1996

An efficient fast-multipole algorithm based on an expansion in the solid harmonics

H.Y. Wang, *Los Alamos National Laboratory*
Richard Alan Lesar, *Los Alamos National Laboratory*

An efficient fastmultipole algorithm based on an expansion in the solid harmonics

H. Y. Wang and R. LeSar

Citation: The Journal of Chemical Physics **104**, 4173 (1996); doi: 10.1063/1.471667
View online: http://dx.doi.org/10.1063/1.471667
View Table of Contents: http://scitation.aip.org/content/aip/journal/jcp/104/11?ver=pdfcov
Published by the AIP Publishing

An efficient fast-multipole algorithm based on an expansion in the solid harmonics

H. Y. Wang and R. LeSar
Center for Materials Science, Los Alamos National Laboratory, Los Alamos, New Mexico 87545

(Received 12 September 1995; accepted 14 November 1995)

We present an efficient variant of the fast-multipole method for calculating long-range interactions in three-dimensional Coulombic systems. Using a multipole expansion based on the solid harmonics instead of the more common spherical harmonics leads to a greater increase in computational efficiency than a recently-reported fast-Fourier transform method, with none of the overhead associated with that approach. © *1996 American Institute of Physics.* [S0021-9606(96)50408-7]

I. INTRODUCTION

The calculation of long-range interactions arises in a number of fields, from simulations of galactic motion[1–3] to ionic materials[4,5] to dislocation microstructures.[6] The simplest approach in the case of two-body interactions is to simply sum all interaction terms. That approach is limited to small system sizes, as the computational burden is $O(N^2)$, where N is the number of particles. Simply cutting off the interactions at some finite range leads to spurious results.[5] Thus, a great deal of work has been done to develop other, lest computationally costly but still accurate, approaches. The goal in this paper is not to review those developments, but rather to discuss an efficient way to implement one of them, the $O(N)$ fast-multipole method of Greengard and Rokhlin.[7]

The fast-multipole method is one of a class of algorithms called hierarchical-tree methods, all of which share two important characteristics; they all utilize a hierarchical-tree data structure and they all compute the force on an individual particle from nearby particles directly and that from remote particles by an approximate, truncated, multipole expansion. The fast-multipole method (FMM) by Greengard and Rokhlin[7] adopts an oct-tree structure, without any restriction on the number of particles in the lowest cells. The major difference between the FMM and other tree methods[1,2,8] is in the way far-field interactions are handled. With a careful management of the interactions, the FMM method makes the computational efficiency essentially $O(N)$. In addition to its computational efficiency, the error in the FMM can be reduced to close to machine accuracy by keeping enough terms in the multipole expansions.

Much development has occurred in the fast-multipole method since the first paper by Greengard and Rokhlin.[7] The method has been extended to three dimensions,[9,10] implemented on parallel computers,[11–13] and implemented for molecular dynamics simulations.[14] Recently, fast-Fourier transform methods have been introduced to increase the computational speed of calculations in three dimensions.[15,16] While increasing computational speed by up to a factor of about 1.6 per energy calculation on serial computers, the use of fast-Fourier transforms introduces a considerable amount of computational and programming complexity.

Here we present an efficient algorithm for performing the FMM in three dimensions that is an alternative to the use of fast-Fourier transforms. By using a multipole expansion based on the solid harmonics rather than the customary spherical harmonics, we develop a much more compact and computationally less costly algorithm with increases in speeds greater than those reported with the fast-Fourier transform approach (up to a factor of 2.4). In Sec. II we outline the FMM with the solid harmonics, in Sec. III we show how the obvious symmetry of the solid harmonics can be used to increase computational efficiency, and in Sec. IV we summarize the results.

II. FAST-MULTIPOLE METHOD

Here we discuss the basic approach taken in the fast-multipole method (FMM). For more details, as well as mathematical justification and error analysis, see Greengard and Rokhlin.[7]

When two sets of particles are ''well-separated,'' as shown in Fig. 1, it is possible to determine the number of terms needed in a multipole expansion of the field of one set of particles to calculate the forces on the particles the other within a given accuracy.[7] This is not to say that being ''well-separated'' is needed for the convergence of a multipole expansion. Convergence only requires that the sets be separated. Note that a slightly different definition is used for ''well-separated'' cells in the multipole methods described next.

The multipole methods work by creating a hierarchy of cells; for clarity we show the two-dimensional case. A typical three-level two-dimensional hierarchy is shown in Fig. 2. The highest level, 0, is the simulation cell itself. That cell is then divided into 4ths to create level-1 cells. The level-1 cells are divided again into 4ths to create level-2, and that in turn is divided again to create level-3. We say that a level-$(n+1)$ cell derived by dividing a level-n cell into 4ths is the child of the level-n cell, which in turn is the parent of the four level-$(n+1)$ cells. For a M-level system, there are a total of $\Sigma_{n=0}^{M} 4^n$ cells. In three dimensions, 8 children are created from each parent and there are a total of $\Sigma_{n=0}^{M} 8^n$ cells for an M-level system. In this type of hierarchy, the next-nearest neighbor cells are the first that are considered ''well-separated;'' this holds at all levels.

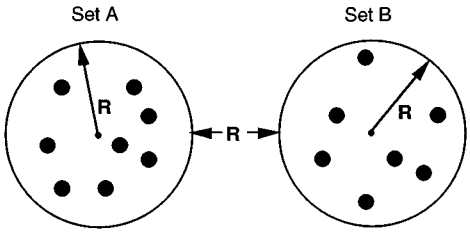

FIG. 1. Well-separated sets of dislocations. In this definition, two regions are well-separated if dislocations in each region are closer to each other than to dislocations in the other region.

The FMM achieves its O(N) computational speed by an efficient calculation of the multipole moments and the interactions. As a concrete example, consider a system of charged particles in a level-3 system (Fig. 2). The interaction of particles in cell a with other particles in a and with those in the nearest-neighbor level-3 cells (marked n) are evaluated as a direct summation of the electrostatic interactions. The next-nearest-neighbor level-3 cells (marked nn) are well-separated from cell a and thus their interactions with particles in cell a are approximated with a multipole expansion. The level-2 cells marked B are well separated from the level-2 parent of cell a, thus the particles in cell a interact with the multipoles of the level-2 B cells. The multipole moments of the various cells are calculated as shown in Fig. 3. The moments of the level-3 cells are computed directly, as discussed in detail below. The moments of the level-2 cells (B in Fig. 3) are computed by *shifting* the multipole expansion from the center of the level-3 cells to that of their parent level-2 cells, yielding simple expressions for the level-2 multipoles that do not require summing over the particles. The multipole moments of the level-1 cells (i.e., parents of the level-2 cells) would be calculated in the same way. The interactions of the level-3 particles (i.e., in a) with the multipoles of a level-2 cell (B) are calculated by shifting the center of the multipole expansion of the potential of the level-2 cell to the well-separated parent (level-2) of the cell in which the particle sits, creating what is usually referred to as a *local* expansion centered at the parent cell (A). This

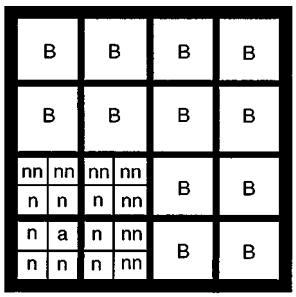

FIG. 2. Schematic diagram for a interactions in a 3-level hierarchy. The dislocations in cell a interact with those in cell a and in all nearest-neighbor cells (n) by direct interactions and with the multipoles of those in the next-nearest-neighbor shells (nn) and with the multipoles of those in the next-nearest-neighbors (B) of the parent of a.

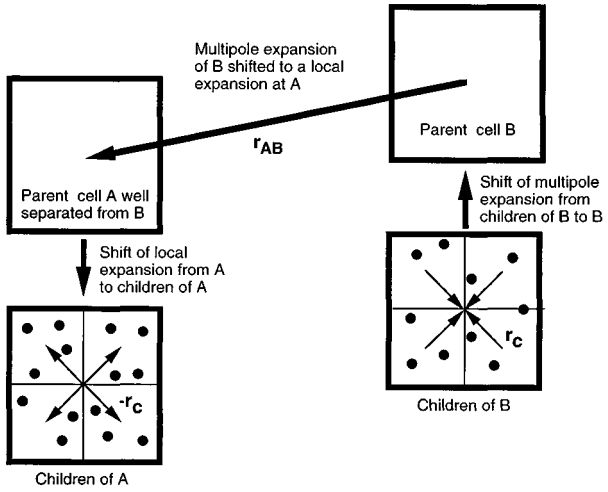

FIG. 3. Example of the steps needed in calculating the interactions of particles in cell a in Fig. 2 with the multipoles of the next-nearest neighbors (B) of the parent of a. The multipoles of the children of B are calculated directly. The multipoles of B are then calculated by shifting up the multipole expansions to be centered on B. The multipole expansions are then shifted to the parent of a and then shifted down to be centered on a.

local expansion is then shifted to be centered at the level-3 cell (a) and the net interaction determined by adding the terms for each particle.

The total electrostatic potential due to a set of point charges is

$$\Phi(\mathbf{r}) = \sum_i \frac{q_i}{|\mathbf{r}-\mathbf{r}_i|}, \tag{1}$$

where q_i is the charge of the ith particle and $\mathbf{r}_i$ its position. In previous implementations of the fast multipole method in three dimensions,[10] this potential was rewritten in a multipole expansion based on the spherical harmonics (defined in the Appendix),

$$\Phi(\mathbf{r}) = 4\pi \sum_{l=0}^{\infty} \sum_{m=-l}^{l} \frac{M_{lm} Y_{lm}(\theta,\varphi)}{(2l+1) r^{l+1}}, \tag{2}$$

where the multipole moments M_{lm} are given by

$$M_{lm} = \sum_i q_i r_i^l Y_{lm}^*(\theta_i, \varphi_i). \tag{3}$$

To evaluate the potential from a well-separated cell (e.g., B) in a cell at the same level (A), we rewrite the multipole expansion from B as an expansion in variables *local* to A, which takes the form

$$\Phi(\mathbf{r}) = 4\pi \sum_{l=0}^{\infty} \sum_{m=-l}^{l} L_{lm} r^l Y_{lm}(\theta,\varphi), \tag{4}$$

where $\mathbf{r}$ is a vector relative to the center of A and the L_{lm} are the *local* moments, which can be written in terms of the multipole moments of B and geometrical factors. This transformation from a multipole series to a local expansion is referred to as a multipole-to-local transformation.

In the notation of Fig. 3, the three transformations for the multipole moments and local moments are[10]

(1) multipole-to-multipole (children of $B \rightarrow B$),

$$M'_{\lambda\mu} = \sum_{l=0}^{\infty} \sum_{m=-l}^{l} \left\{ 4\pi \frac{(-r_c)^{\lambda-1} Y^*_{\lambda-l,\mu-m}(\theta_c,\varphi_c) a_{\lambda-l,\mu-m} a_{lm} (2\lambda-1)}{(2l+1)[2(\lambda-l)+1] a_{\lambda\mu}} \right\} M_{lm}; \tag{5}$$

(2) multipole-to-local ($B \rightarrow A$),

$$L_{\lambda\mu} = \sum_{l=0}^{\infty} \sum_{m=-l}^{l} \left\{ 4\pi \frac{(-1)^{l+m} Y^*_{\lambda+l,\mu-m}(\theta_{AB},\varphi_{AB}) a_{\lambda\mu} a_{lm}}{r_{AB}^{\lambda+l+1} (2l+1)[2(\lambda+l)+1] a_{\lambda+l,\mu-m}} \right\} M_{lm}; \tag{6}$$

(3) local-to-local ($A \rightarrow$children of A),

$$L'_{\lambda\mu} = \sum_{l=0}^{\infty} \sum_{m=-l}^{l} \left\{ 4\pi \frac{r_c^{l-\lambda} Y_{l-\lambda,m-\mu}(\theta_c,\varphi_c) a_{\lambda\mu} a_{l-\lambda,m-\mu}}{(2\lambda+1)[2(l-\lambda)+1] a_{lm}} \right\} L_{lm}; \tag{7}$$

where

$$a_{lm} = (-1)^{l+m} \frac{(2l+1)^{1/2}}{[4\pi(l+m)!(l-m)!]^{1/2}}. \tag{8}$$

An alternative approach to the use of the spherical harmonics is offered by the use of the solid harmonics, as defined in the Appendix. These functions have the benefit of being functions of the vectors $\mathbf{r}$. By taking advantage of addition formulas given in the Appendix, the final expressions are far more compact and require fewer computations. Moreover, using the obvious symmetry of these functions leads to a considerable reduction in the number of calculations needed.

With the definition of $\Re_{lm}(\mathbf{r})$ and $\Im_{lm}(\mathbf{r})$ in Eqs. (A5) and (A6), the multipole expansion in Eq. (2) can be rewritten as

$$\Phi(\mathbf{r}) = \sum_{l=0}^{\infty} \sum_{m=-l}^{l} M_{lm} \Im_{lm}(\mathbf{r}), \tag{9}$$

where the multipole moments are

$$M_{lm} = \sum_i q_i \Re^*_{lm}(\mathbf{r}_i). \tag{10}$$

The local expansion of Eq. (4) is then given by, with $\mathbf{r}$ always taken as the vector relative to the center of the cell in which the expansion is to be evaluated,

$$\Phi(\mathbf{r}) = \sum_{l=0}^{\infty} \sum_{m=-l}^{l} L_{lm} \Re^*_{lm}(-\mathbf{r}). \tag{11}$$

Using the addition formulas in Eqs. (A7) and (A8), the three types of transformations can be easily written as

(1) multipole-to-multipole (children of $B \rightarrow B$),

$$M'_{lm} = \sum_{(\lambda,\mu)=(0,0)}^{l,m} M_{l-\lambda,m-\mu} \Re^*_{\lambda\mu}(\mathbf{r}_c); \tag{12}$$

(2) multipole-to-local ($B \rightarrow A$),

$$L_{lm} = \sum_{\lambda=0}^{\infty} \sum_{\mu=-\lambda}^{\lambda} M_{\lambda\mu} \Im_{l+\lambda,m+\mu}(\mathbf{r}_{AB}); \tag{13}$$

(3) local-to-local ($A \rightarrow$children of A),

$$L'_{lm} = \sum_{\lambda=0}^{\infty} \sum_{\mu=-\lambda}^{\lambda} L_{l+\lambda,m+\mu} \Re^*_{\lambda\mu}(\mathbf{r}_c). \tag{14}$$

The summation in Eq. (12) is defined in Eq. (A9). The form of these expansions is much more compact and somewhat less computer intensive than those based on an expansion in the spherical harmonics [Eqs. (5)–(8)]. As we shall see next, however, the real benefit of using the solid harmonics is that their symmetry can be invoked to reduce the number of computations and to significantly increase the efficiency of the computations.

Note that while we have written the upper limit of all sums over l as ∞, in practice only a finite number of terms are actually used. Specifically, in an order P expansion, all sums take the form $\Sigma_{l=0}^{\infty} \Rightarrow \Sigma_{l=0}^{P-1}$.

III. OPTIMIZATION OF THE ALGORITHM

All of the transformations in Eqs. (5)–(8) or Eqs. (12)–(14) require approximately the same amount of computation. However, while there are at most 8 transformations per cell for the ''vertical'' multipole-to-multipole and local-to-local shifts, there are considerably more for the ''horizontal'' multipole-to-local shifts ($B \rightarrow A$ in Fig. 3). From simple counting arguments, we see that for each cell there are as many as 189 next-nearest-neighbor cells (in three dimensions) at the same level to which its multipoles must be shifted. Thus, the multipole-to-local transformations dominate the computational time to perform the transformations.

Irregular solid harmonics [Eq. (A6)] with field vectors $\mathbf{r}$ and $\mathbf{r}_0$ that differ only by $\Delta\varphi$ (i.e., they lie in the same xy plane) are related by

$$\Im_{lm}(\mathbf{r}) = \Im_{lm}(\mathbf{r}_0) e^{im\Delta\varphi}. \tag{15}$$

We rewrite Eq. (13) for the multipole-to-local expansion moment as

$$L_{lm}(\mathbf{r}_0)=\sum_{\mu=-(P-1)}^{P-1} D_{lm}^{\mu}(\mathbf{r}_0), \tag{16}$$

where

$$D_{lm}^{\mu}(\mathbf{r}_0)=\sum_{\lambda=|\mu|}^{P-1} M_{\lambda\mu}\mathfrak{I}_{l+\lambda,m+\mu}(\mathbf{r}_0). \tag{17}$$

Using Eq. (15), we then have for the local moment at a cell at $\mathbf{r}$,

$$L_{lm}(\mathbf{r})=\sum_{\mu=-(P-1)}^{(P-1)} e^{i(\mu+m)\Delta\varphi} D_{lm}^{\mu}(\mathbf{r}_0). \tag{18}$$

By tabulating $D_{lm}^{\mu}(\mathbf{r}_0)$ for one cell, we can easily evaluate the moment at all other cells with $\mathbf{r}$ differing only by $\Delta\varphi$ from $\mathbf{r}_0$.

Besides the fourfold symmetry in the xy plane, the mirror symmetry can be used to further reduce the computation. In Eq. (A6), if $\mathbf{r}$ is related to $\mathbf{r}_0$ by changing θ to $\pi-\theta$, we have

$$\mathfrak{I}_{lm}(\mathbf{r})=(-1)^{l+m}\mathfrak{I}_{lm}(\mathbf{r}_0). \tag{19}$$

If we define

$$E_{lm}(1)=\sum_{\mu=-(P-1)}^{P-1}\ \sum_{\lambda=|\mu|,|\mu|+2,\ldots}^{P-1} M_{\lambda\mu}\mathfrak{I}_{l+\lambda,m+\mu}(\mathbf{r}_0) \tag{20a}$$

and

$$E_{lm}(2)=\sum_{\mu=-(P-1)}^{P-1}\ \sum_{\lambda=|\mu|+1,|\mu|+3,\ldots}^{P-1} M_{\lambda\mu}\mathfrak{I}_{l+\lambda,m+\mu}(\mathbf{r}_0), \tag{20b}$$

we have

$$L_{lm}(\mathbf{r}_0)=E_{lm}(1)+E_{lm}(2) \tag{21}$$

and

$$L_{lm}(\mathbf{r})=(-1)^{l+m}[E_{lm}(1)-E_{lm}(2)]. \tag{22}$$

In practice, we combine the $\Delta\varphi$ and θ to $\pi-\theta$ transformations to minimize computations. Utilizing these symmetry properties, the total number of independent cells is reduced from 189 to 34. In addition to the 34 full evaluations of the multipole-to-local transformation, we have 92 $\Delta\varphi$ shifts and 63 θ to $\pi-\theta$ shifts for the most general case.

There are P^2 terms in each evaluation of $L_{lm}(\mathbf{r}_0)$. For each multipole-to-local transformation there are P^2 values of $L_{lm}(\mathbf{r}_0)$ to be calculated. Thus, invoking no symmetry, the full multipole-to-local transformation would involve evaluating $189P^4$ terms for each of the $c\Sigma_{n=0}^{M}8^n$ cells. Using the symmetry of the solid harmonics reduces that burden to $34P^4$ terms. In addition, we need 92 evaluations of the $\Delta\varphi$ shift [Eq. (18)] and 63 θ to $\pi-\theta$ shifts [Eq. (22)], which require far fewer operations. For example, the 92 θ to $\pi-\theta$ shifts require no multiplications and just 1 addition, i.e., they contribute $O(P^2)$ to the total operations. We then expect that the overall decrease in the number of operations (i.e., increase in computational speed) should asymptotically approach 189/34 ($\sim$5.6) at large P. In Fig. 4 we show the calculated increase in computational speed for the multipole-to-local transformation as a function of the order of the multipole expansion. We see an increase from about 2.8 at $P=1$ to about 5.6 at $P=16$, as expected.

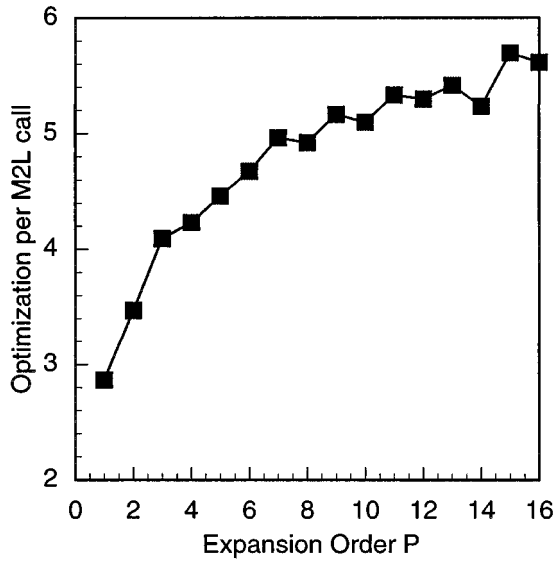

FIG. 4. Optimization (i.e., speedup over the solid-harmonic expansion with no symmetry) of the multipole-to-local transformation as a function of the expansion order P.

We can estimate how increases in computational efficiency in calculating the multipole-to-local shifts affects the computer time for a full potential calculation. Since there are 189 multipole-local shifts for each cell and only 8 multipole-to-multipole and local-to-local shifts, we ignore the latter. We then have the total time in evaluating shifts as approximately $c\Sigma_{n=0}^{M}8^n$, where c is the time for one shift. The other term that dominates the computational time is the direct calculation over the (at most) 27 nearest cells. If we assume that

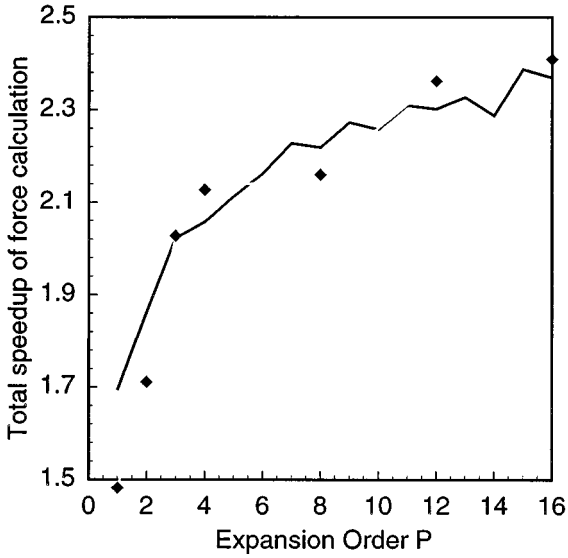

FIG. 5. Optimization (i.e., speedup over the solid-harmonic expansion with no symmetry) of a single total energy calculation as a function of the expansion order P. The diamonds are an analytical approximation to the speedup, as described in the text.

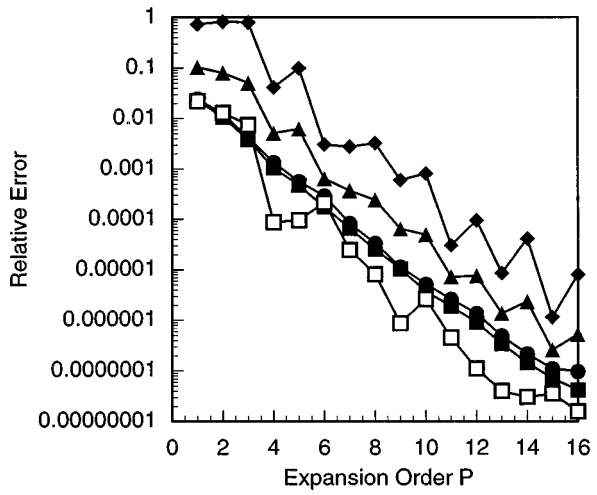

FIG. 6. Error in potential as a function of expansion order P calculated for a system of 512 charges (the direct calculation is too time consuming to use more than this). Shown are five estimates of the error, where

$$\Phi_{\text{direct}}(\mathbf{r}_i)=\sum_{j\neq i=1}^{N}\frac{q_j}{|\mathbf{r}_j-\mathbf{r}_i|}:$$

$$(\blacksquare)\quad \sum_{i=1}^{N}|\Phi(\mathbf{r}_i)-\Phi_{\text{direct}}(\mathbf{r}_i)|\ \Big/\ \sum_{i=1}^{N}|\Phi_{\text{direct}}(\mathbf{r}_i)|,$$

$$(\bullet)\quad \sqrt{\sum_{i=1}^{N}|\Phi(\mathbf{r}_i)-\Phi_{\text{direct}}(\mathbf{r}_i)|^2}\ \Big/\ \sum_{i=1}^{N}|\Phi_{\text{direct}}(\mathbf{r}_i)|^2,$$

$$(\blacktriangle)\quad \frac{1}{N}\sum_{i=1}^{N}\{|\Phi(\mathbf{r}_i)-\Phi_{\text{direct}}(\mathbf{r}_i)|/|\Phi_{\text{direct}}(\mathbf{r}_i)|\},$$

$$(\blacklozenge)\quad \sqrt{\frac{1}{N}\sum_{i=1}^{N}\{|\Phi(\mathbf{r}_i)-\Phi_{\text{direct}}(\mathbf{r}_i)|^2/|\Phi_{\text{direct}}(\mathbf{r}_i)|^2\}},$$

$$(\square)\quad \left|\left\{\sum_{i=1}^{N}[\Phi(\mathbf{r}_i)-\Phi_{\text{direct}}(\mathbf{r}_i)]\right\}\ \Big/\ \sum_{i=1}^{N}\Phi_{\text{direct}}(\mathbf{r}_i)\right|.$$

s is the time for each direct calculation, we can show that the total computational time should scale as $\sqrt{sc}$. Thus, we expect that increases in computational speed for the multipole-to-local transformation from c to c' would reduce the overall computational speed by a factor of $\sqrt{c'/c}$. This approximate relation is seen to be verified in Fig. 5, where we compare the total time per potential calculation with and without using symmetry to reduce the computations. We also show the predicted increase in speed found by using the data in Fig. 4 for c'/c.

There have been many estimates in the literature of convergence properties of the multipole expansions. In Fig. 6, we show a number of these. Plotted are the relative errors of the fast-multipole method relative to the exact, direct, summation as a function of the expansion order P. The main differences between the different measures of error are whether they are relative to the mean total potential (■,●,□) or are they sums of individual relative errors (▲,◆). We prefer the more standard definitions shown by the curves ■ and ●. From those, it is clear that only a relatively low-order multipole expansion is needed; a 0.01% error in total potential is obtained by $P=7$. Note, however, that the total number of operations in the multipole expansions increases as P^4, so a heavy penalty is paid for increased accuracy.

We verify in Fig. 7 that the algorithm is O(N) by plotting cpu time for a total potential calculation as a function of the total number of particles. A critical parameter in determining optimized computational speeds for the FMM is the number of particles in the smallest cell (P/C), which is a function of both the total number of particles and the level of the heirarchy. We show in Fig. 6 results where we have optimized P/C for each N. In Table I, we list the levels and optimal P/C as a function of the number of particles. The calculations were run on an SGI workstation with the 150 MHz RS 4400 processor. In these calculations, all particles had the same magnitude charge ($q=1$) and the sides of the simulation cell were set to $N^{1/3}$ (i.e., the density is one). All calculations were performed with an order seven ($P=7$) multipole expansion.

TABLE I. N is the number of particles, the level refers to the number of divisions made in the heirarchy, P/C is the average number of particles in the smallest cell, and the cpu time is the time (in s) per one force calculation on an SGI workstation with the 150 MHz RS 4400 processor. All calculations were done with a seventh-order multipole expansion.

N	Level	P/C	cpu time (s)
4 000	2	40	0.12
20 480	3	40	0.90
30 720	3	60	1.47
147 456	4	36	8.02
250 000	4	61	14.40

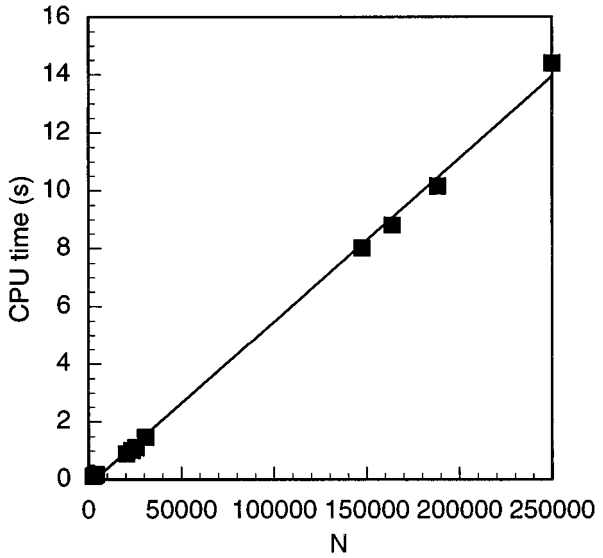

FIG. 7. Total cpu time per potential calculation as a function of the total number of particles. Shown are results where we have optimized the number of particles in the smallest cell. The calculations were run on an SGI workstation with the 150 MHz RS 4400 processor.

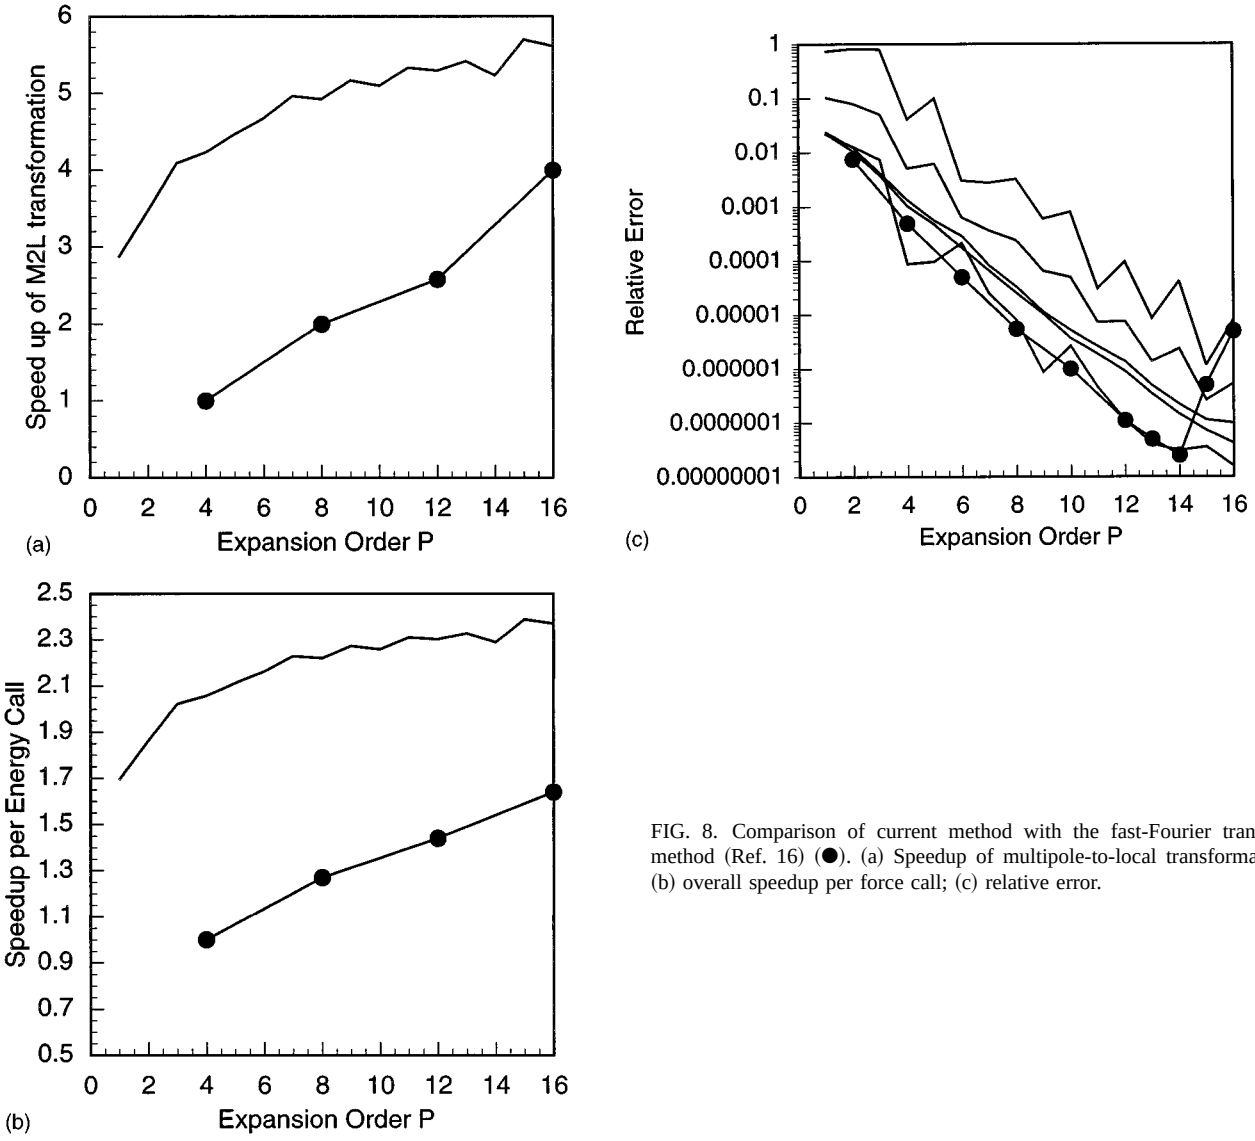

FIG. 8. Comparison of current method with the fast-Fourier transform method (Ref. 16) (●). (a) Speedup of multipole-to-local transformations; (b) overall speedup per force call; (c) relative error.

Because of the convolutionlike form of the multipole-to-local expansion, a fast Fourier transform (FFT) has been utilized to improve the speed,[16] which converts the computation from $O(P^4)$ to $O[P^2 \log_2(P)]$. In Fig. 8, we compare the FFT (Ref. 16) with the present results and see that the current approach offers a greater increase in computational speed than the FFT. For an expansion order of 4 (0.1% error), the current approach offers almost four times the speed up than the FFT for the multipole-to-local transformation. Note that they[16] compare the speedup relative to the base spherical-harmonic expansion while we present increases in computational speed relative to an expansion with the solid harmonics with no symmetry. The computational time for the base calculations (spherical harmonics and solid harmonics with no symmetry) should be similar, so the comparison of the two results is justified.

IV. SUMMARY AND CONCLUSIONS

We have presented an alternative expansion for the fast-multipole method using the solid harmonics. Not only does use of these functions lead to more compact expressions, we can take advantage of the obvious symmetry properties of these functions to greatly reduce the computational burden. Comparison with a recently-reported[16] fast Fourier transform method shows that use of the present approach offers a much greater increase in computational speed.

ACKNOWLEDGMENTS

This work was performed under the auspices of the United States Department of Energy (U.S. DOE) and was supported in part by the Division of Materials Science of the Office of Basic Energy Sciences of the U.S. DOE.

APPENDIX

The properties of the spherical harmonics $Y_{lm}(\theta,\varphi)$ are well known, though phase and normalization factors can differ between authors. Here we use the definitions

$$Y_{lm}(\theta,\varphi)=\left[\frac{2l+1}{4\pi}\frac{(l-m)!}{(l+m)!}\right]^{1/2}P_l^m(\cos\ \theta)e^{im\varphi} \tag{A1}$$

with the associated Legendre functions

$$P_l^m(x)=\frac{(-1)^m}{2^l l!}(1-x^2)^{m/2}\frac{d^{l+m}}{dx^{l+m}}(x^2-1)^l. \tag{A2}$$

The $P_l^m(x)$ satisfy the following relations:

$$P_l^{-m}(x)=(-1)^m\frac{(l-m)!}{(l+m)!}P_l^m(x) \tag{A3}$$

and

$$P_l^m(-x)=(-1)^{l+m}P_l^m(x). \tag{A4}$$

Following Steinborn and Ruedenberg,[17] we define the regular and irregular solid harmonics as follows:

$$\mathfrak{R}_{lm}(\mathbf{r})=\frac{1}{(l+m)!}r^l P_l^m(\cos\ \theta)e^{im\varphi}, \tag{A5}$$

$$\mathfrak{I}_{lm}(\mathbf{r})=\frac{(l+m)!}{r^{l+1}}P_l^m(\cos\ \theta)e^{im\varphi}. \tag{A6}$$

Note that $\mathfrak{R}_{lm}(\mathbf{r})$ corresponds to $\hat{\mathscr{Y}}_{lm}(\mathbf{r})$ and $\mathfrak{I}_{lm}(\mathbf{r})$ to $\hat{\mathscr{S}}_{lm}(\mathbf{r})$ in the notation used by Steinborn and Ruedenberg. These functions satisfy the following addition formulas:

$$\mathfrak{R}_{lm}(\mathbf{r}_1+\mathbf{r}_2)=\sum_{(\lambda,\mu)=(0,0)}^{l,m}\mathfrak{R}_{l-\lambda,m-\mu}(\mathbf{r}_1)\mathfrak{R}_{\lambda\mu}(\mathbf{r}_2), \tag{A7}$$

$$\mathfrak{I}_{lm}(\mathbf{r}_>+\mathbf{r}_<)=\sum_{\lambda=0}^{l}\sum_{\mu=-\lambda}^{\lambda}(-1)^{\lambda+\mu}\mathfrak{R}_{\lambda,-\mu}(\mathbf{r}_<)\times\mathfrak{I}_{l+\lambda,m+\mu}(\mathbf{r}_>)\quad |\mathbf{r}_<|<|\mathbf{r}_>|, \tag{A8}$$

where

$$\sum_{(\lambda,\mu)=(0,0)}^{l,m}\equiv\sum_{\lambda=0}^{l}\sum_{\mu=\max(-\lambda,\lambda-l+m)}^{\mu=\min(\lambda,-\lambda+l+m)}. \tag{A9}$$

[1] J. G. Jernigan, in *IAU Symposium 113, Dynamics of Star Cluster*, edited by J. Goodman and P. Hut (Reidel, Dordrecht, 1985), p. 275.
[2] J. Barnes and P. Hut, Nature **324**, 446 (1986); J. Barnes and P. Hut, Astrophys. J. Suppl. **70**, 389 (1989).
[3] J. K. Salmon and M. S. Warren, J. Comp. Phys. **111**, 136 (1994).
[4] M. P. Allen and D. J. Tildesley, *Computer Simulation of Liquids* (Clarendon, Oxford, 1987).
[5] J.-P. Hansen, in *Molecular Dynamics Simulations of Statistical-Mechanical Systems*, edited by G. Gicotti and W. G. Hoover (North-Holland, Amsterdam, 1986), pp. 89–128.
[6] H. Y. Wang and R. LeSar, Philos. Mag. A **71**, 149 (1995).
[7] L. Greengard and V. Rokhlin, J. Comp. Phys. **73**, 325 (1987).
[8] A. W. Appel, SIAM J. Sci. Stat. Comput. **6**, 85 (1985).
[9] J. Carrier, L. Greengard, and V. Rokhlin, SIAM J. Sci. Stat. Comput. **9**, 669 (1988).
[10] K. E. Schmidt and M. A. Lee, J. Stat. Phys. **63**, 1223 (1991).
[11] L. Greengard and W. D. Gropp, Comp. Math. Applic. **20**, 63 (1990).
[12] F. Zhao and S. Lennart Johnsson, SIAM J. Sci. Stat. Comput. **12**, 1420 (1991).
[13] J. F. Leathrum and J. A. Board, Technical Report, Duke University Department of Electrical Engineering, 1992.
[14] J. A. Board, J. W. Causey, J. F. Leathrum, A. Windemuth, and K. Schulten, Chem. Phys. Lett. **198**, 89 (1992); J. A. Board, Z. S. Hakura, W. D. Elliot, D. C. Gray, W. J. Blanke, and J. F. Leathrum, Technical Report, Duke University Department of Electrical Engineering 94-002, 1994.
[15] L. Greengard and V. Rokhlin, Yale University Research Report No. 602, 1988.
[16] W. D. Elliot and J. A. Board, Technical Report, Duke University Department of Electrical Engineering 94-001, 1994.
[17] E. O. Steinborn and K. Ruedenberg, in *Advances in Quantum Chemistry*, edited by P.-O. Löwdin (Academic, New York, 1973), Vol. 7, pp. 1–81.