] A.-K. Tornberg and L. Greengard, *A fast multipole method for the three-dimensional Stokes equations*, J. Comput. Phys. **227** (2008), no. 3, 1613–1619. MR 2009g:76110 Zbl 05248608

[57] L. N. Trefethen and D. Bau, III, *Numerical linear algebra*, Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, 1997. MR 98k:65002 Zbl 0874.65013

[58] H. Wang, T. Lei, J. Li, J. Huang, and Z. Yao, *A parallel fast multipole accelerated integral equation scheme for 3D Stokes equations*, Internat. J. Numer. Methods Engrg. **70** (2007), no. 7, 812–839. MR 2008a:76120 Zbl 1194.76221

[59] C. Whitea, B. Johnson, P. M. W. Gill, and M. Head-Gordon, *The continuous fast multipole method*, Chem. Phys. Lett. **230** (1994), 8–16.

[60] L. Ying, G. Biros, and D. Zorin, *A kernel-independent adaptive fast multipole algorithm in two and three dimensions*, J. Comput. Phys. **196** (2004), no. 2, 591–626. MR 2005d:65235 Zbl 1053.65095

[61] L. Ying, G. Biros, D. Zorin, and M. H. Langston, *A new parallel kernel-independent fast multipole method*, SC'2003 Conference CD, IEEE/ACM SIGARCH, 2003.

M. HARPER LANGSTON: harper@cims.nyu.edu
Courant Institute, New York University, 251 Mercer Street, New York 10012, United States
http://cs.nyu.edu/~harper/

LESLIE GREENGARD: greengard@cims.nyu.edu
Courant Institute, New York University, 251 Mercer Street, New York NY 10012, United States
http://math.nyu.edu/faculty/greengar/

DENIS ZORIN: dzorin@cims.nyu.edu
Courant Institute, New York University, 251 Mercer Street, New York 10012, United States
http://mrl.nyu.edu/~dzorin/

Communications in Applied Mathematics and Computational Science

pjm.math.berkeley.edu/camcos

Communications in Applied Mathematics and Computational Science