Boundary Element Methods 15 (BEM15)
Worcester, MA, 8/93

Multipole-Accelerated Preconditioned Iterative Methods For Three-Dimensional Potential Problems

F. T. Korsmeyer and D. K. P. Yue
Department of Ocean Engineering
K. Nabors and J. White
Department of Electrical Engineering and Computer Science
Massachusetts Institute of Technology, Cambridge, MA 02139 U.S.A.

Abstract

In this paper a multipole-accelerated iterative algorithm for solving the matrices generated by panel or boundary-element method discretizations of three-dimensional integral equations is described. The method reduces the memory and computational cost of solving these equations from N^2 to nearly order N, where N is the number of degrees of freedom in the discretization. Experimental results from a problem in potential flow is presented to demonstrate that the method has sufficiently low overhead that discretizations with more than four thousand unknowns can be solved in minutes on a scientific workstation.

1 Introduction

For a wide variety of problems in computer-aided engineering analysis, the most computationally expensive task is solving Laplace's equation for the potential on a two-dimensional surface embedded in a three-dimensional problem domain. An integral equation for the surface potential can be derived using Green's theorem, from which it follows that for each point x on a piecewise smooth surface S, the potential, $\psi(x)$, must satisfy

$$2\pi\psi(x) + \int_S \psi(x')G_n(x,x')da' - \int_S \psi_n(x')G(x,x')da' = 0. \qquad (1)$$

where $G(x,x') = \frac{1}{\|x-x'\|}$. Given (1), the surface potential can be determined uniquely if for each point x on S, the potential, $\psi(x)$, and its normal derivative,

$\psi_n(x) \equiv \vec{\nabla}\psi(x) \cdot \hat{n}$, are constrained to satisfy

$$\beta(x)\psi(x) + \gamma(x)\psi_n(x) = f(x). \tag{2}$$

The simplest, and still commonly used, approach to numerically solving (1) and (2) is to divide the surface into N triangular or quadralateral panels, and then assume that ψ and ψ_n are constant over each panel. Insisting that this piecewise constant appoximation satisfies (1) and (2) at a collection of N collocation points, denoted $\{x_i\}$, leads to a system of equations of the form

$$2\pi + Pp - Dp_n = 0 \tag{3}$$

and

$$\Upsilon p + \Gamma p_n = f \tag{4}$$

where $p, p_n \in \Re^N$ are the vectors of coefficients of the piecewise constant approximations to ψ and ψ_n respectively, $\Upsilon, \Gamma \in \Re^{N \times N}$ are diagonal matrices whose diagonal elements are given by $\Upsilon_{jj} = \beta(x_j)$ and $\Gamma_{jj} = \gamma(x_j)$. The entries in $P, D \in \Re^{N \times N}$ are given by

$$P_{i,j} = \frac{1}{a_j} \int_{S_j} \frac{1}{\|x' - x_i\|} da', \tag{5}$$

and

$$D_{i,j} = \frac{1}{a_j} \int_{S_j} \vec{\nabla} \frac{1}{\|x' - x_i\|} \cdot \hat{n} \; da', \tag{6}$$

where S_j is the surface of the j^{th} panel, and a_j is the j^{th} panel's surface area.

The system composed of (3) and (4) can be solved by direct factorization, or with an iterative technique like a Krylov-subspace method [11]. As P and D in (3) are dense matrices, the cost of directly factoring (3) and (4) grows like N^3, and even a rapidly converging iterative method will still require computation time and memory which grows like N^2.

In this paper, we will describe an accelerated approach, based on using a fast-multipole algorithm, which reduces the time and memory required to solve (3) and (4) to nearly order N. In the next section, we will briefly describe the fast-multipole algorithm, and then in Section 3 we will give explicit formulas for determining the multipole expansions of piecewise constant panels. In Section 4, we will give experimental results demonstrating the accuracy and efficiency of our multipole-accelerated algorithm applied to a problem in potential flow. Finally, in Section 5 we give conclusions and acknowledgements.

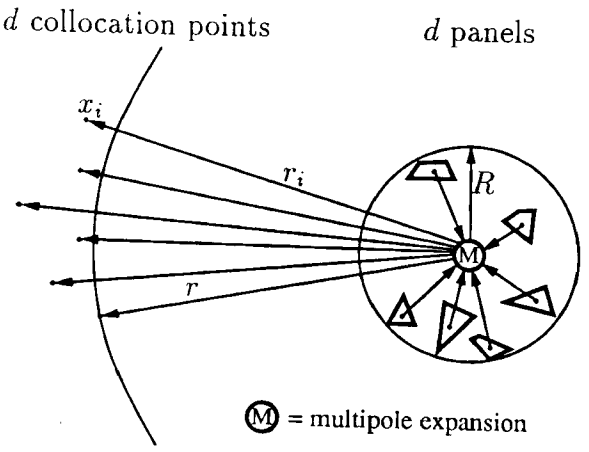

Figure 1: Evaluation of d distant panel potentials at d collocation points with the multipole expansion.

2 Multipole Acceleration

Krylov-Subspace based iterative methods for solving the N equation system $Ax = b$ select solutions from the space of vectors $\{r, Ar, A^2r, ...A^kr\}$, where k is the iteration index. If A is dense, then the cost of computing a $(k+1)^{st}$ iteration for most Krylov-Subspace methods is dominated by the N^2 cost of forming the dense matrix-vector product $A^{k+1}r = A(Ar^k)$. In the case of equation (3), computing the dense matrix-vector product is equivalent to evaluating potentials at N points, $\{x_1, ..., x_N\}$ (the panel centroids), due to monopole and dipole distributions on the N panels. Therefore, it is possible to reduce the cost of a Krylov subspace method for solving (3) and (4) by accelerating the potential calculation.

There are a wide variety of $NlogN$ algorithms for accelerating the evaluation of N potentials due to N distributions, mostly based on some form of hierarchical panel clustering [2, 3]. The basic idea is depicted in Figure 1. Here, the potential due to a cluster of panels is evaluated at some distant point x_j by first accumulating the panel influences into a multipole expansion, and then evaluating the single expansion. As we will describe briefly below, the fast-multipole algorithm [6, 9] reduces the cost yet futher, to order N, by efficiently distributing the accumulated multipole expansions through the use of local expansions. More specifically, the use of multipole and local expansions are orchestrated by a tree-structured hierarchy of panel clusters; multipole expansions for clusters of panels are accumulated from the *leaves* of the tree to the *root*, and local expansions are distributed from the root to the leaves for evaluation at collocation points. This is accomplished in order N operations while maintaining a uniform precision.

The implementation of a multipole expansion is shown in Figure 1. Here, the potential due to a cluster of panels is represented by a truncated multipole

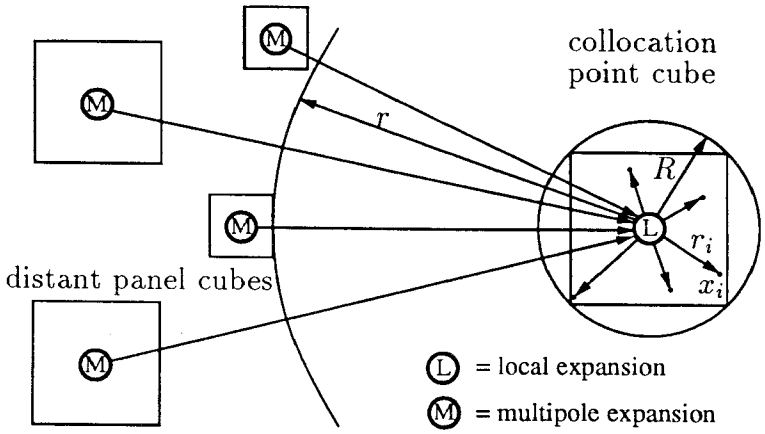

Figure 2: The conversion of several distant multipole expansions into one local expansion for evaluation at the local collocation points.

expansion:

$$\psi(r_i, \theta_i, \phi_i) \approx \sum_{n=0}^{l} \sum_{m=-n}^{n} \frac{M_n^m}{r_i^{n+1}} Y_n^m(\theta_i, \phi_i), \tag{7}$$

where l is the expansion order and r_i, θ_i, and ϕ_i are the spherical coordinates of the i-th evaluation point relative to the center of the multipole expansion. The $Y_n^m(\theta_i, \phi_i)$ are the surface spherical harmonics [7, 8] and the M_n^m are the complex multipole coefficients given by

$$M_n^m = \sum_{j=1}^{d} \frac{q_j}{a_j} \int_{S_j} (\rho')^n Y_n^{-m}(\alpha', \beta') da_c', \tag{8}$$

where there are d panels in the cluster with constant density strengths q_j/a_j. The variables of integration are the spherical coordinates (ρ', α', β') of the differential panel surface area da_c' relative to the center of the multipole expansion.

The error in this approximation is [5],

$$\left| \psi(r_i, \theta_i, \phi_i) - \sum_{n=0}^{l} \sum_{m=-n}^{n} \frac{M_n^m}{r_i^{n+1}} Y_n^m(\theta_i, \phi_i) \right| \leq K_1 \left(\frac{R}{r_i}\right)^{l+1} \leq K_1 \left(\frac{R}{r}\right)^{l+1}. \tag{9}$$

The quantities r and R are as in Figure 1 and K_1 is a constant independent of the multipole expansion order, l.

In the fast multipole algorithm, most of the multipole expansions are not evaluated at collocation points. Instead, multipole expansions are either combined to form multipole expansions which represent distributions in a greater portion of the domain, or multipole expansions are transformed into local expansions.

The implementation of a local expansion is shown in Figure 2. Here, a number of multipole expansions are transformed into a single local expansion. Local

expansions have the form:

$$\psi(r_i, \theta_i, \phi_i) \approx \sum_{n=0}^{l} \sum_{m=-n}^{n} L_n^m Y_n^m(\theta_i, \phi_i) r_i^n, \tag{10}$$

where l is the order of the expansion, r_i, θ_i and ϕ_i are the spherical coordinates of the i-th collocation point with respect to the center of the local expansion, and the L_n^m are the complex local expansion coefficients.

The error in this approximation is:

$$\left| \psi(r_i, \theta_i, \phi_i) - \sum_{n=0}^{l} \sum_{m=-n}^{n} L_n^m Y_n^m(\theta_i, \phi_i) r_i^n \right| \leq K_2 \left(\frac{r_i}{r}\right)^{l+1} \leq K_2 \left(\frac{R}{r}\right)^{l+1} \tag{11}$$

where K_2 is independent of l, and r and R are as in Figure 2.

Local and multipole expansions must be carefully applied to insure that the potential is accurately approximated everywhere in the problem domain. The structure which makes this possible is the hierarchical partitioning of the domain. Consider the smallest cube which contains the entire domain, that is, all of the panels. We refer to this cube as the level 0 or root cube. This *parent* cube is subdivided into eight *child* cubes, and the panels are divided among these level 1 child cubes. This process is repeated down to some *finest* level (the leaves), designated level L. The number of levels, L, is usually selected so that no finest level cube contains more than some fixed small number of panels. After setting up this hierarchical spatial decomposition, the fast multipole algorithm begins with the finest level, where each panel distribution is represented by a multipole expansion. These expansions are then shifted to the centers of the finest-level cubes and combined, so that a single expansion represents all of the panels in the cube. During an *upward pass* through the tree to the root, each child cube's multipole expansions is shifted to the child cube's parent's center, to generate a single expansion which represents all of the panels in the parent cube. In an *interaction phase* , at each level a local expansion is created for each cube by accumulating multipole expansions representing distant cubes at that level. In a *downward pass,* the local expansions in the parent cubes are shifted to the centers of their children. Finally, in an *evaluation phase,* the local expansions and direct contributions from nearby panels are evaluated at the points at which the potential is required: the collocation points.

It is important to note that all expansion shifting and transforming can be represented as translation matrices whose elements depend only on the geometry of the surface and the cube hierarchy. If the fast-multipole algorithm is used to evaluate matrix-vector products as part of an iterative procedure, efficiency can be improved by forming the translation matrices once and then reusing them for each matrix-vector product. Based on this observation, the fast multipole algorithm can be summarized as in Figure 3, where $Q2M$ denotes the matrix which maps panel distribution coefficients to multipole expansion coefficients, $M2M$ denotes

the matrix which shifts a multipole expansion's center, $M2L$ denotes the matrix which transforms multipole expansion coefficients to local expansion coefficients, $L2L$ denotes the matrix which shifts a local expansion's center, and $L2P$ denotes the matrix which evaluates the local expansion at the collocation points.

Finally, note that as panels are generated by discretizing a surface in the three-dimensional problem domain, for practical examples most of the cubes at level 2 and finer in the spatial hierarchy will be empty. For this reason, adaptive versions of the multipole algorithm are commonly used [4]. An adaptive strategy particularly well suited to panel or boundary element methods is detailed in [9].

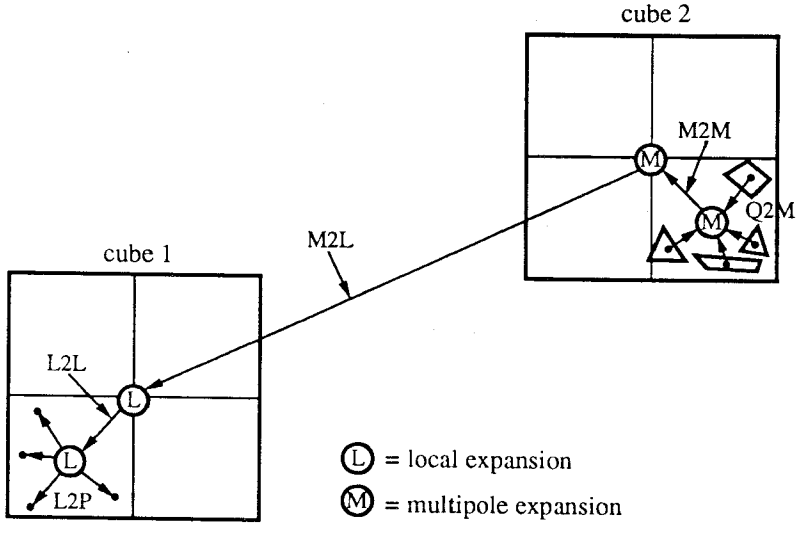

Figure 3: The basic operations of the fast-multipole algorithm.

3 Expansions for Singularity Distributions

Unlike N particle problems, panel or boundary element methods require multipole expansions for the panel *distributions* . In principle, these distributions may be any order of approximation of the unknown, but higher-order panels are more complicated to represent, and since the overall algorithm is nearly order N, the use of many simple panels has little penalty. We do not in fact evaluate the finest level multipole coefficients by (8) directly but rather compute the multipole coefficients for the singularity distributions on each panel in local panel coordinates, and then shift these coefficients with an $M2M$ operation to the centers of the finest-level cubes. A local cartesian coordinate system, with the origin coincident with the panel centroid and the x_1, x_2 plane coplanar with the panel, is chosen for convenience. This necessitates a rotation of the expansion as well as a translation when panel expansions are accumulated at the centers of the finest-level cubes. Formulas for rotating spherical harmonics are already well known, particularly in the context of quantum angular momentum (for example, see [12]).

We require the coefficients of a multipole expansion such that:

$$\sum_{n=0}^{\infty}\sum_{m=-n}^{n}\frac{M_n^m}{r^{n+1}}Y_n^m(\theta,\phi)=\iint_{S_j}\frac{1}{r'}dx_1'dx_2' \tag{12}$$

The definition of these multipole coefficients for the constant source distribution on a planar panel is derived in [9]:

$$M_n^m=K_n^m\sum_{j=0}^{\frac{n-|m|}{2}}\binom{\frac{n-|m|}{2}}{\frac{n-|m|}{2}-j}\sum_{k=0}^{|m|}(-\mathrm{sgn}(m)i)^{|m|-k}\binom{|m|}{|m|-k}\mathcal{I}_{2j-k,n-(2j+k)}, \tag{13}$$

where

$$K_n^m=\sqrt{\frac{(n-|m|)!}{(n+|m|)!}}P_n^{|m|}(0), \tag{14}$$

in which the $P_n^{|m|}$ are the Legendre functions, and the terms

$$\mathcal{I}_{j,k}\equiv\int_Q x^j y^k da \tag{15}$$

are the moments of the panel, recursion relations for which can be found in [10].

In the local panel coordinate system described above, we require for the dipole distribution:

$$\sum_{n=0}^{\infty}\sum_{m=-n}^{n}\frac{N_n^m}{r^{n+1}}Y_n^m(\theta,\phi)=\iint_{S_j}\frac{\partial}{\partial x_3'}\frac{1}{r'}dx_1'dx_2'. \tag{16}$$

Taking this derivative of (13) is straightforward because of the choice of coordinate system and simply requires the recurrence relations for the derivatives of the Legendre functions [1]. The result is:

$$N_n^m=\sqrt{n^2+m^2}M_n^m \tag{17}$$

4 An Example Problem

To demonstrate the efficiency of the adaptive, multipole-accelerated, preconditioned iterative algorithm for solving the boundary integral equations with a combination of Neumann and Dirichlet boundary conditions, a model problem is selected to which the solution is known in closed form. The problem is the representation of a sphere translating in an ideal, infinite fluid.

For the unit sphere translating at unit velocity in an ideal fluid, in this case parallel to the x_3 axis and in the direction of its positive sense, the potential is known to be:

$$\psi(x)=-\frac{1}{2}\frac{x_3}{||x||^3}. \tag{18}$$

To test the algorithm for the direct formulation with both types of boundary conditions, a Dirichlet boundary condition is imposed on the leading hemisphere,

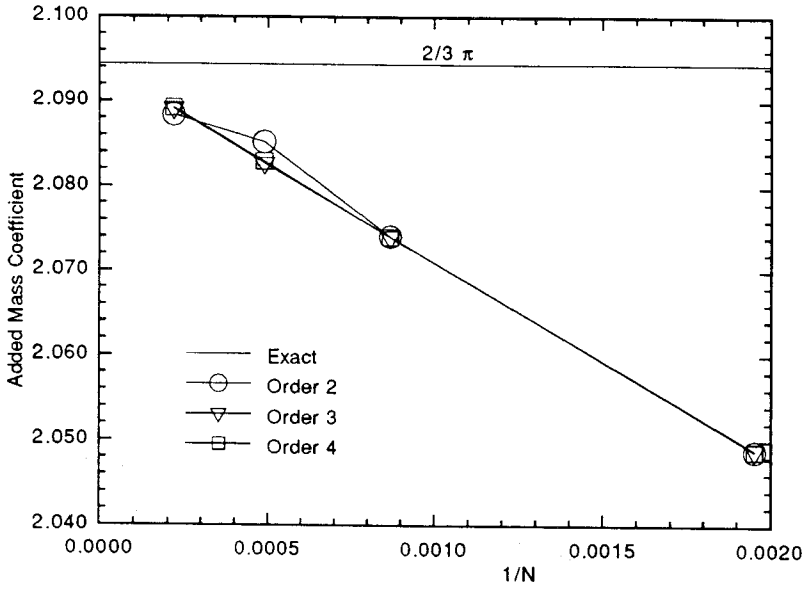

Figure 4: The convergence of the added-mass coefficient for a translating sphere as the discretization is refined. The added-mass of a sphere in an infinite fluid is known in closed form to be equal to $\frac{2}{3}\pi$.

and a Neumann boundary condition is imposed on the trailing hemisphere. The solution, then, is the complimentary function on each hemisphere.

The convergence of the computed solution to the exact solution is demonstrated in Figure 4. This figure plots the computed added-mass coefficient as a function of the number of panels used to discretize the sphere. The added-mass coefficient is defined as:

$$a_{33} = \int\int_{S_D \cup S_N} \psi(x') n_3 da', \tag{19}$$

where n_3 is the component of the unit surface normal in the direction of sphere translation. From equation (18), it is easy to see that (19) may be written as

$$a_{33} = \int\int_{S_D \cup S_N} \psi(x') \psi_n(x') da'. \tag{20}$$

Therefore, from (20), the added-mass may be computed as a sum over all the panels of the product of the boundary condition, the solution, and the panel area. In Figure 4 the convergence of the added-mass is linear in $\frac{1}{N}$ if the multipole expansions are carried out to a sufficiently high order, which for the coarser two discretzations of 512 and 1152 panels is order 2, and for the finer two discretizations of 2048 and 4608 panels is order 3. It should be noted that the increased accuracy of the order 2 result for 4608 panels is fortuitous.

Number of Spheres	Number of Panels	Depth	Number of Iterations	CPU Time Seconds	Memory Mb
1	1152	4	6	58	16.8
2	2304	6	13	105	31.2
4	4608	6	14	228	65.9
1	1152	0	6	104	11(est.)

Table 1: Computation requirements on a DECStation 5000/240 for a fictitious Neumann-Dirichlet problem of multiple spheres. The spheres each have 1152 panels. The order of the multipole expansions is 3, and the convergence tolerance on the GMRES algorithm is 0.001. The *Depth* column gives the value of L, the number of levels in the spatial hierarchy. For reference to an order N^2 solution, in the last line of the table, where the depth is zero, the solution has no multipole acceleration. The storage typically required by an order N^2 approach is estimated simply as $N \times N \times 8$ bytes, while the storage reported for the solutions by the multipole accelerated algorithm are the actual total memory allocations.

Table 4 shows the linear growth in the computational effort required for a fictitious problem of multiple spheres. The problems are posed by setting the Neumann and Dirichlet conditions for the single translating sphere on hemispheres of one, two, and four spheres, which are arrayed with one half radius separation in the cases of the multiple spheres. This demonstrates the increase in computational effort as problems increase in complexity, rather than simply increase in the number of unknowns on a fixed geometry.

5 Conclusions and Acknowledgements

In this paper, we demonstrated that multipole-accelerated iterative algorithms are extremely effective for solving three-dimensional surface integral equations. The results show that the computation time of the method grows linearly with problem size, and that the fast-multipole algorithm has an advantage even on problems with as few as 500 panels, and is nearly an order of magnitude faster on typical 4000 panel problems. Currently we are using the algorithm on a workstation to investigate the physics of free-surface waves. The simulation of the evolution of gravity waves on a free-surface can require hundreds of solutions of large Laplace problems with a mix of Neumann and Dirichlet boundary conditions. Hitherto this work has only been done on supercomputers. Future work is on incorporating higher-order panels, as well as trying to exploit the multipole algorithm's structure to efficiently adapt the surface discretization.

This work is supported by the Defense Advanced Research Projects Agency (N00014-91-J-1698), the Office of Naval Research (N00014-90-J-1085), the Na-

tional Science Foundation (MIP-8858764 A02), the F.B.I. (J-FBI-88-067,) and grants from Digital Equipment Corporation and I.B.M.

References

[1] M. Abramowitz and I. A. Stegun. *Handbook of Mathematical Functions.* Dover Press, New York, 1964.

[2] R. W. Appel. An efficient program for many-body simulation. *SIAM J. Sci. Stat. Comp.*, 6:85–103, 1985.

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

[4] J. Carrier, L. Greengard, and V. Rokhlin. A fast adaptive multipole algorithm for particle simulations. *SIAM Journal on Scientific and Statistical Computing*, 9(4):669–686, July 1988.

[5] L. Greengard. *The Rapid Evaluation of Potential Fields in Particle Systems.* M.I.T. Press, Cambridge, Massachusetts, 1988.

[6] L. Greengard and V. Rokhlin. A fast algorithm for particle simulations. *Journal of Computational Physics*, 73(2):325–348, December 1987.

[7] E. W. Hobson. *The Theory of Spherical and Ellipsoidal Harmonics.* Chelsea, New York, 1955.

[8] O. D. Kellogg. *Foundations of Potential Theory.* Dover Publications, New York, 1953. Reissue of the 1929 edition.

[9] K. Nabors, F. T. Korsmeyer, F. T. Leighton, and J. White. Preconditioned, adaptive, multipole-accelerated iterative methods for three-dimensional potential integral equations of the first-kind. To appear: *SIAM J. Sci. Comp.*

[10] J. N. Newman. Distributions of sources and normal dipoles over a quadralateral panel. *Journal of Engineering Mathematics*, 20:113–126, 1986.

[11] Y. Saad and M. H. Schultz. GMRES: A generalized minimal residual algorithm for solving nonsymmetric linear systems. *SIAM J. Sci. Stat. Comp.*, 7:856–869, 1986.

[12] J. J. Sakurai. *Modern Quantum Mechanics.* Addison-Wesley, Reading, 1985.