

# The Block-P<sup>3</sup>M Algorithm for Fast Integral Equation Solvers

**Abstract** — The Block-P<sup>3</sup>M solver calculates action-at-a-distance convolution sums that arise in electromagnetics and, unlike the plane-wave fast multipole method (FMM), it does not suffer from low-frequency breakdown. The solution uses  $\mathcal{O}(N \log N)$  operations and  $\mathcal{O}(N)$  storage for surface as well as volume distributions of sources. It differs from single-mesh P<sup>3</sup>M [1, 2] primarily in that sparse meshes are used and the mesh convolution sum is evaluated using Fast Fourier Transforms on a nested hierarchy of blocks. Block-P<sup>3</sup>M gives optimal accuracy, and is applicable from statics to high frequency electromagnetics.

## I. INTRODUCTION

This article describes the Block-P<sup>3</sup>M generalisation of the P<sup>3</sup>M algorithm. The work was motivated by the needs of large Method of Moments (MoM) calculations, but is also applicable to low frequency electromagnetic calculations that use integral formulations. Block-P<sup>3</sup>M has advantages over multi-level fast multipole (MLFMA) and multigrid algorithms [3, 4, 5].

Block-P<sup>3</sup>M evaluates the convolution sum

$$\phi_i = \sum_{j=1}^N G(\mathbf{x}_i; \mathbf{x}_j) \sigma_j \quad \forall i \in [1, N] \quad (1)$$

in  $\mathcal{O}(N \log N)$  operations using  $\mathcal{O}(N)$  storage. The distribution of sources  $\{\sigma_j\}$  may be on surfaces or throughout volumes and the displacement-invariant Green's function  $G(\mathbf{x}_i; \mathbf{x}_j)$  may take any functional form, for example, the Helmholtz kernel (for scattering) or the Laplace kernel (for potential fields).

Block-P<sup>3</sup>M is used in solving action-at-a-distance integral formulations of the equations of electromagnetism. Fields described by the differential equation

$$L\phi = \sigma \quad (2)$$

where  $L$  is the differential operator,  $\phi$  is the field and  $\sigma$  is the source are equivalently given by the integral equation

$$\phi(\mathbf{x}) = \int G(\mathbf{x}; \mathbf{x}') \sigma(\mathbf{x}') d\mathbf{x}' \quad (3)$$

where  $\mathbf{x}$  is the field point,  $\mathbf{x}'$  is the source point and the integral is over all volumes containing sources. The Green's function satisfies the differential equation

$$LG = \delta(\mathbf{x} - \mathbf{x}') \quad (4)$$

In applications where Block-P<sup>3</sup>M is applicable, the boundary conditions are such that  $G$  takes either a displacement invariant form  $G = G(\mathbf{r})$  where  $\mathbf{r} = (\mathbf{x} - \mathbf{x}')$ , or may be written as the sum of displacement invariant functions.

The simplest instance of action-at-a-distance formulation is in electrostatics, where Eq. (2) is Poisson's equation, Eq. (3) is Coulomb's Law and the Green's function becomes the Laplace kernel  $G = 1/4\pi\epsilon_0|\mathbf{r}|$ . In magnetostatics, the differential Ampere's Law can be transformed to the integral Biot-Savart Law. Maxwell's equations have action-at-a-distance equivalents of the

Liénard-Weichart equations in the time domain and the Stratton-Chu equations for time harmonic signals [6].

Historically, computational complexity has favoured the differential equations as the starting point for numerical computations. These are approximated using volume-filling finite difference or finite element nets within a computational volume on which boundary conditions are applied. The resulting sparse matrix equations are then solved using either direct or iterative solvers. For example, given a volume distribution of  $N$  charges, rapid elliptic solvers can find potentials in  $\mathcal{O}(N \log N)$  operations [2, Chap. 6], whereas a direct Coulomb sum takes  $\mathcal{O}(N^2)$ . Fast solvers such as P<sup>3</sup>M [1, 2] can reduce the complexity of the Coulomb sum to better than  $\mathcal{O}(N \log N)$ , making the action-at-a-distance approach faster than differential equation approaches for large problems. The Block-P<sup>3</sup>M method applied to the Helmholtz kernels that arise in MoM again gives complexity gains over volume mesh methods. Fast solvers plus the absence of body conforming volume meshes makes integral equation methods the preferred choice for complex geometries, especially if there are moving conductors.

## A. Background

All fast solver methods for evaluating the so-called 'N-body' problem (Eq. (1)) rely on

1. decomposing the Green's function into the sum of short-range and long-range parts;
2. explicitly evaluating the convolution sums for the short-range parts;
3. approximately evaluating the long-range parts on meshes.

P<sup>3</sup>M variants use continuous finite element or spline approximations of the Green's function, whereas the fast multipole methods use discontinuous point expansions about cell centres.

The P<sup>3</sup>M method [1, 2] uses additive splitting of the Green's function into a short-range part and a part that can be accurately approximated on a mesh. It works for any displacement-invariant kernel, provided that the mesh spacing is smaller than the scale-length of the long-range Green's function. The Laplace Green's function is scale-free, and this has been successfully exploited by Couchman in developing the Adaptive-P<sup>3</sup>M algorithm [7, 8]. He showed that extra speedup is obtained by recursively applying P<sup>3</sup>M over reducing scale lengths in astrophysical problems where there are large density contrasts. The scale-free nature of the Laplace kernel is also key to multigrid methods (c.f. Section III.).

Rokhlin [9] introduced the FMM for astrophysical  $N$ -body calculations. The FMM relies on analytic expansions that allow the long-range Green's function to be extrapolated from lattice points using spherical wave expansions [10]. Recursive application of the extrapolation on increasing scale-lengths led to the tree-based MLFMA [4, 5]. Recently, Dehnen has shown that replacing the spherical-wave expansion by a Cartesian Taylor expansion offers further speedup gains [11].

The FMM was extended to electromagnetic scattering [12, 13] by replacing the spherical-wave expansion by a plane-wave expansion. Fast multipole methods have been the subject of intensive research over the past decade. The multi-level fast multiple algorithm further reduces the complexity of the FMM.

Another fast method for scattering is the Adaptive Integral Method (AIM) [14, 15]. This uses a multipole expansion on a lattice, and like P<sup>3</sup>M uses Fast Fourier Transforms (FFTs) to evaluate the long-range part of the interaction. The operations count scaling for AIM (and for single-mesh P<sup>3</sup>M) degrades to  $\mathcal{O}(N^{3/2} \log N)$  for surface distributions of sources.

## B. P<sup>3</sup>M and MoM

In MoM, the Helmholtz kernel  $G \propto \exp(-ik_0|\mathbf{x}|)/|\mathbf{x}|$  has scale length  $\lambda = 2\pi/k_0$  and this limits to low frequency the use of fast methods that rely on the Green's function being scale-free.

The quadrature points for the integrals of MoM play the same role as 'particles' in particle applications of P<sup>3</sup>M. For example, the scattered field in the electric field integral equation (EFIE)  $\mathbf{E}^s = -\nabla\phi - i\omega\mathbf{A}$  is evaluated using the integral equations for  $\phi$  and  $\mathbf{A}$ . Currents are approximated by finite elements on the elements and integrals are evaluated using quadrature, allowing the approximations for  $\phi$  and  $\mathbf{A}$  to be written in the form of Eq. (1)

$$\phi(\mathbf{x}) = \frac{1}{\epsilon} \sum_j G(\mathbf{x}; \mathbf{x}_j) Q_j; \quad \mathbf{A}(\mathbf{x}) = \mu \sum_j G(\mathbf{x}; \mathbf{x}_j) \mathbf{I}_j \quad (5)$$

where  $Q_j$  and  $\mathbf{I}_j$  are the charge and current associated with quadrature point  $j$ .

The building blocks for a fast P<sup>3</sup>M-MoM solver are MoM integration routines to evaluate sources and fields at quadrature points, a Krylov solver [16] and a Block-P<sup>3</sup>M solver. These allow fast computation of the voltage excitation vector  $V$  and the product  $ZI$  in the iterative solution for the current vector  $I = Z^{-1}V$ .

Block-P<sup>3</sup>M for MoM uses a Helmholtz kernel with the singularity removed, and is used in the following steps for the fast computation of the  $ZI$  product:

- map MoM currents to the quadrature points;
- use Block-P<sup>3</sup>M to find fields (c.f. Eq. (1));
- add in singular integral corrections;
- map fields at quadrature points to the MoM basis functions;
- sum field contributions to obtain the  $ZI$  product vector.

## II. SINGLE MESH P<sup>3</sup>M

The original P<sup>3</sup>M method used a mesh Poisson solver with short-range corrections [17]. Eastwood [1, 2] generalised and optimised it by introducing the FFT solver, Green's function formulation, spline function assignment/interpolation and least-squares optimisation. P<sup>3</sup>M using point collocation is also known as the PME (Particle-Mesh Ewald) method [18].

The Particle-Particle/Particle-Mesh (P<sup>3</sup>M) method splits the potential and Green's function into a short-ranged (P-P) part and a smoothly varying (P-M) part

$$\phi_i = \phi_i^{sr} + \phi_i^{mesh}; \quad G = G^{sr} + G^{mesh} \quad (6)$$

$\phi_i^{sr}$  is evaluated using direct pairwise sums over the set  $S_i$  of near neighbours of  $i$

$$\phi_i^{sr} = \sum_{j \in S_i} G^{sr}(\mathbf{x}_i - \mathbf{x}_j) \sigma_j; \quad i \in [1, N] \quad (7)$$

If the spectral content of  $G^{mesh}$  is negligible beyond wavenumber  $k_c$ , then by the sampling theorem [19] it can be accurately represented by a sample of nodal values at an interval  $\Delta \leq \pi/k_c$ . This allows  $\phi^{mesh}$  to be accurately computed using the FFT on a mesh.  $G^{mesh}$  is approximated at general points using finite elements:

$$G^{mesh}(\mathbf{x}_i - \mathbf{x}_j) \simeq W_p(\mathbf{x}_i) G^{fe}(\mathbf{x}_p - \mathbf{x}_{p'}) W_{p'}(\mathbf{x}_j) \quad (8)$$

Nodal amplitudes  $G^{fe}$  are precomputed using a Galerkin (or point collocation) method and their FFTs  $G^k$  are stored. The use of the FFT on a regular lattice favours the use of spline functions for  $W_p$  rather than the more common finite element choices of Lagrange or Hermite polynomials. The P-M calculation reduces to

$$\text{assign:} \quad \sigma_{p'} = \sum_j W_{p'}(\mathbf{x}_j) \sigma_j \quad (9)$$

$$\text{solve:} \quad \phi_p = \sum_{p'=0}^{P-1} G_{p-p'}^{fe} \sigma_{p'} \quad (10)$$

$$\text{interpolate:} \quad \phi_i^{mesh} = \sum_p W_p(\mathbf{x}_i) \phi_p \quad (11)$$

where  $p$  and  $p'$  label mesh nodes and  $i$  and  $j$  are quadrature points in the MoM model. Figure 1 illustrates the assign and interpolate steps for quadratic splines in 2D. In 3D, the currents from the quadrature points are assigned to the nearest 27 mesh points, and fields are interpolated from the nearest 27 mesh points. Cubic splines use 64 points, and order  $p$  splines use the nearest  $(p+1)^3$  points. The correction for the smoothing effect of the spline functions is incorporated in the FFT solve.

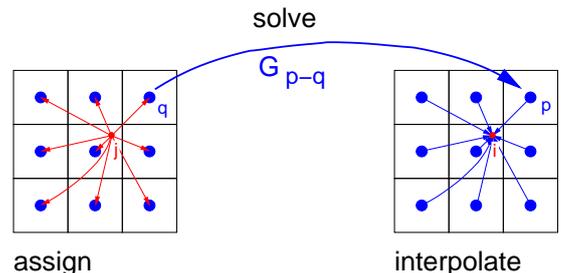


Figure 1: The assign-solve-interpolate steps of the mesh calculation.

The mesh convolution sum (Eq. (10)) is evaluated using FFTs; the FFT of source values  $\sigma_p$  gives harmonics  $\sigma^k$ . Multiplication by Green's function harmonics  $G^k$  gives potential harmonics  $\phi^k = G^k \sigma^k$ , and the inverse FFT converts these to values  $\phi_p$ . If periodic boundary conditions are used, Eq. (10) can be transformed using a length  $P$  FFT and short range contributions from periodic images need adding to the fields. If isolated boundary conditions are required, then zero padding sources and FFT lengths greater than  $2P - 1$  [20] are needed to diagonalise Eq. (10). The field  $\phi_i$  is given by combining the contributions of the P-P and P-M calculations.

P<sup>3</sup>M in two and three dimensions is given by interpreting  $\mathbf{x}$  and  $p$  as 2D or 3D vectors and indices, and  $W_p$  as products of 1D polynomial functions [2]. For spatially uniform distributions of sources, the operations count scales as  $\mathcal{O}(N\sqrt{\log N})$ , but this increases to  $\mathcal{O}(N^{3/2} \log N)$  for surface distributions in a 3D volume. The Block-P<sup>3</sup>M algorithm described below recovers the more favourable scaling for non-uniform and surface distributions that are encountered in MoM and other integral equation applications.

P<sup>3</sup>M avoids low-frequency breakdown by using interpolation for approximating the Green's function, and this should also lead to better accuracy at a given order. Figure 2 gives a low-order 1D illustration. The finite element approximation with  $\mathcal{O}(p)$  spline functions  $W_p$  gives continuity of value and derivatives to  $\mathcal{O}(p-1)$ , whereas FMM uses independent power expansions about each cell centre. In Figure 2 the exact  $\cos(r)/r$  Green's function is approximated using 2 degrees of freedom in each 0.1 interval using linear splines (as in P<sup>3</sup>M) or point expansions about cell centres (as in FMM). The accuracy of the latter degrades rapidly with decreasing  $r$ .

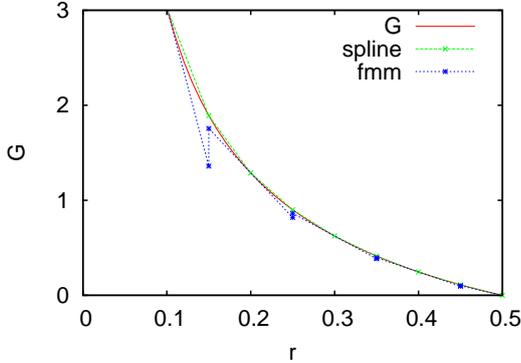


Figure 2: A 1D illustration of the FMM and P<sup>3</sup>M representations of the Green's function. See text for details.

### III. RECURSIVE INTERPOLATION

If the scale length of the Green's function is much larger than the size of the mesh, then the assign and interpolate steps can be repeated to reduce the size of the solve step mesh.

Recursive interpolation repeats the decomposition of the convolution as described above in Eqs. (6)-(11). Eq. (10) has the same form as the original problem, Eq. (1), but now for sources on the mesh. It can be decomposed in the same way as the original problem but on a coarser mesh. For the Laplace kernel, this can be repeated to the size of the system, since the problem is self-similar, but for the Helmholtz kernel, the decompositions are limited to cases where the mesh size is less than  $\lambda$ . The next section describes how to overcome this limitation.

The recursion can be terminated at any level, leaving the solve step at the coarsest mesh level to be completed as for the single-mesh P<sup>3</sup>M algorithm.

### IV. BLOCK-P<sup>3</sup>M

Recursive interpolation decomposes the solve step by repeating the short-range/long-range split of the Green's function and the assign/solve steps on increasingly coarse spatial meshes. A key feature of Block-P<sup>3</sup>M is to decompose the solve step by repeatedly grouping the mesh into blocks but now applying the FFT solve only to the non-empty subset of blocks.

Block decomposition splits 1D into 2D convolution sums (i.e., 3D to 6D in  $\mathbb{R}^3$ ) and introduces inter-block propagators  $H$ . Given Eq. (10), where  $P = QR$ , we can set  $p = q + Qr$  and rewrite it as a 2D convolution

$$\phi_{q,r} = \sum_{r'=0}^{R-1} \sum_{q'=0}^{Q-1} G_{q-q',r-r'} \sigma_{q',r'} \quad (12)$$

The convolution over index  $q$  is not periodic. However, zero padding [20]  $\sigma$  with  $(1+w)Q$  zeroes (to give  $\rho$ ) and setting  $H_{q,r} = F_q G_{q,r}$  gives a periodic convolution of period  $(2+w)Q$  that yields  $M = (2+w)Q$  values of  $\psi_{q,r}$  where  $\phi_{q,r} \equiv \psi_{q,r}$  for  $q \in [0, Q-1]$ .

$F_q$  is a filtering function which is 1 for  $|q| < Q$  and goes to zero over the interval  $Q \leq |q| < (1+w)Q$ . If the padding width  $w$  is zero, then  $F$  is a rectangular window of width  $2Q$ , but there are advantages in taking  $w \approx 1$  and using smooth transitions such as cosine bells to reduce spectral infilling.

The periodic convolution (period  $M = (2+w)Q$ ) and its Fourier transform over index  $q$  are

$$\psi_{q,r} = \sum_{r'=0}^{R-1} \sum_{q'=0}^{M-1} H_{q-q',r-r'} \rho_{q',r'} \quad (13)$$

$$\psi_r^k = \sum_{r'=0}^{R-1} H_{r-r'}^k \rho_{r'}^k \quad (14)$$

where  $k$  is the harmonic index corresponding to  $q$ , and  $q$  indices are modulo  $M$ . Any block  $r$  where  $\rho_{q,r} = 0$  for all  $q$  is ignored.

Inspection of Eq. (14) shows that the original problem (Eq. (10)) has been recovered, but now for different harmonics  $k$  on the large blocks with indices  $r$ . These harmonics are split into three groups, depending on the range of  $H_r^k$ . If  $|H_r^k|$  are below the  $k$ -space truncation amplitude they are ignored, if  $H_r^k$  are short ranged (i.e., non-zero only for small  $|x|$  so that direct evaluation is faster than FFT evaluation), Eq. (14) is evaluated by direct summation, otherwise a further level of decomposition is used, setting  $R = ST$  and repeating the above sequence.

At the final level, all remaining harmonics are short ranged. The computation of the potential is completed by repeatedly computing inverse FFTs (IFFTs) at each level and combining the results with the short-range contributions until the first level is reached.

#### A. One-level algorithm

The P-P part of the one-level Block-P<sup>3</sup>M algorithm is a direct application of Eqs. (13)-(14). Figure 3 summarises the algorithm. If  $R = 1$ , it reduces to single-mesh P<sup>3</sup>M described above in Sec. II, but in general covers the situation with  $R > 1$  illustrated in Figure 4.

In Figure 4, the surface distribution of sources is represented by the oval curve. The computational region containing the curve is divided into coarse blocks  $r$  and these in turn are subdivided into fine blocks  $q$ . Data are stored only for non-empty (shaded) blocks. Starting with source  $\sigma_j$ , the direct P-P contribution to  $\phi_i$  is given by Eq. (7), and mesh sources  $\sigma_{qr}$  are found by assignment to the fine blocks in each coarse block. Each non-empty coarse block is zero-padded and transformed to give  $\rho_r^k$ . Direct evaluation of Eq. (14) followed by inverse Fourier transformation gives  $\phi_{qr}$ , and this is interpolated to test point positions to give  $\phi_i^{mesh}$ . If Eq. (14) is evaluated by padding and FFT, then it would become a two-level scheme.

A naïve estimate of the operation count for summing Eq. (14) is  $\mathcal{O}(N_r^2 N_k)$ , where  $N_r$  is the number of non-empty coarse blocks and  $N_k$  is the number of  $k$  values. However, the nature of the inter-block propagators reduces this to close to  $\mathcal{O}(N_r)$ .

Figure 5(a) shows an example of contours of constant real part of  $H_{qr}$  that arise from the filtered section of the Helmholtz Green's function; the larger box is the range over which it is non-zero, and the smaller is the FFT periodic box. Figure 5(b)

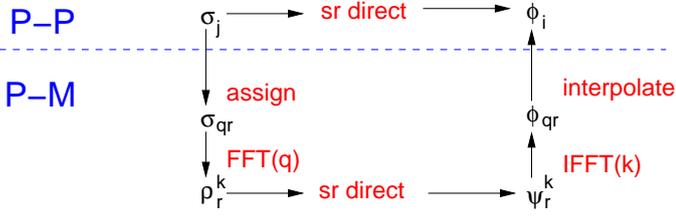


Figure 3: The one-level Block-P<sup>3</sup>M algorithm.

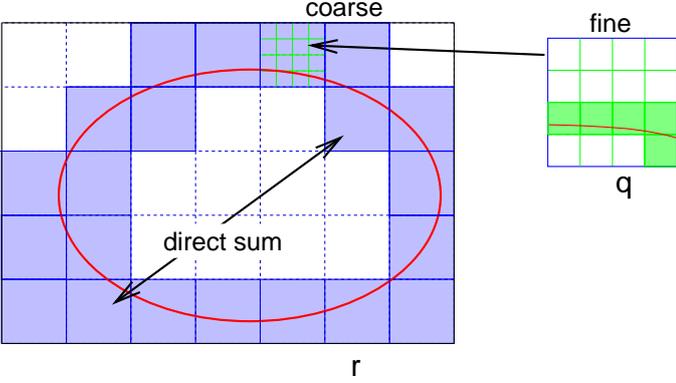
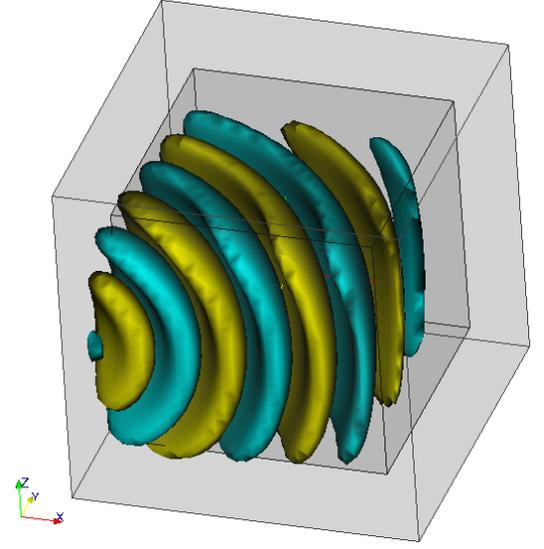
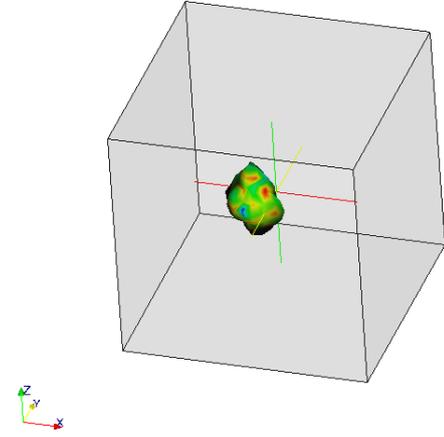


Figure 4: One-level Block-P<sup>3</sup>M blocks.



(a)  $\Re(H_{pq})$  contours



(b)  $|H_r^k|$  contour

Figure 5: (a) the real part of the level 1 propagator for  $r = (2, 2, 0)$  and (b) the 1% peak amplitude contour of the transform of the propagator for coarse blocks of side one wavelength.

shows the 1% of peak level contour of amplitude of the transformed propagator  $H_r^k$ , shaded by  $\Re(H_q^k)$ . The spectral content is strongly peaked about  $|k| = k_0$  along the vector direction between the source and test blocks, where  $k_0$  is the signal wavenumber. For larger  $r$ ,  $H_{qr}$  becomes closer to a plane wave pulse, and so its transform becomes even more localised. For  $|k|$  away from  $k_0$ ,  $H_r^k$  are only non-negligible for near-neighbour blocks. These properties substantially reduce the operations count for evaluating the sum Eq. (14). If blocks are larger in terms of wavelengths, the peaking of the propagator transform is even more pronounced.

The benefit of using a filter width  $w > 0$  is that spectral infilling from the sharp cutoff of  $H$  is avoided, and so more inter-block propagator harmonics can be discarded. For example, using a cosine bell truncation with  $w = 1$  for blocks of  $32^3$  cells with cell width of  $\lambda/8$  and a threshold of  $10^{-4}$  peak amplitude for discarding harmonics retains less than 3% of harmonics for blocks beyond second neighbour and gives an rms error  $\approx 10^{-5}$ ; a threshold of  $10^{-2}$  retains  $< 0.5\%$  and gives an error  $\approx 10^{-3}$ .

## B. Multi-level algorithm

In multilevel Block-P<sup>3</sup>M the assign and interpolate steps may be single level (c.f. Sec. II) or multilevel (c.f. Sec. III). The solve step goes through  $L$  levels of blocks computing harmonics from the smallest blocks at level 1 to the largest at level  $L$ , and then from largest to smallest computing field values. The one-level case described in the previous section generalises to:

For level 1 to  $L$ :

1. For all harmonics at the current level where  $G$  is short ranged, directly evaluate the sum, Eq. (10);
2. Terminate if all harmonics have been processed, otherwise, group sub-blocks into blocks of side  $Q$ , yielding the convolution sum, Eq. (12);
3. Zero-pad and FFT the finer mesh index  $q$  for non-empty  $r$  blocks giving for each harmonic  $k$  Eq. (14);

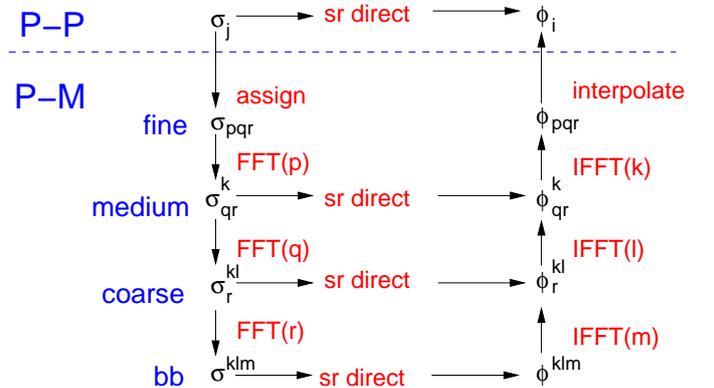


Figure 6: Flow diagram for three-level Block-P<sup>3</sup>M.

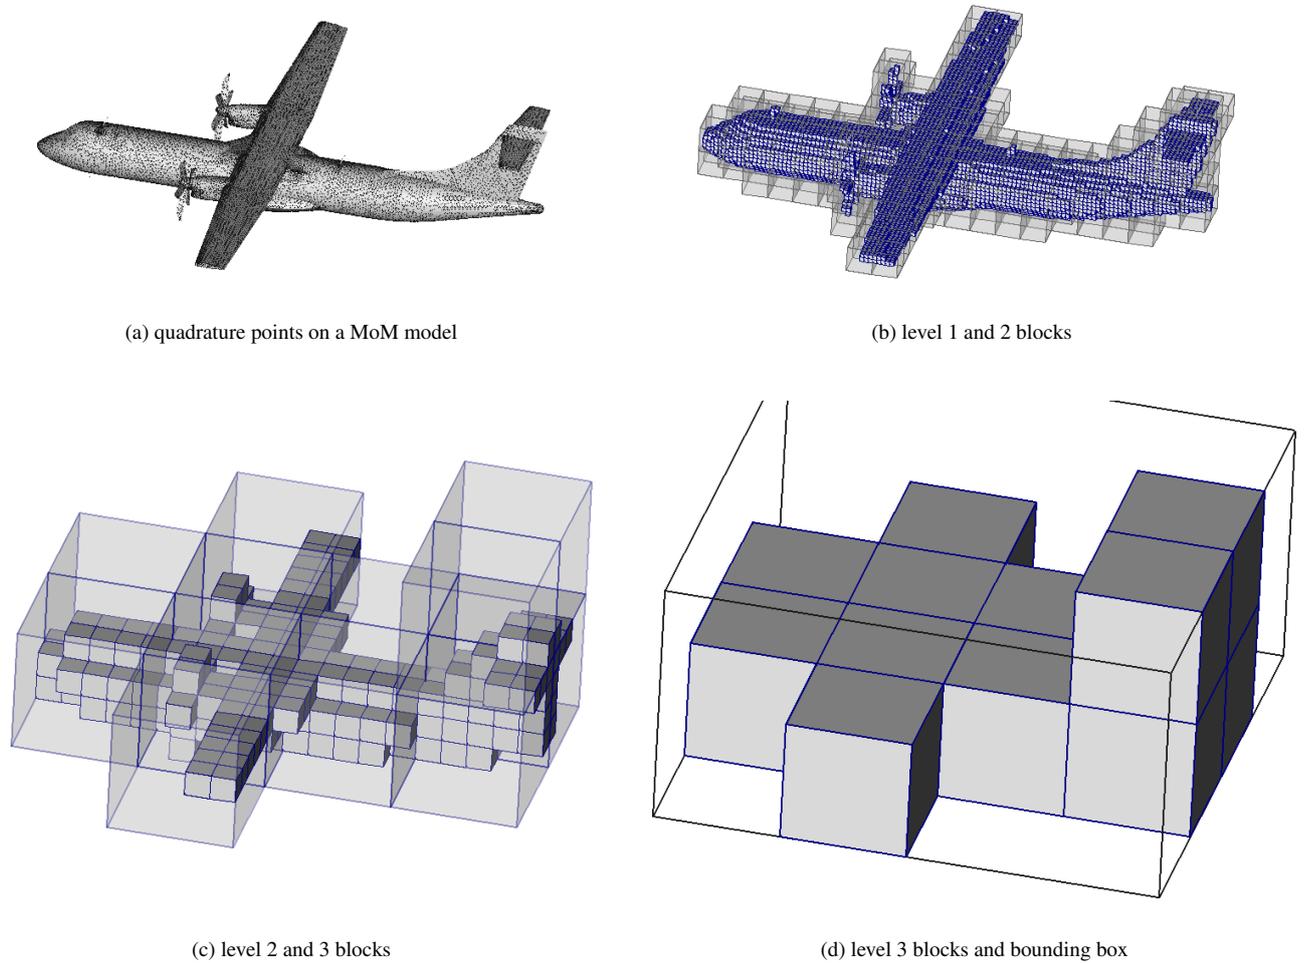


Figure 7: A MoM model and nested blocks for three-level Block-P<sup>3</sup>M.

4. Set  $r \rightarrow p$ ,  $R \rightarrow P$ ,  $\psi \rightarrow \phi$ ,  $H \rightarrow G$ ,  $\rho \rightarrow \sigma$  in Eq. (14) to recover Eq. (10) and go to the next level.

For level  $L$  to 1:

1. IFFT harmonics to get field values at the lower level;
2. discard field values for padding index values;
3. combine with directly-summed terms at this level.

Figures 6 and 7 illustrate a three-level scheme for the mesh field calculation with single assign/interpolate levels. Figure 6 gives the flow diagram and Figure 7 illustrates the steps for a real MoM model [21]. Sources at quadrature points (Figure 7(a)) on the MoM model are assigned to the fine blocks (Figure 7(b)) to give  $\sigma_{pqr}$ . These are grouped in the medium-level blocks (larger blocks in Figure 7(b)), and the data are Fourier transformed to give  $\sigma_{qr}^k$  on the medium blocks. Zero-padding is used to avoid aliasing. The result is a mesh of values for each retained  $k$  on the finer mesh in Figure 7(c).

For wavenumbers  $k$  where  $H_{qr}^k$  is short-ranged, the convolution is summed directly, and where it is long-ranged, the data are Fourier transformed to give  $\sigma_r^{kl}$  on the coarse mesh. The result is a mesh of values for each retained  $kl$  on the mesh in Figure 7(d). The same procedure as is used for lower levels is followed to get  $\sigma^{klm}$  on the bounding box.

The direct sum on the bounding box gives  $\phi_r^{klm}$ . This is inverse transformed and combined with the directly calculated  $\phi_r^{kl}$ .

The process is repeated to get  $\phi_{qr}^k$  and then  $\phi_{pqr}$  and finally the result is interpolated to the test point positions.

Cubic blocks are used to allow storage saving for the precomputed filtered Green's functions by exploiting the 48-fold symmetry:

$$H(\mathbf{i}) \equiv H(\text{Perm}\{\pm i_x, \pm i_y, \pm i_z\}) \quad (15)$$

where indices  $\mathbf{i} = (i_x, i_y, i_z)$  are triplets of  $pqr$ ,  $kqr$ ,  $klr$  or  $klm$  depending on the level.

## V. FINAL REMARKS

The Block-P<sup>3</sup>M algorithm overcomes the problem of the reduced performance of the original P<sup>3</sup>M algorithm for non-uniform source distributions that are typically found in surface integral calculations. It gives  $\mathcal{O}(N \log N)$  complexity and  $\mathcal{O}(N)$  storage for non-uniform distributions of sources and for any displacement-invariant Green's function.

The complexity for surface distributions of sources and test points is achieved through avoiding empty blocks at all levels and through discarding harmonics where the amplitudes of the propagator harmonics are below the chosen threshold level. Direct summation of convolutions is used where the range of the propagator is small, and Fourier transforms are used where the FFT is computationally faster. Storage reduction is achieved by storing only blocks that are non-empty. When blocks do contain non-empty sub-blocks, then sparse storage methods are used to store only those sub-blocks. Computations and storage needed

for the precomputed inter-block propagators are reduced by exploiting 48-fold symmetry on cubic lattices and by storing only those harmonics of the propagators with amplitudes above the cutoff threshold.

The algorithm is built on FFTs on uniform lattices and works equally well for Laplace and Helmholtz kernels. Unlike FMM, it does not rely on truncated divergent series, and so does not suffer breakdown when point separations become small compared to the wavelength.

If a Galerkin approximation is used in Block-P<sup>3</sup>M then it gives optimal least-squares accuracy. The trade-off between computational cost and accuracy remains optimal even when  $k$ -space truncation is used. How to make the best choice for combining the recursive interpolation, recursive block decomposition and  $k$ -space truncation is currently being investigated. However, Block-P<sup>3</sup>M is expected to provide a competitively fast method that exploits optimised FFT routines and is amenable to parallelisation.

Block-P<sup>3</sup>M is also of value for volume-filling source distributions. The block decomposition offers a method for breaking a global mesh in large computations on distributed-memory parallel computers into a set of spatially-localised blocks on which the FFT can be used without communication between processors. The lower levels of the calculation are embarrassingly parallel, and higher levels involve few harmonics so there is a much reduced interprocessor communication when compared to applying the FFT to the global mesh.

The Block-P<sup>3</sup>M algorithm is particularly attractive for electromagnetic scattering calculations using the method of moments. It has the same complexity and storage scaling as the MLFMA, but superior properties: MLFMA has a point-matched discontinuous approximation to  $G$  that is recursively extrapolated using a mixture of rectangular and polar meshes and suffers from low-frequency breakdown. In contrast, Block-P<sup>3</sup>M has a least-squares fit, continuous (to the order of the splines used) approximation to  $G$  that is recursively interpolated on rectangular meshes, and does not suffer low-frequency breakdown.

## VI. REFERENCES

- [1] J.W. Eastwood. Optimal P<sup>3</sup>M algorithms for molecular dynamics simulations. In M.B. Hooper, editor, *Computational methods in classical and quantum physics*, pages 206–228. Advance Publications Ltd, 1976.
- [2] R.W. Hockney and J.W. Eastwood. *Computer simulation using particles*. McGraw-Hill, 1981. (Student Ed., Adam-Hilger 1988).
- [3] A. Brandt. Multi-level adaptive solutions to boundary-value problems. *Math. Comp.*, 31(138):333–390, April 1977.
- [4] A. Brandt. Multilevel computations of integral transforms and particle interactions with oscillatory kernels. *Comput. Phys. Commun.*, 65(1–3):24–38, April 1991.
- [5] W.C. Chew, J-M. Jin, E. Michielssen, and J. Song, editors. *Fast and efficient algorithms in computational electromagnetics*. Artech House, 2001.
- [6] J.A. Stratton. *Electromagnetic Theory*. McGraw-Hill, 1941.
- [7] H.M.P. Couchman. Mesh-refined P<sup>3</sup>M: A fast adaptive  $N$ -body algorithm. *Astrophys. J.*, 368:L23–L26, February 1991.
- [8] R.J. Thacker and H.M.P. Couchman. A parallel adaptive P<sup>3</sup>M code with hierarchical particle reordering. *Comput. Phys. Commun.*, 174(7):540–554, April 2006.
- [9] V. Rokhlin. Rapid solution of integral equations of classical potential theory. *J. Comput. Phys.*, 60(2):187–207, September 1985.
- [10] M.A. Epton and B. Dembart. Multipole translation theory for the three-dimensional Laplace and Helmholtz equations. *SIAM J. Sci. Comput.*, 16(4):865–897, July 1995.
- [11] W. Dehnen. A hierarchical  $\mathcal{O}(N)$  force calculation algorithm. *J. Comput. Phys.*, 179(1):27–42, June 2002.
- [12] N. Engheta, W.D. Murphy, V. Rokhlin, and M.S. Vassiliou. The fast multipole method (FMM) for electromagnetic scattering problems. *IEEE Trans. Antennas Propag.*, 40(6):634–641, June 1992.
- [13] R. Coifman, V. Rokhlin, and S. Wandzura. The fast multipole method for the wave equation: a pedestrian description. *IEEE Antennas Propag. Mag.*, 35(3):7–12, June 1993.
- [14] E. Bleszynski, M. Bleszynski, and T. Jaroszewicz. A fast integral-equation solver for electromagnetic scattering problems. *IEEE AP-S Int. Symp. Dig.*, 1:416–419, June 1994.
- [15] E. Bleszynski, M. Bleszynski, and T. Jaroszewicz. AIM: Adaptive integral method for solving large-scale electromagnetic scattering and radiation problems. *Radio Sci.*, 31(5):1225–1251, September-October 1996.
- [16] H.A. van der Vorst. Krylov subspace iteration. *Computing in Sci. and Engng*, 2(1):32–37, January-February 2000.
- [17] R.W. Hockney, S.P. Goel, and J.W. Eastwood. A 10000 particle molecular dynamics model with long range forces. *Chem. Phys. Lett.*, 21(3):589–591, September 1973.
- [18] T. Darden, D. York, and L. Pedersen. Particle mesh Ewald: An  $N\log(N)$  method for Ewald sums in large systems. *J. Chem. Phys.*, 98(12):10089–10092, June 1993.
- [19] R. Bracewell. *The Fourier transform and its applications*. McGraw-Hill, 1965.
- [20] J.W. Eastwood and D.R.K. Brownrigg. Remarks on the solution of Poisson’s equation for isolated systems. *J. Comput. Phys.*, 32:24–38, July 1979.
- [21] D. Leclerc. ATR, Toulouse, private communication, 2008.

## AUTHORS NAME AND AFFILIATION

James Eastwood,  
Culham Electromagnetics Ltd,  
Culham Science Centre,  
Abingdon, Oxfordshire,  
OX14 3DB, UK

Telephone: +44 1865 408326  
email: James.Eastwood@CulhamEM.co.uk