This article was downloaded by:
On: *14 January 2011*
Access details: *Access Details: Free Access*
Publisher *Taylor & Francis*
Informa Ltd Registered in England and Wales Registered Number: 1072954 Registered office: Mortimer House, 37-41 Mortimer Street, London W1T 3JH, UK

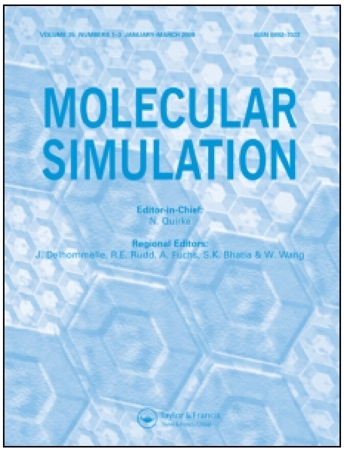

Molecular Simulation

Publication details, including instructions for authors and subscription information:
http://www.informaworld.com/smpp/title~content=t713644482

A Comparison of Particle-Particle, Particle-Mesh and Ewald Methods for Calculating Electrostatic Interactions in Periodic Molecular Systems

Brock A. Luty[a]; Malcolm E. Davis[b]; Ilario G. Tironi[a]; Wilfred F. Van Gunsteren[a]
[a] Laboratorium für Physikalische Chemie, Zürich, Switzerland [b] Bristol-Myers Squibb Pharmaceutical Research Institute, Princeton, NJ, USA

To cite this Article Luty, Brock A. , Davis, Malcolm E. , Tironi, Ilario G. and Van Gunsteren, Wilfred F.(1994) 'A Comparison of Particle-Particle, Particle-Mesh and Ewald Methods for Calculating Electrostatic Interactions in Periodic Molecular Systems', Molecular Simulation, 14: 1, 11 — 20

To link to this Article: DOI: 10.1080/08927029408022004

URL: http://dx.doi.org/10.1080/08927029408022004

PLEASE SCROLL DOWN FOR ARTICLE

Full terms and conditions of use: http://www.informaworld.com/terms-and-conditions-of-access.pdf

This article may be used for research, teaching and private study purposes. Any substantial or systematic reproduction, re-distribution, re-selling, loan or sub-licensing, systematic supply or distribution in any form to anyone is expressly forbidden.

The publisher does not give any warranty express or implied or make any representation that the contents will be complete or accurate or up to date. The accuracy of any instructions, formulae and drug doses should be independently verified with primary sources. The publisher shall not be liable for any loss, actions, claims, proceedings, demand or costs or damages whatsoever or howsoever caused arising directly or indirectly in connection with or arising out of the use of this material.

Molecular Simulation, 1994, Vol. 14, pp. 11–20
Reprints available directly from the publisher
Photocopying permitted by license only

A COMPARISON OF PARTICLE-PARTICLE, PARTICLE-MESH AND EWALD METHODS FOR CALCULATING ELECTROSTATIC INTERACTIONS IN PERIODIC MOLECULAR SYSTEMS

BROCK A. LUTY*[,1], MALCOLM E. DAVIS[2], ILARIO G. TIRONI[1] and
WILFRED F. VAN GUNSTEREN[1]

[1] *Laboratorium für Physikalische Chemie, ETH Zentrum CH-8092 Zürich, Switzerland*
[2] *Bristol-Myers Squibb Pharmaceutical Research Institute, P.O. Box 4000, Princeton, NJ 08543-4000 USA*

(Received July 1994, accepted July 1994)

We compare the Particle-Particle Particle-Mesh (PPPM) and Ewald methods for calculating electrostatic interactions in periodic molecular systems. A brief comparison of the theories shows that the methods are very similar differing mainly in the technique which is used to perform the "k-space" or mesh calculation. Because the PPPM utilizes the highly efficient numerical Fast Fourier Transform (FFT) method it requires significantly less computational effort than the Ewald method and scales almost linearly with system size.

KEY WORDS: Electrostatic interactions, particle-particle particle-mesh, Ewald summation, molecular simulation.

1 INTRODUCTION

Given finite computational resources, there will always be a limit to the number of atoms which can be explicitly included in a molecular dynamics simulation and, accordingly, there will be a surface or boundary of the system which must be considered. In the simplest approximation, one could consider the explicit system to be surrounded by a vacuum. Under this assumption, the atoms in the neighborhood of the surface will feel an imbalanced force which will distort their structure from that typically found in the interior of an extended system. If, as in many cases, we are interested in macroscopic systems, then the surface effects are undesirable artifacts which should be eliminated if possible.

The distortive effects of the vacuum outside the simulated system can be reduced by conceptually considering the system to be enclosed by an infinite implicit bath. The influence of the bath is then approximated by altering the motions of the outer layers of explicit atoms and by appropriately modifying the forces which the explicit atoms feel.

---

* Author to whom correspondence should be addressed.

Typically, the positions of the outer layer of explicit atoms are restrained and the dynamics of these atoms may be coupled with a heat bath to simulate energy exchange with the surroundings [1, 2]. These stochastic boundary methods coupled with the newly developing computational continuum electrostatic methods [3–5], for accurately estimating the forces exerted by the implicit bath on the explicit atoms, look very promising. Unfortunately at present, these methods are still computationally quite expensive and there are number of important unknown parameters which are not easily estimated *a priori* [6].

A second alternative, which is well established, is to consider the simulated region to be infinitely repeated on a periodic lattice to completely fill space. Using Periodic Boundary Conditions (PBCs) we are able to reduce the effects of having "free surfaces" at the expense of introducing a periodicity into the system. Whether the benefits of reducing the surface effects outweigh the imposed periodicity is dependent on the type of system under study. For instance, in simulations of molecules in the crystal phase the periodic boundary conditions are a good approximation and there is little question that an attempt should be made to model them accurately. In a homogeneous system, such as a solution of small ions dissolved in water, it appears that benefits of eliminating the free surfaces may outweigh the penalty of the artificial periodicity. Also, as discussed in the next paragraph, there are methods to minimize the effects of the periodicity. In non-homogeneous systems, such as the very interesting case of a macromolecule surrounded by a solvent bath, the question becomes even more difficult and, perhaps, only answerable by numerical experimentation.

Of the interactions present in a molecular system, perhaps the most important, due to their long range nature, are the electrostatic interactions. In an infinite periodic system, it is, of course, impossible to calculate the interaction between all atoms and periodic images using a straight forward application of Coulomb's law. As an alternative, one could consider including only the interactions of an atom with atoms and images out to a certain cutoff radius. For homogeneous systems, the interactions beyond the cutoff radius can be estimated very accurately using newly derived expressions for the Poisson-Boltzmann generalized reaction field [7]. For truly periodic or non-homogeneous systems, one must either approximate the true correction by the correction used for a homogeneous system or, as is commonly done, ignore the interactions outside the cutoff radius. For the non-periodic system, this technique has the benefit of reducing the effects of artificial periodicity. However, computational constraints usually limit the cutoff radius in the range of 9–10 Å, which may introduce significant errors into the calculation. There is some evidence that cutoff radii even up to 14 Å may be insufficient to avoid sizable distortions of the spatial distribution of ions or structures of small polypeptides [7, 8].

To include all interactions in a periodic system, an alternative to the direct application of Coulomb's law is to split the single function into the sum of two functions. The first function is constructed to be a rapidly varying short-ranged function which is best handled by direct summation. The second function is a smooth periodic long-ranged function which is well suited for solution by Fourier methods. Of these methods, the most popular is the Ewald method which has been applied quite successfully to molecular simulations [8, 9]. In this paper we consider, as an alternative to the Ewald method, the Particle-Particle Particle-Mesh (PPPM) method which has been fully developed by Hockney and Eastwood [10]. This method has been applied in the simulation of plasmas but has not been extensively used in molecular simulations. In constrast to the Ewald method which performs the Fourier Transforms analytically, the PPPM method relies on the very efficient numerical Fast Fourier Transform (FFT) techniques. A simple comparison of the similarities and differences of the two methods is the subject of the next section.

The remaining forces in a molecular simulation typically contain terms to account for the covalent interaction of bonded atoms and terms to account for the dispersion and repulsion between non-bonded atoms. The computational costs of evaluating the covalent interactions scale as the number of atoms in the system and is a small percentage of the total computational demand. The dispersion/repulsion terms are usually modeled using a Lennard-Jones (6–12) type potential energy function. Contrary to the Coulombic interaction, the Lennard-Jones interaction is short-ranged and a direct sum can be performed over a sufficient number of neighboring atoms. From a practical point of view, it would be desirable to have the short-range part of the Coulombic interaction operate over the same region as the Lennard-Jones interaction. A comparison of the computational efficiency of the two methods under this restriction is the subject of the third section. In the final section, we present some conclusions which can be drawn from these comparisons.

2 THEORY

Here we describe briefly the theory behind the Ewald and PPPM methods. This description is only intended to highlight the similarities and differences in the two methods and was derived from the development presented by Berendsen [11]. For rigorous derivations of the individual methods we refer the reader to the literature [10, 12].

Both the Ewald method and PPPM method separate the interaction between point charges into the sum of a short-range function and a long-range function. To accomplish this, a "shielding" charge distribution (of equal magnitude) is subtracted from the original point charge. The resulting interaction from the combination of the point charge and shielding charge is a short-range function which is best handled by a direct summation over nearest neighbors. Effects of the shielding charges, which are long-range but which will be of a much smoother nature than those of the original point charges, are then removed using Fourier methods.

2.1 Ewald Method

The Ewald method splits the electrostatic potential into the sum of a short-range potential and long-range potential. The standard Ewald method uses a three-dimensional Gaussian shaped shielding distribution:

$$\gamma_l(\vec{r}) = \frac{\alpha^3}{\pi^{(3/2)}} \exp(-\alpha^2 r^2) \tag{1}$$

The short-range electrostatic potential generated by a unit point charge minus the normalized Gaussian shielding charge distribution is [12]:

$$\phi_s(r) = \frac{1}{4\pi\varepsilon_0}\left(\frac{1}{r} - \frac{\mathrm{erf}(\alpha r)}{r}\right) \tag{2}$$

The original potential of the point charges is then found by adding the potential due to all the shielding charges. The potential of the shielding charges is found by analytically Fourier transforming Poisson's equation, solving the equation for the Fourier transformed potential, and then applying the inverse Fourier transform to yield:

$$\phi_l(\vec{r}) = \frac{1}{\varepsilon_0 V}\sum_{\vec{k}} \frac{\exp(-k^2/(4\alpha^2))}{k^2} \sum_i q_i \exp(i\vec{k}\cdot(\vec{r} - \vec{r}_i)) \tag{3}$$

2.2 *Particle-Particle Particle-Mesh Method*

We define the interaction potential as the electrostatic interaction energy between two charge distributions divided by the product of the net charges (i.e. $\psi(r_{ij}) = U(r_{ij})/(q_i q_j)$). For two charge distributions, the interaction potential can be determined by integrating the electrostatic potential generated by the first distribution over the charge distribution of the second. For point charges the interaction potential reduces to the familiar electrostatic potential.

The PPPM method separates the interaction potential into the sum of a short-range function and a long-range function. The short-range component is the sum of the interaction potential of two point charges minus the interaction potential between the two shielding charge distributions. The standard charge distribution suggested by Hockney and Eastwood is the S2 function [10]:

$$\gamma(\vec{r}) = \frac{48}{\pi a^4}\left(\frac{a}{2} - r\right) \quad r < a/2 \tag{4}$$

$$\gamma(\vec{r}) = 0 \quad r > a/2 \tag{5}$$

The short range interaction potential between two particles, each with a S2 charge distribution $\gamma(\vec{r})$, can be expressed as a polynomial in $\xi_{ij} = (2r_{ij}/a)$:

$$\psi_s(\xi_{ij}) = \frac{1}{4\pi\varepsilon_0}\left(\frac{1}{r_{ij}} - \frac{1}{70a}\sum_{n=-1}^{7} C_n \xi_{ij}^n\right) \quad \xi_{ij} < 2 \tag{6}$$

where the constants $C_{-1} \ldots C_7$ are for $0 < \xi_{ij} < 1$, $C = (0, 208, 0, -112, 0, 56, -14, -8, 3)$ and for $1 < \xi_{ij} < 2$, $C = (12, 128, 224, -448, 280, -56, -14, 8, -1)$.

The correct interaction potential between the original point charges is then found by adding the interactions between the shielding distributions. Again, by using the method of Fourier transforms, the long-range interaction can be written as:

$$\hat{\psi}_l(\vec{k}) = \frac{\hat{\gamma}(\vec{k})^2}{\varepsilon_0 k^2}\hat{\rho}(\vec{k}) = \hat{G}(\vec{k})\hat{\rho}(\vec{k}) \tag{7}$$

To solve this equation numerically, we perform the following steps:

1. Assign the charges from the atoms to a grid which encompasses the simulation region.
2. Fast Fourier Transform (FFT) the gridded charge density to obtain $\hat{\rho}(\vec{k})$.
3. Multiply $\hat{\rho}(\vec{k})$ by $\hat{G}(\vec{k})$ at each grid point to get $\hat{\psi}_l(\vec{k})$.
4. Apply the FFT again, this time to $\hat{\psi}_l(\vec{k})$ which yields $\psi_l(\vec{r})$ at the grid points.

The electrostatic fields can be found at the grid points by numerically differencing the potentials. The electrostatic potentials and fields at any point in space can then be interpolated from the known values on the grid.

In Hockney and Eastwood's "Q-minimizing method" [10], $\hat{G}(\vec{k})$ is not simply taken equal to $\hat{\gamma}(\vec{k})^2/(\varepsilon_0 k^2)$ but is calculated to give a force from the mesh which is as close as possible, in a least squares sense, to the force expected from two $\gamma(\vec{r})$ shaped charge distributions. The optimal influence function, $\hat{G}(\vec{k})$, for a given system size, charge shape and assignment/interpolation function, needs to be constructed only once before the simulation is performed, thereafter the interactions are calculated by simply applying the above steps. Again, for the details of the method we refer the reader to Hockney and Eastwood [10].

3 RESULTS

In both methods, the short range interaction is significant only over a limited distance R_c which can be arbitrarily set by adjusting the value of α or a for the Ewald or PPPM methods, respectively. To evalulate the short range interactions for the system, one could consider each atom i and loop over all other atoms j selecting only those which lie within a distance R_c. This algorithm obviously grows as the number of atoms in the system squared (N^2) and is a very large computational burden to a dynamics simulation. To reduce computational demands, most molecular dynamics programs employ a pair list paradigm [13]. The pair list is first constructed using the algorithm above or perhaps a more sophisticated neighbor locating method. This list is subsequently used for several time steps n of the system, thus reducing the computational requirements per step by a factor of n.

As discussed in the Introduction, in a molecular system there are dispersion/repulsion interactions between atoms. For these interactions, a cutoff of 8 Å is typically sufficient. From a practical point of view, it would be convenient if a similar cutoff distance could be used for the short-range electrostatic interactions. Requiring a shorter distance for the electrostatics would have negligible computational savings since similar calculations must be performed for the dispersion/repulsion interaction. Allowing the distance to become significantly larger, requires the generation of a much larger pair list which can require significant amounts of memory.

To compare relative computational efficiencies of the PPPM and Ewald methods, we calculated the electrostatic forces in three systems of different size. A single configuration, typical for homogeneous distribution of sodium and chloride ions in simple point charge (SPC) water, was selected for each system. The number of ions and water molecules and the box length for each system are summarized in Table 1. The systems correspond to a concentration of approximately 1.286 Molar in each of the ions and a Debye inverse screening length of 0.3682 Å^{-1}.

For both methods, *a priori* force error estimates are non-trivial. For the Ewald method, Kolfka and Perram [14] present some simple formulas based on approximate local charge neutrality. For a desired Root Mean Square (RMS) error and a given cutoff radius R_c, these formulas can be used to estimate a reasonable α and $l_{\max}$ (i.e. the integer cutoff radius for the number of required k-space vectors). Using a cutoff $R_c = 8$ Å, the predicted α and l_{max} for a desired accuracy of 10^{-4} and 10^{-3} kcal/mol/Å are given, for each system, in columns 5 and 6 of Table 1.

Hockney and Eastwood give elaborate methods for determining the origins of errors in the Particle-Mesh (PM) calculation: charge undersampling, too large harmonic content in the reference force, finite differencing errors, etc. These measurements help to determine the necessary modifications to improve accuracy, and give a global measure of the error in the PM calculation. The error in the Particle-Particle (PP) calculation can also be estimated using similar assumptions to those proposed by Kolfka and Perram [14]. In particular, since the charge distribution with the S2 function extends only over a finite region, the error in the Particle-Particle (PP) interaction can be made identically zero by requiring R_c to be greater than a. In practice, because increasing the value of a makes the distribution smoother and therefore increases the accuracy of the PM calculation, it may be beneficial to make a larger than R_c and tolerate some error in the PP calculation to reduce overall error.

Practically, there are many choices to be made in a PPPM calculation and it may be easier to perform some experiments numerically. We have experimented with two schemes. The first scheme is the combination suggested by Hockney and Eastwood [10]. The charge from each atom is allocated to the nearest 27 grid points using a "Triangle-Shaped Charge (TSC)" function. The potential is interpolated using the same TSC function and the electrostatic fields are found using a 4-point central finite difference stencil. The grid size h was approximately 1 Å (see Table 1). In the second scheme, the charge is allocated to the nearest 8 grid points using a simple tri-linear weighting. The potential is interpolated using the same weighting and a simple 2-point central finite difference stencil is used to estimate the fields. To compensate for the errors caused by using a coarse assignment/interpolation and potential differencing, a finer grid of about 0.5 Å spacing was used (see Table 1).

Table 1 The three molecular systems used in this study consisted of approximately homogeneous mixtures of water and sodium and chloride ions in a cubic box of the specified length. There were an equal number of positive and negative ions in each system (i.e. for system 1 there were 5 sodium and 5 chloride ions) such that the entire system had no net charge. Parameters for the Ewald method ($\alpha, l_{\max}$) were estimated using the formula proposed by Kolfka and Perram. The number of grid points (N_g) in each dimension and the grid spacing (h) for the two PPPM methods described in the text are also given.

Sys	*Number of water*	*Number of ion pairs*	*Box Length*	*Ewald(1)*		*Ewald(2)*		*PPPM(1)*		*PPPM(1)*	
			[Å]	α [Å^{-1}]	l_{max}	α [Å^{-1}]	l_{max}	N_g	h [Å]	N_g	h [Å]
1	206	5	18.6206	0.335	5	0.276	4	16	1.164	32	0.582
2	1648	40	37.2412	0.335	10	0.276	7	32	1.164	64	0·582
3	5562	135	55.8618	0.335	15	0.276	11	48	1.164	96	0.582

As with all methods which rely on FFTs, the PPPM method is restricted in grid sizes to preferably powers of 2, although most FFT's have been optimized to perform well also with radix 3 and 5. Although this is inconvenient, it is usually not too difficult to find an appropriate grid size and spacing for any given system of interest.

In Table 2, we present measures of the errors in the forces for system 3 (results for systems 1 and 2 were similar). The errors are calculated using a very conservative Ewald calculation ($\alpha = 0.633\,\text{Å}^{-1}$ and $l_{max} = 58$, estimated error in the force of 10^{-12} kcal/mol/Å) as a reference. In the first row of Table 2, the absolute error of the force in kcal/mol/Å is presented measured either as RMS error, $(1/N)[\Sigma(f_i^{ex} - f_i)^2]^{1/2}$, or in parentheses as maximum error, $\text{MAX}\{|f_i^{ex} - f_i|\}$, where f_i is the magnitude of the force on atom i and f_i^{ex} is the magnitude of the force from the "exact" reference Ewald calculation. In the second row, the corresponding measures of the percentage relative error, $(1/N)[\Sigma(f_i^{ex} - f_i)^2/(f_i^{ex})^2]^{1/2}$, MAX $\{|(f_i^{ex} - f_i)/f_i^{ex}|\}$, are given. Finally, in the third row, the directional errors, $(1/N)\Sigma\theta_i$, and $\text{MAX}\{|\theta_i|\}$, where θ_i is the angle between the calculated force and the "exact" reference force, in degrees, are given. From Table 2, it is clear that all methods give reasonable errors typically with a very small percentage relative error and a maximum error on the order of 1 kcal/mol/Å. Values from the least accurate method proposed, PPPM(2), which uses the coarse assignment/interpolation and 2-point finite differencing, are perhaps too large for practical simulations. Ewald calculation (1) has an RMS error very close to the prescribed 10^{-4} kcal/mol/Å used to estimate the parameters while the second Ewald calculation, Ewald(2), has noticeably smaller RMS than the prescribed 10^{-3} kcal/mol/Å. The PPPM(1) combination favored by Hockney and Eastwood, using a TSC assignment/interpolation and 4-point finite differencing, has errors intermediate to the two Ewald calculations.

In Table 3, the CPU timings for the four methods are presented. Timings were performed on a Sun Sparc 2 workstation and on an IBM RISC/6000. All calculations were done using double precision. Both the Ewald and PPPM methods require calculation of short range interactions which can be accomplished by constructing a pair list and then interpolating the interaction from a pre-constructed table for each pair. The CPU time for constructing the pair list, using the simple double loop algorithm described in the Introduction, and the CPU time for calculating the short range interaction for each pair from a table look up are presented in the third and fourth columns of Table 3, respectively. As expected, the simple pair list algorithm, scales as the number of atoms squared and if constructed at every step in a dynamics

Table 2 Errors in the force calculated for system 3 relative to reference values calculated using a very conservative Ewald calculation (estimated error comparable to numerical round off). The first row gives the absolute error of the magnitude of the force in kcals/mol/Å measured as RMS or as the maximum error (in parentheses). The second row gives the corresponding percentage relative error in the magnitude of the force and the final row gives the corresponding measures for the absolute directional error in degrees.

Error	*Ewald(1)*	*Ewald(2)*	*PPPM(1)*	*PPPM(2)*
Absolute	0.106e-3 (0.058)	0.654e-3 (0.351)	0.194e-3 (0.098)	1.819e-3 (1.995)
Relative [%]	0.439e-3 (0.367)	2.881e-3 (2.831)	0.853e-3 (1.280)	7.693e-3 (12.52)
Direction	0.611e-3 (0.007)	3.736e-3 (0.043)	1.346e-3 (0.027)	9.289e-3 (0.146)

simulation would dominate the computational effort for a large system. If, as in practice, the list is re-constructed only every 10–20 time steps, the computational cost can be significantly reduced. Moreover, with sophisticated algorithms [15], the cost of generating a pair list can be made to scale linearly with the number of atoms. Since both implementations which we are comparing require the construction of a pair list, a discussion of these algorithms is beyond the scope of this paper. Given a pair list, the CPU time spent calculating the short range interaction, as expected, scales linearly with the number of atoms (column 4, Table 3).

Listed in columns 5 and 6 are the CPU times required for performing the "k-space" calculations in the Ewald sums. It should be noted that the code used to calculate these interactions was derived from that presented by W. Smith [16] and takes full advantage of obvious simplifications such as exploiting the symmetry in k-space and using trigonometric identities to reduce the double summation over particles to a single summation. This particular Ewald code has also been used for production runs in our laboratory, and therefore a more thorough optimization, including the pre-calculation and storage of common factors has also been performed. However, even when carefully optimized, the Ewald code for a *fixed* value of the cutoff grows as the number of particles squared. This is due to the requirement that the k-space cutoff ($l_{\max}$) grows with the system size and that we require a summation over k-space vectors for each particle. It is well known that if we allow the cutoff radius to vary, the Ewald method can be made to scale as the number of particles to the 3/2 power [17]. As mentioned above, for practical reasons we prefer to fix the value of the cutoff radius.

Listed in columns 7 and 8 are the CPU times required for performing the Particle-Mesh calculations. On the Sun workstation, we used the standard NAG Mark 14 [18] subroutine C06FJF to perform the FFTs and on the IBM workstation we used the routines DRCFT3/DCRFT3 from the IBM ESSL library [19] to perform the FFTs. Because the assignment/interpolation scales linearly with the number of particles and the FFT scales only slightly worse as $N \log_2(N)$, the computational cost of the PM

Table 3 CPU times for calculating forces using the methods described in the text. Timings were performed using either a call to the standard Fortran dtime subroutine on the Sun Sparc 2 workstation or the IBM timef subroutine on the RISC/6000 workstations. The columns labeled Pair List and Short Range are the CPU times required to perform the pair list construction and calculate the short range interaction using a table look-up. Both the Ewald and the PPPM method require these computations. Ewald (1) and Ewald (2) are the CPU times required to calculate the "k-space" contribution to the electrostatic forces using the Ewald method with the parameters listed in Table 1. PPPM (1) and PPPM (2) are the CPU times for calculating the Particle-Mesh contribution to the electrostatic forces for the PPPM method described in the text. All times are in seconds.

	Sys	*Number of Atoms*	*Pair List*	*Short Range*	*Ewald(1)*	*Ewald(2)*	*PPPM(1)*	*PPPM(2)*
Sun	1	216	3.17	2.68	2.85	1.50	1.45	4.74
Sparc 2	2	5024	201.43	20.73	177.13	61.56	11.25	36.35
	3	16956	2915.62	77.38	2023.22	801.08	38.67	129.39
IBM	1	216	0.62	0.69	0.35	0.19	0.10	0.12
RISC/6000	2	5024	64.66	6.21	44.42	15.13	0.97	1.00
	3	16956	731.03	20.89	489.75	198.87	3.69	3.32

calculation grows approximately linearly with the number of particles. More important than the theoretical scaling is the result that a PPPM calculation can be a competitive alternative to the Ewald calculation even for relatively small systems.

It is also interesting to compare the PPPM method to a traditional calculation which uses a cutoff and ignores the long range component of the electrostatic interactions outside this distance. For System 2 described above, the PPPM(1) method requires two 32^3, preferably double precision, grids of values on top of the storage required for the 8 Å pair list and other variables. The pair list requires approximately 520,000 integers, which, if stored in a standard 4-byte representation, requires almost 4 fold more memory than the double precision grids. More interestingly, if we increase the cutoff to 10 Å, which is more typical value for the traditional method, the size of the pair list approximately doubles and the CPU time for the calculation of the Coulomb interaction increases by almost 17 seconds on the Sun workstation and by almost 5 seconds on the IBM workstation. This leads to the non-obvious result that the PPPM method requires less computer time and on the order of half of the memory to perform a calculation including all periodic images, than is required by a traditional calculation using a 10 Å cutoff radius.

4 CONCLUSIONS

Relative to the maximum possible time step in a dynamics simulation, the time scales necessary for a careful investigation of dynamic phenomena of molecular systems can be daunting. In order to reach these time scales, every effort needs to be made to make the simulation program as efficient as possible while retaining the required accuracy. It is apparent that even for a relatively small molecular system, the PPPM method developed by Hockney and Eastwood can be an attractive alternative to the Ewald method and even to the relatively inaccurate cutoff methods. The trade off for efficiency is the extra effort required to construct computer code to perform the "Q-minimization" procedure and obtain the optimal influence functions $\hat{G}(\vec{k})$. Construction of this code is tedious, but since the code is run only a single time prior to simulation, efficiency is less of a burden and the code can be written in a straightforward fashion. Writing code to assign charges and interpolate forces is not difficult and is well described by Hockney and Eastwood.

The wide utilization of FFT algorithms virtually guaranties that highly optimized and robust library routines will be available for any desired platform. Thus, with only minor modifications in the existing code, we can expect to reach very high efficiency on a variety of computers just by linking to existing mathematical libraries.

As discussed in the Introduction, periodic boundary conditions are applicable in simulations of a variety of molecular systems. By using methods such as the Ewald or PPPM, we are able to include all periodic images in the electrostatic force calculation. Although this may impose an artificial periodicity on the system, it is arguably a better approximation than the currently popular methods which use a cutoff distance [9]. To numerically investigate whether the enhanced periodicity will have significant effects, we are currently introducing the PPPM code into the GROMOS suite of programs [20]. With the high computational efficiency of the PPPM method, it will be possible to study systems of sufficient size, for reasonable lengths of time, to hopefully draw some conclusion regarding the extent of these effects.

Acknowledgement

B.A.L. is grateful to the International Human Frontier Science Program Organization (HFSPO) for providing a Post-Doctoral fellowship during which this work was completed.

References

[1] M. Berkowitz and J. A. McCammon, "Molecular dynamics with stochastic boundary conditions", *Chemical Physics Letters*, **90**, 215 (1982).

[2] C. L. Brooks III, A. Brünger and M. Karplus, "Active site dynamics in protein molecules: A stochastic boundary molecular dynamics approach." *Biopolymers*, **24**, 843 (1985).

[3] M. K. Gilson, K. A. Sharp and B. H. Honig, "Calculating the electrostatic potential of molecules in solution: Method and error assessment", *Journal of Computational Chemistry*, **9**, 327 (1987).

[4] R. J. Zauhar and R. S. Morgan, "The rigorous computation of the molecular electric potential", *Journal of Computational Chemistry*, **9**, 171 (1988).

[5] M. E. Davis and J. A. McCammon, "Solving the finite difference linearized Poisson-Boltzmann equation: A comparison of relaxation and conjugate gradient methods", *Journal of Computational Chemistry*, **10**, 386 (1989).

[6] A. H. Juffer, *On the Modelling of Solvent Mean Force Potentials*, Ph.D. thesis, University of Groningen (1993).

[7] I. G. Tironi, R. Sperb, P. E. Smith and W. F. van Gunsteren, "A generalized reaction field method for molecular dynamics simulations", *Journal of Chemical Physics* (1994), submitted for publication.

[8] H. Schreiber and O. Steinhauser, "Cut-off size does strongly influence molecular dynamics results on solvated polypeptides.", *Biochemistry*, **31**, 5856 (1992).

[9] P. E. Smith and B. M. Pettitt, "Peptides in ionic solutions: A comparison of Ewald and switching function techniques". *Journal of Chemical Physics*, **95**, 8430 (1991).

[10] R. W. Hockney and J. W. Eastwood, *Computer Simulation using Particles*. IOP Publishing Ltd, Bristol, England (1988).

[11] H. J. C. Berendsen, "Electrostatic interactions", in W. F. van Gunsteren, P. K. Weiner and A. J. Wilkinson (eds.), *Computer Simulation of Biomolecular Systems, volume* 2, pp. 161–181. ESCOM Science Publishers B. V., Leiden, The Netherlands (1993).

[12] D. M. Heyes, "Electrostatic potentials and fields in infinite point charge lattices" *Journal of Chemical Physics*, **74**, 1924 (1981).

[13] M. P. Allen and D. J. Tildesley, *Computer Simulation of Liquids*. Oxford University Press, New York (1987).

[14] J. Kolfka and J. W. Perram, "Cutoff error in the Ewald summation formulae for point charge systems," *Molecular Simulations*, **9**, 351 (1992).

[15] Z. A. Rycerz and P. W. M. Jacobs, "On the efficiency of vectorized molecular dynamics algorithms of order n", *Molecular Simulations*, **8**, 249 (1992).

[16] W. Smith, "Fortran code for the Ewald summation method", *CCP5 Information Quarterly*, pp. 37–53 (1986).

[17] J. W. Perram, H. G. Petersen and S. W. de Leeuw, "An algorithm for the simulation of condensed matter which grows as the 3/2 power of the number of particles." *Molecular Physics*, **65**, 875 (1988).

[18] *NAG Fortran Library Manual, Mark 14* (1990).

[19] *Engineering and Scientific Subroutine Library Version 2, Guide and Reference* (1994).

[20] W. F. van Gunsteren and H. J. C. Berendsen, *Groningen Molecular Simulation* (*GROMOS*). Biomos, Groningen, The Netherlands (1987).