

Particle-particle, particle-scaling function algorithm for electrostatic problems in free boundary conditions

Alexey Neelov, S. Alireza Ghasemi, and Stefan Goedecker

Citation: *The Journal of Chemical Physics* **127**, 024109 (2007); doi: 10.1063/1.2746328

View online: <http://dx.doi.org/10.1063/1.2746328>

View Table of Contents: <http://scitation.aip.org/content/aip/journal/jcp/127/2?ver=pdfcov>

Published by the [AIP Publishing](#)

Articles you may be interested in

[A divide-conquer-recombine algorithmic paradigm for large spatiotemporal quantum molecular dynamics simulations](#)

J. Chem. Phys. **140**, 18A529 (2014); 10.1063/1.4869342

[Extending molecular simulation time scales: Parallel in time integrations for high-level quantum chemistry and complex force representations](#)

J. Chem. Phys. **139**, 074114 (2013); 10.1063/1.4818328

[Electrostatic interactions in finite systems treated with periodic boundary conditions: Application to linear-scaling density functional theory](#)

J. Chem. Phys. **135**, 204103 (2011); 10.1063/1.3662863

[Boundary conditions in local electrostatics algorithms](#)

J. Chem. Phys. **128**, 214103 (2008); 10.1063/1.2918365

[Electrostatic free energy calculations using the generalized solvent boundary potential method](#)

J. Chem. Phys. **117**, 7381 (2002); 10.1063/1.1507108

A row of several tablets or e-readers is shown, each displaying the cover of the journal 'Computing: Science & Engineering'. The covers feature a colorful, abstract, circular pattern. The text 'computing' is in a large, bold, orange font, with 'SCIENCE & ENGINEERING' in a smaller, black font below it. At the bottom of the image, there is a dark banner with white text that reads 'AIP's JOURNAL OF COMPUTATIONAL TOOLS AND METHODS. AVAILABLE AT MOST LIBRARIES.' The 'computing' logo is also present in the bottom right corner of the banner.

Particle-particle, particle-scaling function algorithm for electrostatic problems in free boundary conditions

Alexey Neelov, S. Alireza Ghasemi, and Stefan Goedecker

Institute of Physics, University of Basel, Klingelbergstrasse 82, CH-4056 Basel, Switzerland

(Received 23 February 2007; accepted 11 May 2007; published online 12 July 2007)

An algorithm for fast calculation of the Coulombic forces and energies of point particles with free boundary conditions is proposed. Its calculation time scales as $N \log N$ for N particles. This novel method has lower crossover point with the full $O(N^2)$ direct summation than the fast multipole method. The forces obtained by our algorithm are analytical derivatives of the energy which guarantees energy conservation during a molecular dynamics simulation. Our algorithm is very simple. A version of the code parallelized with the Message Passing Interface can be downloaded under the GNU General Public License from the website of our group. © 2007 American Institute of Physics.

[DOI: 10.1063/1.2746328]

I. INTRODUCTION

The computation of the electrostatic or gravitational interaction of a large number N of point particles is a central problem in many fields of physics, such as molecular dynamics¹ and astrophysics.²

If the charge distribution is nonperiodic, one frequently employs the fast multipole method (FMM).^{3–8} Its computation time scales as $O(N)$. Its point of the crossover with full direct calculation depends on the accuracy and is estimated to be 10^3 – 10^4 in Ref. 5 or 10^4 – 10^5 in Ref. 9. The Barnes and Hut hierarchical tree method¹⁰ is advocated for the use with low accuracy (0.1%–1%) in Ref. 11.

Although the cutoff methods still attract attention^{12,13} since they scale as $O(N)$, the corresponding prefactor is immense. These methods also introduce artifacts. Citing Ref. 14 (see the references therein), “The severe artifacts caused by cutoff truncation in the simulation of liquids and solvated ions, ion pairs, and biomolecules are well documented.” More recent studies, e.g., Refs. 15–17 show the same conclusion.

If the charge distribution is periodic, one can use the Ewald method^{18–20} that scales as $O(N^{3/2})$. It is good for up to 10^3 – 10^4 particles and can achieve high accuracy; it is clearly faster than full direct summation of the periodic particle images.

The Ewald method was further improved by the use of fast Fourier transform that led to the particle-mesh schemes.^{21–32} These algorithms scale as $O(N \log N)$. The same scaling was obtained by combining the Ewald method with the FMM for the calculation of the short range forces³³ and, alternatively, by using nonuniform FFT in the Ewald method.³⁴ The particle-mesh schemes are suitable¹¹ for relatively high rms force error of about 10^{-4} and $N=10^4$ – 10^5 .

The FMM method can be applied in the periodic case too^{6,26,35–37} especially if one needs high precision. Due to its better scaling, it becomes preferable in the limit of large N . However, for not so big N the particle-mesh schemes are faster. The exact location of the crossover point depends both

on the system studied and the computer used.^{11,37} The usual estimate for the crossover is around $N=10^5$ – 10^6 .

It is mentioned in Refs. 1 and 38 that in the FMM the forces are not equal to the negative analytical gradients of the energy. Therefore if one uses the FMM for a MD simulation, the total energy will change during the simulation, unless very high precision is used, see, e.g., Fig. 1 of Ref. 36. Thus it is impossible to do simulations in the microcanonical ensemble.

In the particle-mesh schemes it is possible to have conservation of either energy or momentum, but not the two together.^{24,25} The particle-mesh schemes are usually also easier to code, compared to the FMM.

Much attention has been given to the parallelization of the particle-mesh^{26,29,39–41} and FMM^{42–45} algorithms. Recent papers^{46,47} cite speedups of over 100 for 10^5 particles for a state of the art parallel implementation of FMM. The most well-known massively parallel version of the particle-mesh algorithm is that in the NAMD molecular dynamics package.⁴⁸ It cites speedups of over 500 for 10^5 particles and 10^4 processors.⁴⁹ Similar achievements for the PMEMD code of the Amber suite⁵⁰ can be found on its website.^{51,52} A comparison of performance of several such molecular dynamics packages can be found in Ref. 53.

Parallelization of particle-mesh codes has been tried also by the astrophysics community,⁵⁴ but there the main concern is the clustering that results from the attractive gravitational forces. This calls for adaptive schemes^{55,56} which are beyond the scope of the present paper.

Many systems ranging from those in the electronic structure calculations of molecules to those of astrophysics require the use of free boundary conditions (BC). Imposition of artificial periodicity that is needed for the Ewald-type methods may lead to artifacts, e.g., in the implicit solvent systems.^{14,57} On the other hand, for the explicit solvent case no significant artifacts were found.⁵⁸

In general the Ewald artifacts are considered less harmful than those of the cutoff schemes^{48,59,60} but it would still be useful to have a particle-mesh-like method that could handle free BC.

Such algorithms were set forth in Refs. 61–63. The method of Hockney and Eastwood²¹ was intended for applications in plasma physics where high precision is not required. Martyna and Tuckerman⁶² used the smooth particle-mesh Ewald (SPME) method of Essman *et al.*²⁴ along with a special scheme to handle free BC. That scheme was also used for the solution of Poisson equation for smooth charge distributions occurring in electronic structure calculations. However, the method of Ref. 62 has the disadvantage that the accuracy decreases if one approaches the border of the cell. Thus one needs to take bigger cells that leads to an increase of the CPU time needed.

Another line of the particle-mesh-like algorithms stems from the fast Fourier Poisson method (FFP).^{23,29} In that case one solves the Poisson equation for an auxiliary charge distribution that is made up from Gaussians centered at the original particle positions. The resulting charge density is then projected onto a grid and the thus discretized Poisson equation is solved via a FFT. The forces are obtained by analytical differentiation of the energy expression; this solves the aforementioned problem of energy conservation. This method is believed to be accurate but slow¹ compared to the standard particle-mesh schemes such as the particle-particle, particle-mesh (P3M) one of Ref. 21 or SPME.²⁴

Developing the idea of Ref. 64, the authors of Ref. 29 proposed to use multigrid⁶⁵ for the solution of the Poisson equation, instead of the fast Fourier transform. This technique, now known as lattice Gaussian multigrid (LGM), was further developed in Refs. 31 and 32. The CPU time in the LGM algorithm scales as $O(N)$; however, due to large prefactor it becomes preferable to the FFT-based methods only for a very large N and/or on a massively parallel computer.

Multigrid is in general not applicable with free BC. However, in the new method of Sutmann and Steffen⁶³ the free BC were accounted for by calculating the potential at the boundaries with the FMM algorithm. This algorithm may be a competitor to the FMM algorithm in the future. Another particle-mesh-like algorithm that uses a multiscale algorithm instead of FFT is described in Ref. 30.

In a recent paper⁶⁶ we proposed a new algorithm for the solution of the Poisson equation for smooth charge distribution with free BC. It is also parallelized with Message-Passing Interface (MPI).

In the present paper we propose a combination of the Poisson solver with free BC (Ref. 66) with the fast Fourier Poisson method.²³ We call it particle-particle particle-scaling function (P3S) algorithm to emphasize the relation to the particle-mesh schemes. Our method can achieve high accuracy and speed and compete with the FMM schemes in the range of particle numbers $N=10^3-10^5$ that is important in many applications. The approximate forces resulting from our method are exact (negative) analytic gradients of the approximate energy, which allows the conservation of energy during molecular dynamics (MD) runs.

To calculate the short range forces, we make a linked list.¹⁹ Then, following Refs. 67 and 68, we rearrange the particles to optimize the cache performance.

A MPI parallelized version of our code can be downloaded under the GNU General Public License from the website of our group.⁷⁴

The paper is organized as follows. Section II describes the Ewald construction in free BC. In Sec. III we briefly review the Poisson solver of Ref. 66 for the calculation of the long range Ewald forces and energies for free BC. In Secs. IV and V we discuss the cutoffs for the short and long range forces, respectively. In Sec. VI we describe the linked cell list for the acceleration of the short range force calculation. Section VII contains the final formulas for the interparticle forces. In Sec. VIII we give the results of our code for systems of 1000–20 000 particles and give the graphs of the optimal parameter values for a given accuracy. Section IX describes the results of a MD simulation of a 1000-particle NaCl crystal that demonstrates the energy conservation property of our algorithm. Finally, in Sec. X we describe an efficient parallelization of our algorithm and present the parallel speedup results on a CRAY XT3.

II. THE EWALD CONSTRUCTION IN FREE BOUNDARY CONDITIONS (REF. 25)

The total electrostatic energy of N point charges in free BC is given by

$$U = \frac{1}{2} \sum_{i=1}^N \sum_{j \neq i}^N \frac{Q_i Q_j}{r_{ij}}, \quad r_{ij} \equiv |\mathbf{r}_i - \mathbf{r}_j|.$$

Adding and subtracting the term corresponding to the electrostatic energy of smooth point charges with density $\rho_i(\mathbf{r})$, we get

$$\begin{aligned} U = & \frac{1}{2} \sum_{i=1}^N \sum_{j \neq i}^N \left[\frac{Q_i Q_j}{r_{ij}} - \int \frac{\rho_i(\mathbf{r}) \rho_j(\mathbf{r}')}{|\mathbf{r} - \mathbf{r}'|} d\mathbf{r} d\mathbf{r}' \right] \\ & + \frac{1}{2} \sum_{i=1}^N \sum_{j=1}^N \int \frac{\rho_i(\mathbf{r}) \rho_j(\mathbf{r}')}{|\mathbf{r} - \mathbf{r}'|} d\mathbf{r} d\mathbf{r}' \\ & - \frac{1}{2} \sum_{i=1}^N \int \frac{\rho_i(\mathbf{r}) \rho_i(\mathbf{r}')}{|\mathbf{r} - \mathbf{r}'|} d\mathbf{r} d\mathbf{r}'. \end{aligned}$$

The Ewald choice for the screening charge distribution is

$$\rho_i(\mathbf{r}) = \gamma(\mathbf{r} - \mathbf{r}_i), \quad \gamma(\mathbf{r}) \equiv Q_i (G^2/\pi)^{3/2} \exp[-G^2 \mathbf{r}^2]. \quad (1)$$

We have also experimented with the variant $\gamma(\mathbf{r}) = A(r_{\text{cut}}^2 - r^2)^m$, $m=4, 8, 16$, where the factor A normalizes the charge to one. However, the CPU time spent by the program with this screening distribution was slightly longer than with the Gaussian for the same accuracy.

The sum of the screening charge distribution is

$$\rho(\mathbf{r}) = \sum_{i=1}^N \rho_i(\mathbf{r}). \quad (2)$$

Then, Eq. (1) assumes the form

$$U = E_{\text{short}} + E_{\text{long}} - E_{\text{self}}, \quad (3)$$

$$E_{\text{short}} = \frac{1}{2} \sum_{i=1}^N \sum_{j \neq i}^N \frac{Q_i Q_j}{r_{ij}} \operatorname{erfc} \left(\frac{G r_{ij}}{\sqrt{2}} \right), \quad (4)$$

$$E_{\text{long}} = \frac{1}{2} \int \frac{\rho(\mathbf{r}) \rho(\mathbf{r}')}{|\mathbf{r} - \mathbf{r}'|} d\mathbf{r} d\mathbf{r}', \quad E_{\text{self}} = \frac{G}{\sqrt{2\pi}} \sum_{i=1}^N Q_i^2. \quad (5)$$

III. CALCULATION OF THE LONG RANGE ENERGY USING THE INTERPOLATING SCALING FUNCTIONS

The long range term E_{long} in Eq. (5) is nothing but the electrostatic energy of the smooth charge distribution $\rho(\mathbf{r})$. Therefore it can be evaluated by a suitable Poisson solver. We make use of the one from Ref. 66 with free BC. It amounts to expanding the screening charge density (2) in a real space basis defined on the grid with spacing h ,

$$\rho(\mathbf{r}) \approx \tilde{\rho}(\mathbf{r}) = \sum_{\mathbf{i}} \rho_{\mathbf{i}} \Phi_{\mathbf{i}}(\mathbf{r}), \quad \mathbf{i} \equiv (i_1, i_2, i_3),$$

$$\Phi_{\mathbf{i}}(\mathbf{r}) = \Phi(x/h - i_1) \Phi(y/h - i_2) \Phi(z/h - i_3). \quad (6)$$

It was suggested in Ref. 66 to take the interpolating scaling functions^{69,70} of high order (up to 100) as the basis functions $\phi(x)$. The scaling functions of the order N interpolate the polynomials of the order N exactly and are reasonably localized. Therefore they can interpolate a Gaussian very well.

On the other hand, since the scaling functions are cardinal, we obtain for the coefficient in Eq. (6),

$$\rho_{\mathbf{i}} = \rho(\mathbf{i}h). \quad (7)$$

The action of calculating this screening distribution on the grid is called the charge assignment in the standard P3M schemes.

Consider the potential that arises from the approximate charge distribution $\tilde{\rho}(\mathbf{r})$ in Eq. (6),

$$\phi(\mathbf{r}) = \int \frac{\tilde{\rho}(\mathbf{r}')}{|\mathbf{r} - \mathbf{r}'|} d\mathbf{r}'. \quad (8)$$

At a grid point \mathbf{j} , this potential has the form

$$\phi(\mathbf{j}h) \equiv \phi_{\mathbf{j}} = h^2 \sum_{\mathbf{i}} K_{\mathbf{i}-\mathbf{j}} \rho_{\mathbf{i}},$$

where the kernel

$$K_{\mathbf{i}} \equiv \int \frac{\Phi_{\mathbf{i}}(\mathbf{r}) d\mathbf{r}}{|\mathbf{r}|}$$

is computed in Ref. 66.

From this moment on we can use the grid sum approximation to the long range energy,

$$E_{\text{long}} \approx \frac{h^3}{2} \sum_{\mathbf{i}} \rho_{\mathbf{i}} \phi_{\mathbf{i}} = \frac{h^5}{2} \sum_{\mathbf{ij}} \rho_{\mathbf{i}} \rho_{\mathbf{j}} K_{\mathbf{i}-\mathbf{j}}. \quad (9)$$

The latter sum is a convolution that can be calculated via FFT techniques.⁶⁶ The energy is a product of FFT transform of $\rho_{\mathbf{i}}$ squared times the FFT transform of $K_{\mathbf{i}}$, times a constant.

The kernel $K_{\mathbf{i}}$ is calculated only once at the beginning of a calculation and does not change. Thus one only needs one FFT to compute the energy.

It also follows from the above that the use of high order of interpolation does not lead to a significant increase of calculation time.⁶⁶

IV. THE LONG RANGE PART CUTOFF

The density array in Eq. (9) is defined via Eq. (7),

$$\rho_{\mathbf{i}} = \rho(\mathbf{i}h) = \sum_{j=1}^N \gamma(\mathbf{i}h - \mathbf{r}_j).$$

However, the Gaussian (1) is a quickly decaying function. Therefore, one can make the summation in the above equation only for the charges within the distance x_{cut} from the grid point,

$$\rho_{\mathbf{i}} \approx \sum_{|\mathbf{i}h - \mathbf{r}_j| < x_{\text{cut}}} \gamma(\mathbf{i}h - \mathbf{r}_j). \quad (10)$$

The cutoff distance is chosen so that the function (1) has a small value at the radius x_{cut} .

Actually, the sum in Eq. (10) is calculated slightly differently in our program. The charge position \mathbf{r}_j is replaced by the position $h\mathbf{i}_j$ of the grid point that is nearest to it,

$$\rho_{\mathbf{i}} \approx \sum_{(\mathbf{i} - \mathbf{i}_j)^2 < x_{\text{cut}}^2/h^2} \gamma(h(\mathbf{i} - \mathbf{i}_j)). \quad (11)$$

The error entailed by this is negligible. Then we precalculate and store the square roots of integers in order to determine to which gridpoints a given charge contributes. This is faster than calculating square roots of real numbers and truncating to integers in Eq. (10).

As a result the calculation time of the charge spreading and long range force interpolation is reduced almost by a factor of 2 compared to a simple summation over a rectangular domain. Note that this is only possible because in the FFP algorithm the charge density assigned to the grid is made of Gaussians. In other particle-mesh schemes B splines are used instead, so our method would be inapplicable there.

Making the cutoff approximation in Eq. (9) gives us, finally,

$$E_{\text{long}} \approx \frac{h^5}{2} \sum_{\mathbf{ij}} \rho_{\mathbf{i}} \rho_{\mathbf{j}} K_{\mathbf{i}-\mathbf{j}}. \quad (12)$$

The long range energy of our algorithm is given only by the formulas (11) and (12). Our algorithm is thus as simple in formulation as the fast Fourier Poisson method.²³ On the other hand, the properties of interpolating scaling functions allow us to combine fast Gaussian charge assignment and very high order interpolation.

To illustrate the accuracy of our approximation, we calculate the self-energy of a single unit Gaussian (1) by the formula (12). The Gaussian is initially at a grid point and then gradually shifted by two grid constants to see the dependence of the error on its position relative to the grid. We have $x_{\text{cut}} = 4.5h$ and $h = 0.7$; these are typical parameter values that we use in other tests.




FIG. 1. Relative error in the self-interaction energy of a single Gaussian.

The graph of the resulting energy error versus the shift distance is given in Fig. 1. We see that the error is rather smooth; there is no visible jumps coming from the cutoff approximation. The energy oscillation is due to the grid discretization; its amplitude goes to zero in the limit $h \rightarrow 0$.

V. THE SHORT RANGE PART CUTOFF

The short range pairwise potential in Eq. (4) is a rapidly decaying function. Therefore, it is standard to introduce the cutoff radius r_{cut} beyond which the short range potential is approximated by zero,

$$E_{\text{short}} = \frac{1}{2} \sum_{i,j:r_{ij} < r_{\text{cut}}} \frac{Q_i Q_j}{r_{ij}} \operatorname{erfc}\left(\frac{Gr_{ij}}{\sqrt{2}}\right). \quad (13)$$

To summarize, the Coulombic energy of N charges in our approximation has the form (3) with the approximate long and short contributions given by Eqs. (12) and (13), correspondingly.

The evaluation of the error function in Eq. (13) is computationally inefficient, compared to algebraic operations. Therefore, we replace the term that corresponds to error function by a spline approximation in the interval $(0, r_{\text{cut}})$.

The function $\operatorname{erfc}(Gr/\sqrt{2})/r$ has a singularity at 0. Therefore we rewrite it as

$$\operatorname{erfc}(Gr/\sqrt{2})/r = 1/r - \operatorname{erf}(Gr/\sqrt{2})/r, \quad (14)$$

and approximate with a spline only the second term which is regular. This approximation holds for $r < r_{\text{cut}}$. For $r > r_{\text{cut}}$, the whole potential $\operatorname{erfc}(Gr/\sqrt{2})/r$ is set to 0.

VI. CALCULATING THE FORCES

The force acting on the i th particle is given by the gradient of Eq. (3),

$$\mathbf{F}_i = \frac{\partial U}{\partial \mathbf{r}_i} = \mathbf{F}_i^{\text{short}} + \mathbf{F}_i^{\text{long}}, \quad (15)$$

$$\mathbf{F}_i^{\text{short}} = Q_i \sum_{j:r_{ij} < r_{\text{cut}}} Q_j \left(\frac{G}{r_{ij}^2} \sqrt{\frac{2}{\pi}} \exp\left[-\frac{G^2 r_{ij}^2}{2}\right] + \frac{1}{r_{ij}^3} \operatorname{erfc}\left[\frac{Gr_{ij}}{\sqrt{2}}\right] \right) \mathbf{r}_{ij},$$

where $\mathbf{r}_{ij} = \mathbf{r}_i - \mathbf{r}_j$; the prime at the sum indicates that it is performed over $j \neq i$. The short range forces (15) are equal to the negative analytical gradients of the short range energy (4). They are also calculated via a spline approximation. Actually, the spline that we use for Eq. (14) is obtained by integration of the spline for Eq. (15).

Following Ref. 29, we also get the long range forces as the negative analytical gradients of Eq. (12),

$$\mathbf{F}_i^{\text{long}} \approx h^5 \sum_{ij} \frac{\partial \rho_i}{\partial \mathbf{r}_i} \rho_j K_{i-j} \equiv h^5 \sum_{ij} \mathbf{q}(lh - \mathbf{r}_i) \rho_j K_{i-j}, \quad (16)$$

where

$$\mathbf{q}(\mathbf{r}) = \frac{\partial \gamma(\mathbf{r})}{\partial \mathbf{r}} = -2\mathbf{r}G^2 \gamma(\mathbf{r}) \quad (17)$$

is the vector of derivative Gaussians. The calculation of the long range forces (16) is called the force back interpolation.

Taking into account the cutoff (11), one can rewrite the long range force as

$$\mathbf{F}_i^{\text{long}} = h^5 \sum_{j:i:(i-j)^2 < x_{\text{cut}}^2/h^2} \mathbf{q}(lh - \mathbf{r}_i) \rho_j K_{i-j}.$$

The convolution $\sum_j \rho_j K_{i-j}$ is done using two scalar FFTs. The sum over \mathbf{i} can then be calculated fast because each component of Eq. (17) is a product of three one-dimensional functions.

VII. LINKED CELL LIST

In order to avoid the scanning of all particle pairs in Eq. (13) we make a linked subcell list.¹⁹ The subcell size is smaller or equal to r_{cut}/M where M is a positive integer ranging, in practice, from 1 to 3. The standard linked cell list algorithm corresponds to $M=1$; higher values of M correspond to smaller cells such as those of Refs. 67 and 68. To account for the free BC we add M layers of empty cells outside the original cell. The sum of forces for a given particle is then done over all particles in subcells which are within r_{cut} of the subcell where the original particle is located.

The sum over subcells is done along the stripes in the direction x . In particular, following Refs. 54, 55, 67, and 68 we rearrange the particles so that the particles in the same subcell have consecutive numbers and the subcells are sorted in the x direction. The actual sorting can be avoided because we already have the linked subcell list.

The reordering of the particles has the additional advantage that the cache performance of the long range part is optimized. The reason is that the particles that have close indices are also close physically: $|\mathbf{r}_i - \mathbf{r}_j|$ is small when $|i - j|$ is small.

The loop in the charge assignment and force back interpolation goes over the particle numbers. The particles in the

same subcell have overlapping screening charge densities, so the elements of the density/potential array are reused again and again, as the loop index goes over the particles in the same subcell.

This makes the program significantly faster compared to the case without reordering, especially if the particles are distributed randomly.

VIII. THE RESULTS OF THE SERIAL CODE

We chose the following two test systems:

- N particles in the unit cube with random coordinates and charges equal to ± 1 , the total charge being zero (henceforth called the random system).
- N particles with charges ± 1 forming a rock salt crystal, filling the unit cube with $M \approx N^{1/3}$ nodes at each side. The lattice constant is then $d = (M-1)^{-1}$. Each particle is then shifted away from its node by a vector with random coordinates in the range $(-d/3, d/3)$. This system will be called the crystal one.

In both cases, the particle number $N = 1000 \times 10^{j/3}$, $j = 0, \dots, 4$. In other words, we consider the following values of N : 1000, 2154, 4642, 10 000, and 21 544.

We always use the scaling functions of the order 100. We found that using lower orders leads to a decline in accuracy, while the order of the scaling functions only affects the time of the calculation of the kernel which is done only once for a simulation.

The values of x_{cut} , r_{cut} , G , and h are free parameters. As explained, e.g., in Ref. 72, such parameters should be chosen such that the accuracy is fixed to some value, and the CPU time spent is minimal for the given accuracy.

As the measure of the accuracy we use the mean square force error

$$\delta F \equiv \sqrt{\frac{\sum_{i=1}^N (\mathbf{F}_i - \mathbf{F}_i^{\text{direct}})^2}{\sum_{i=1}^N (\mathbf{F}_i^{\text{direct}})^2}},$$

where

$$\mathbf{F}_i^{\text{direct}} = Q_i \sum_{j \neq i} Q_j \frac{\mathbf{r}_i - \mathbf{r}_j}{|\mathbf{r}_i - \mathbf{r}_j|}$$

are the forces obtained by the full direct calculation.

Unfortunately we cannot give the *a priori* error estimates of the accuracy. The error in the short range forces is the same as that for the Ewald method and is estimated in Refs. 71 and 72. However, our calculation of the long range forces involves the use of a finite basis and the approximation in Eq. (9), the accuracy of which is difficult to estimate.

Therefore, for our two test systems we made tests with all reasonable values of the parameters. The most important parameter is G , the width of the Gaussian in Eq. (1). We also introduced the following dimensionless quantities: the dimensionless Gaussian cutoff Gx_{cut} , the cutoff ratio $r_{\text{cut}}/x_{\text{cut}}$, the dimensionless grid constant Gh , and




FIG. 2. (Color online) The Pareto frontiers for the random system on Opteron 244.

$$N_{\text{short}} \equiv \frac{4/3 \pi r_{\text{cut}}^3}{V_{\text{box}}} N.$$

The latter is the average number of particles inside a sphere of radius r_{cut} . It turns out that the rms force error depends stronger on these dimensionless quantities and G than on the number of particles in the system.

We checked seven to ten possible values for each parameter. From the resulting pool of results we selected the so-called Pareto frontier. A point is on the Pareto frontier if there is no other point which has both smaller CPU time and smaller error.

The Pareto frontiers for the values of $N = 1000 \times 10^{j/3}$, $j = 0, \dots, 4$ for the random system and the crystal systems are given in Figs. 2 and 3. The tests were done on an 1.8 GHz AMD Opteron 244.

For the purpose of illustration we present here the optimal parameter values for the crystal system. The values for the random one do not differ significantly. From Figs. 4 and 5 we see that the optimal values of Gx_{cut} and Gh are determined by the accuracy level. In contrast, the optimal value of N_{short} depends on the number of particles too (Fig. 6). There-





FIG. 3. (Color online) The Pareto frontiers for the crystal system.



FIG. 4. (Color online) The optimal values of $G_{x_{\text{cut}}}$, crystal system.

fore in an actual calculation it must be adjusted by the trial and error method to give optimal accuracy and CPU time. This is similar to finding the optimal value of the Gaussian width in the standard particle-mesh schemes.²⁸ Finally, the optimal ratio $r_{\text{cut}}/x_{\text{cut}}$ was observed to be independent on the required accuracy but slightly decreased with rising N , as seen in Fig. 7.

The Pareto frontiers allow us to determine the crossover points for each N : the values of MSQ force error for which our calculation takes the same time as the full direct one. They are presented in Fig. 8. We have plotted the crossover curves for the random system and for the crystal system. For comparison we have also plotted the crossover curve of the fast multipole method taken from Ref. 5, where one of the best implementations of FMM is described. In that paper the same charge distribution was used as in our random system. Of course, the random positions of the particles were different in our tests and in those of Ref. 5. However, we still see that our method has lower crossovers than the FMM for the same accuracy.

IX. THE ENERGY CONSERVATION

In contrast to other methods such as FMM,^{1,38} our algorithm has the advantage that the approximate forces are exact

FIG. 5. (Color online) The optimal values of G_h , crystal system.FIG. 6. (Color online) The optimal values of N_{short} , crystal system.


analytic derivatives of the approximate energy. This allows for energy conservation during a MD run. To illustrate this, we make a MD simulation of a rock salt crystal formed by 1000 Na and Cl atoms. The particle positions and velocities are updated by the velocity Verlet algorithm.

To get physically reasonable results, we made the particles interact through the Born-Mayer-Huggins-Fumi-Tosi (BMHFT) rigid-ion potential⁷³ that has bonding terms in addition to the Coulombic force.

At first we made the system equilibrate for 300 oscillation periods. We then monitored the potential and total energies for another 100 periods using the full direct algorithm. Then the last 100 periods were repeated using our P3S algorithm.

On Figs. 9 and 10 we plot the absolute values of deviations of the potential and total energy from their mean values. The ratio of the mean square deviation of the total energy to that of the potential one is found to be equal to 1.4×10^{-3} .

The P3S results are shown on the graphs; those of the full direct calculation are indistinguishable from the P3S ones.

FIG. 7. Values of $r_{\text{cut}}/x_{\text{cut}}$ for both random and crystal systems.

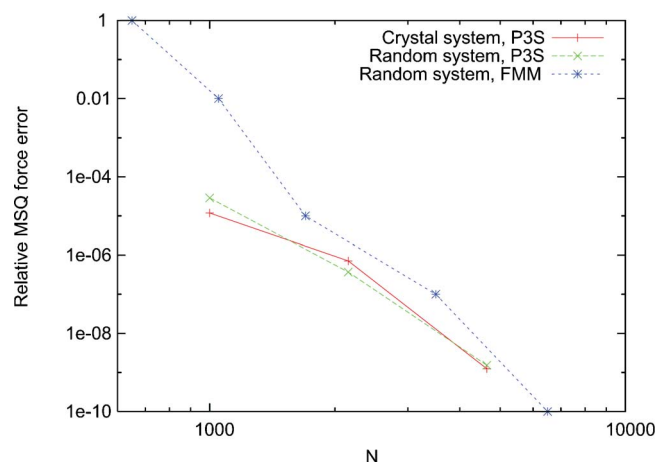


FIG. 8. (Color online) The crossover curves.

X. PARALLELIZATION

The parallelization of the calculation of the short range energy and forces is straightforward: for the particle i and processor IPROC, we calculate the force only if $\text{MODULO}(I, \text{NPROC}) = \text{IPROC}$, where NPROC is the total number of processors. In the end, an MPI_ALLREDUCE command sums all the contributions.

For the long range part, we rely on the parallel structure of the Poisson solver.⁶⁶ The charge density on the grid is divided into slabs in the x - y plane. Each such slab is the input density at a separate node. The Poisson equation (for the whole cell) is solved for that slab density. The output at a given node again contains only the corresponding slab of the grid, this time it is a piece of the potential array. Of course one needs global interprocessor communications, including MPI_ALLTOALL for the solution of the Poisson equation.

Then in the original algorithm of Ref. 66, an MPI_ALLREDUCE command sums up the potentials of all the slabs.

In our program for point particles, we kept the parallelization of the charge assignment as above: each processor receives only the grid charges from the corresponding slab. However, we do not use the MPI_ALLREDUCE command to get the potential. Instead, we used the slabwise structure




FIG. 9. The potential energy fluctuations with the P3S method.




FIG. 10. The total energy fluctuations with the P3S method.

of the output of the Poisson solver. For each node, the corresponding slab potential contributes only to the forces on particles that are close to it.

In the end, we add up the *forces* with the MPI_ALLREDUCE command. This MPI_ALLREDUCE is actually merged in the program with the one for the short range forces.


In this way we minimize the interprocessor communication considerably compared, e.g., to Ref. 26 since it is much easier to send the components of the forces than the pieces of the enormous potential array.

To test the parallel program we ran it for the random system (the same as in the serial tests) with 10^5 particles. The relative accuracy of the forces was kept around 10^{-6} .

The result is given in Fig. 11. At the y axis we have the ratio of the CPU time spent on several processors to that spent by one processor.

XI. CONCLUSION

We have developed a point particle Poisson solver algorithm that has a lower crossover point than the FMM. It can be considered as a generalization of the particle-mesh solvers for free boundary conditions. It can also achieve high precision.

FIG. 11. The parallel speedup results on a CRAY XT3 for 10^5 particles.

The forces obtained by our program are analytical derivatives of the energy; this is an advantage in the context of MD simulations. An MPI parallelized version of the algorithm is presented that scales well on a moderate number of processors.

ACKNOWLEDGMENTS

The authors acknowledge financial support from the Swiss National Science Foundation. The numerical calculations were done on a parallel supercomputer of the Swiss National Supercomputing Center.

- ¹C. Sagui and T. A. Darden, *Annu. Rev. Biophys. Biomol. Struct.* **28**, 155 (1999).
- ²K. Spurzem and A. Kugel, in *Molecular Dynamics on Parallel Computers*, Workshop of the John von Neumann-Institute for Computing (NIC) Jülich, edited by R. Esser, P. Grassberger, J. Grotendorst, and M. Lewerenz (World Scientific, Singapore, 1999).
- ³L. Greengard and V. Rokhlin, *J. Comput. Phys.* **73**, 325 (1987).
- ⁴H.-Q. Ding, N. Karasawa, and W. Goddard, *J. Chem. Phys.* **97**, 4309 (1992).
- ⁵C. A. White and M. Head-Gordon, *J. Chem. Phys.* **101**, 6593 (1994).
- ⁶J. Shimada, H. Kaneko, and T. Takada, *J. Comput. Chem.* **15**, 28 (1994).
- ⁷L. F. Greengard, *Science* **265**, 90914 (1994).
- ⁸S. Pfalzner and P. Gibbon, *Many-body Tree Methods in Physics* (Cambridge University Press, Cambridge, 1996).
- ⁹K. Esselink, *Comput. Phys. Commun.* **87**, 375 (1995).
- ¹⁰J. Barnes and P. Hut, *Nature (London)* **324**, 446 (1986).
- ¹¹P. Gibbon and G. Sutmann, in *Quantum Simulations of Complex Many-Body Systems: From Theory to Algorithms*, NIC Series Vol. 10, edited by J. Grotendorst, D. Marx, and A. Muramatsu (John von Neumann Institute for Computing, Jülich, 2002), pp. 467–506.
- ¹²D. Wolf, P. Kéblinski, S. R. Phillpot, and J. Eggebrecht, *J. Chem. Phys.* **110**, 8255 (1999).
- ¹³C. J. Fennell and J. D. Gezelter, *J. Chem. Phys.* **124**, 234104 (2006).
- ¹⁴P. H. Hünenberger and J. A. McCammon, *Biophys. Chem.* **78**, 69 (1999).
- ¹⁵M. Patra, M. Karttunen, M. T. Hyvnen, P. Lindqvist, E. Falck, and I. Vattulainen, *Biophys. J.* **84**, 3636 (2003).
- ¹⁶C. Anezo, A. H. de Vries, H.-D. Höltje, D. P. Tieleman, and S.-J. Marrink, *J. Phys. Chem. B* **107**, 9424 (2003).
- ¹⁷M. Patra, M. Karttunen, M. T. Hyvnen, E. Falck, and I. Vattulainen, *J. Phys. Chem. B* **108**, 4485 (2004).
- ¹⁸P. P. Ewald, *Ann. Phys.* **64**, 253 (1921).
- ¹⁹M. P. Allen and D. J. Tildesley, *Computer Simulation of Liquids* (Clarendon, New York, 1988).
- ²⁰A. Toukmaji and J. A. Board, *Comput. Phys. Commun.* **95**, 78 (1996).
- ²¹R. W. Hockney and J. W. Eastwood, *Computer Simulation Using Particles* (IOP, Bristol, 1988).
- ²²T. A. Darden, D. M. York, and L. G. Pedersen, *J. Chem. Phys.* **98**, 10089 (1993).
- ²³D. York and W. Yang, *J. Chem. Phys.* **101**, 3298 (1994).
- ²⁴U. Easman, L. Perera, M. L. Berkowitz, T. Darden, H. Lee, and L. Pedersen, *J. Chem. Phys.* **103**, 8577 (1995).
- ²⁵P. MacNeice, NASA Contractor Report No. 4666, 1995 (unpublished).
- ²⁶E. L. Pollock and J. N. Glosli, *Comput. Phys. Commun.* **95**, 93 (1996).
- ²⁷M. Deserno and C. Holm, *J. Chem. Phys.* **109**, 7678 (1998).
- ²⁸P. H. Hünenberger, *J. Chem. Phys.* **113**, 10464 (2000).
- ²⁹C. Sagui and T. Darden, *J. Chem. Phys.* **114**, 6578 (2001).
- ³⁰R. D. Skeel, I. Tezcan, and D. J. Hardy, *J. Comput. Chem.* **23**, 673 (2002).
- ³¹S. Banerjee and J. A. Board, *J. Comput. Chem.* **26**, 957 (2005).
- ³²Y. Shan, J. L. Klepeis, M. P. Eastwood, R. O. Dror, and D. E. Shaw, *J. Chem. Phys.* **122**, 054101 (2005).
- ³³Z.-H. Duan and R. Krasny, *J. Chem. Phys.* **113**, 3492 (2000).
- ³⁴F. Hedman and A. Laaksonen, *Chem. Phys. Lett.* **425**, 142 (2006).
- ³⁵K. E. Schmidt and M. A. Lee, *J. Stat. Phys.* **63**, 1223 (1991).
- ³⁶F. Figuerido, R. M. Levy, R. Zhou, and B. Berne, *J. Chem. Phys.* **106**, 9835 (1997).
- ³⁷M. Challacombe, C. White, and M. Head-Gordon, *J. Chem. Phys.* **107**, 10131 (1997).
- ³⁸T. C. Bishop, R. D. Skeel, and K. Schulten, *J. Comput. Chem.* **18**, 1785 (1997).
- ³⁹A. Toukmaji, D. Paul, and J. A. Board, Jr., Duke University Technical Report No. 96-002 (unpublished).
- ⁴⁰I. F. Sbalzarini, J. H. Walther, M. Bergdorf, S. E. Hieber, E. M. Kotsalis, and P. Koumoutsakos, *J. Comput. Phys.* **215**, 566 (2006).
- ⁴¹A. V. Philippov, V. V. Vlutin, and R. A. Veselov, *Rev. Sci. Instrum.* **75**, 1499 (2004).
- ⁴²J. A. Board, J. W. Causey, J. F. Leathrum, A. Windemuth, and K. Schulten, *Chem. Phys. Lett.* **198**, 89 (1992).
- ⁴³H. Schwichtenberg, G. Winter, and H. Wallmeier, *Parallel Comput.* **25**, 535 (1999).
- ⁴⁴J. A. Lupo, Z. Wang, A. M. McKenney, R. Pachter, and W. Mattson, *J. Mol. Graphics Modell.* **21**, 89 (2002).
- ⁴⁵S. Ogata, T. J. Campbell, R. K. Kalia, A. Nakano, P. Vashishta, and S. Vemparala, *Comput. Phys. Commun.* **153**, 445 (2003).
- ⁴⁶J. Kurzak and B. M. Pettitt, *J. Comput. Phys.* **203**, 731 (2005).
- ⁴⁷J. Kurzak and B. M. Pettitt, *J. Parallel Distrib. Comput.* **65**, 870 (2005).
- ⁴⁸J. C. Phillips, R. Braun, W. Wang, J. Gumbart, E. Tajkhorshid, E. Villa, C. Chipot, R. D. Skeel, L. Kale, and K. Schulten, *J. Comput. Chem.* **26**, 1781 (2005).
- ⁴⁹J. C. Phillips, G. Zheng, S. Kumar, and L. V. Kal, Proceedings of the IEEE/ACM SC2002 Conference (IEEE, New York, 2002). Technical Paper No. 277.
- ⁵⁰D. A. Case, T. E. Cheatham III, T. Darden, H. Gohlke, R. Luo, K. M. Merz, Jr., A. Onufriev, C. Simmerling, B. Wang, and R. J. Woods, *J. Comput. Chem.* **26**, 1668 (2005).
- ⁵¹<http://amber.scripps.edu/amber9.bench1.html>
- ⁵²<http://amber.scripps.edu/amber9.bench2.html>
- ⁵³J. Hein, F. Reid, L. Smith, I. Bush, M. Guest, and P. Sherwood, *Philos. Trans. R. Soc. London, Ser. A* **363**, 1987 (2005).
- ⁵⁴P. P. Brieu and A. E. Evrard, *New Astron.* **5**, 163 (2000).
- ⁵⁵T. MacFarland, H. M. P. Couchman, F. R. Pearce, and J. Pichlmeier, *New Astron.* **3**, 687 (1998).
- ⁵⁶R. J. Thacker and H. M. P. Couchman, *Comput. Phys. Commun.* **174**, 540 (2006).
- ⁵⁷P. H. Hünenberger and J. A. McCammon, *J. Chem. Phys.* **110**, 1856 (1999).
- ⁵⁸M. A. Villarreal and G. G. Montich, *J. Biomol. Struct. Dyn.* **23**, 135 (2005).
- ⁵⁹G. Hummer, L. Pratt, and A. Garsia, *J. Phys. Chem.* **100**, 1206 (1996).
- ⁶⁰S. E. Feller, R. W. Pastor, A. Rojnuckarin, S. Bogusz, and B. R. Brooks, *J. Phys. Chem.* **100**, 17011 (1996).
- ⁶¹R. W. Hockney, *Methods Comput. Phys.* **9**, 135 (1970).
- ⁶²G. J. Martyna and M. E. Tuckerman, *J. Chem. Phys.* **110**, 2810 (1999).
- ⁶³G. Sutmann and B. Steffen, *Methods Comput. Phys.* **169**, 343 (2005).
- ⁶⁴J. Beckers, C. Loewe, and S. de Leeuw, *Mol. Simul.* **6**, 369 (1998).
- ⁶⁵U. Trottenberg, C. W. Oosterlee, and A. Schuller, *Multigrid* (Academic, San Diego, 2001).
- ⁶⁶L. Genovese, T. Deutsch, A. Neelov, S. Goedecker, and G. Beylkin, *J. Chem. Phys.* **125**, 074105 (2006).
- ⁶⁷T. N. Heinz and P. H. Hünenberger, *J. Comput. Chem.* **25**, 1474 (2004).
- ⁶⁸Z. Yao, J.-S. Wang, G.-R. Liu, and M. Cheng, *Comput. Phys. Commun.* **161**, 27 (2004).
- ⁶⁹G. Deslauriers and S. Dubuc, *Constructive Approx.* **5**, 49 (1989).
- ⁷⁰S. Goedecker, *Wavelets and Their Application* (Presses Polytechniques et Universitaires Romandes, Lausanne, 1998).
- ⁷¹J. Kolafa and J. W. Perram, *Mol. Simul.* **9**, 351 (1992).
- ⁷²H. G. Petersen, *J. Chem. Phys.* **103**, 3668 (1995).
- ⁷³F. G. Fumi and M. P. Tosi, *J. Phys. Chem. Solids* **25**, 45 (1964).
- ⁷⁴<http://pages.unibas.ch/comphys/comphys/SOFTWARE/index.html>