Direct-Space Corrections Enable Fast Accurate Lorentz-Berthelot Combination Rule Lennard-Jones Lattice Summation

Christian L. Wennberg, Teemu Murtola, Szilárd Páll, Mark J. Abraham, Berk Hess, and Erik Lindahl*

Swedish e-Science Research Center & Dept. Theoretical Physics, KTH Royal Institute of Technology,

Box 1031, 17121 Solna, Sweden &

Center for Biomembrane Research, Dept. Biophysics & Biochemistry,

Stockholm University, 106 91 Stockholm, Sweden

E-mail: erik.lindahl@scilifelab.se

Abstract

Long-range lattice summation techniques such as the Particle-Mesh Ewald(PME) algorithm for electrostatics have been a revolution to the precision and accuracy of molecular simulations in general. Despite, a performance penalty, few biomolecular simulations today are performed without lattice summation electrostatics. There are increasingly strong arguments for moving in the same direction for Lennard-Jones (LJ) interactions, and we recently made a new fast LJ-PME implementation available in the Gromacs package, where we relied on approximations of the combination rules in reciprocal space to reach high performance. Here, we propose a new way to correct for these approximations that achieve exact treatment of Lorentz-Berthelot combination rules within the cutoff, and only a very small approximation remains outside the cutoff in the reciprocal space component. Not only does this improves accuracy by almost an order of magnitude, but it also achieves absolute biomolecular simulation performance that is an order of magnitude faster than alternatives. The implementation in Gromacs

*To whom correspondence should be addressed

includes both CPU and GPU acceleration, and combined with improved scaling LJ-PME simulations now provide performance close to truncated potentials, and much higher accuracy.

1 Introduction

The Particle-Mesh Ewald (PME) algorithm developed by Darden and coworkers[1,2] has been the method of choice for long-range electrostatic interactions in Molecular Dynamics (MD) simulations of biological systems the last decade. Previously, these interactions were treated with cutoffs after a fixed distance, which has been shown to yield severe artifacts in calculated properties.[3–8] While these artifacts were avoided with PME, the initial implementations were computationally expensive (typically 2-3x times slower), and it was first with the improved smooth PME[2] and later accelerated implementations by others and us[9] the performance improved to the point where PME electrostatics largely replaced truncated or reaction-field Coulomb interactions. In fact, already in the early papers Darden *et al.* noted how the PME method could be extended to include any interaction for which the potential decayed as $r^{-\alpha}$ (assuming $\alpha \geq 1$), and one obvious such extension of the method would be to apply it to calculate the dispersive component of Lennard-Jones interactions. However, the implementation of Lennard-Jones PME (LJ-PME) turned out to be very challenging since the combination rules applied in order to derive the interaction coefficients for mixed particle types made it several times slower. The penalty was considerably larger than for electrostatics PME, and since the trunation errors are both smaller and it is possible to apply analytical dispersion corrections, LJ-PME never caught on. Given the increasing accuracy requirements for sensitive, heterogeneous and anisotropic systems such as membranes, we argue it is time to change this practice.

Currently, most simulations of biologically relevant systems containing lipids (e.g. membranes) utilize cutoffs in the range of 10-14 Å for the Lennard Jones (LJ) interactions, and correct for the missing energy by an analytical dispersion correction

$$U_{\mathrm{dc}} = \frac{N}{2}\rho \int_{r_c}^{\infty} 4\epsilon \left[\left(\frac{\sigma}{r}\right)^{12} - \left(\frac{\sigma}{r}\right)^{6} \right] g(r) 4\pi r^2 dr. \tag{1}$$

Here, N is the number of particles in the system, ρ the average number density, and σ, ϵ the interaction coefficients that are approximated as average parameters. $g(r)$ is the pair distribution function that is assumed to be equal to unity at a sufficiently long cutoff, which enables us to evaluate the integral as

$$U_{\mathrm{dc}} = 8N\pi\rho\epsilon\sigma^3 \left[\frac{1}{9}\left(\frac{\sigma}{r_c}\right)^{9} - \frac{1}{3}\left(\frac{\sigma}{r_c}\right)^{3} \right]. \tag{2}$$

However, the use of this correction term in simulations of different types of interfaces (e.g. lipid bilayers) is troublesome since the assumption that $g(r)$ is equal to unity beyond the cutoff inherently assumes that the system is homogeneous and isotropic, something that most definitely is not true for these systems. The implication of this for simulations containing lipid bilayers has not been given much attention in the past, but coupled to the

development of alternative ways of treating long-range LJ interactions[10–17] the problem has started to become more apparent the last few years. In light of this, we recently presented an implementation in the Gromacs simulation package where the PME method was used to calculate long-range LJ interactions.[18] In this implementation we proposed to circumvent the expensive calculations for the reciprocal sum by approximating the reciprocal-space interaction coefficients with the geometric combination rules[19] even where the force field uses Lorentz-Berthelot rules[20] (LB). While this introduced slight errors in the entire reciprocal space component (which extends inside the cutoff too), we argued this was much smaller than the errors historically accepted with analytical dispersion correction.

Here, we present a new idea to correct for these approximations, where we modify the evaulation of the direct-space potential in order to properly account for the approximated combination rules in reciprocal space. This direct space modification completely cancels the approximation errors in the reciprocal space combination rules inside the cutoff, which means the resulting total interaction is *i* exact even for Lorentz-Berthelot combination rules up to the cutoff (both direct and recpirocal space components), and there is only a very small approximation remaining outside the cutoff. Further, we have improved the performance and scaling

of the previous implementation so that it now is comparable to simulations using cutoffs with dispersion correction to handle the long-range LJ-interactions.

2 Theory

2.1 Lattice summation of dispersion interactions

Given N particles at positions $\mathbf{r}_1, \mathbf{r}_2, ..., \mathbf{r}_n$ within the unit cell U, we can describe a general $r^{-\alpha}$-interaction between them as

$$U_\alpha = \frac{1}{2} \sum_{\mathbf{n}}{}' \sum_j \sum_k \frac{C(j,k)}{|\mathbf{r}_{jk} + \mathbf{n}|^\alpha} \tag{3}$$

where the outer sum is over the lattice translation vectors $\mathbf{n}$, $C(j,k)$ is the combined interaction coefficient, $\mathbf{r}_{jk} = \mathbf{r}_j - \mathbf{r}_k$ is the interparticle distance, and $'$ indicates that terms with $j = k$ are omitted when $\mathbf{n} = 0$.

For interactions with $\alpha \leq 3$ (such as electrostatics) this sum does not even converge if the system has a net charge. For neutral systems it can be made conditionally convergent (i.e., if the terms are added in the right order) but even then the convergence is very slow. For this reason a lot of work has been invested in further developing the general solution, derived by Ewald,[21] in order to also be applicable to molecular simulations.[1,2,15,17,22–24]

Following the derivation of Essman,[2] the expression above, applied to dispersion inter-

actions, can be rewritten as

$$\begin{aligned} U_{\text{dispersion}} &= \frac{1}{2}\sum_{\mathbf{n}}{}'\sum_{j}\sum_{k}\frac{C(j,k)g(\beta|\mathbf{r}_{jk}+\mathbf{n}|)}{|\mathbf{r}_{jk}+\mathbf{n}|^6} \\ &+ \frac{\pi^{3/2}\beta^3}{2\mathbf{V}}\sum_{\mathbf{m}} f(\pi|\mathbf{m}|/\beta)\times\sum_{j}\sum_{k}C(j,k)\exp\left[-2\pi i\mathbf{m}\cdot\mathbf{r}_{jk}\right] \\ &- \frac{\beta^6}{12}\sum_{j}C(j,j) \end{aligned} \tag{4}$$

with $\mathbf{m}$ being the reciprocal lattice vectors and β the ewald splitting parameter. The first and second term are the direct and reciprocal interaction potentials, and the third term constitutes the self energy. The expressions for $f(\pi|\mathbf{m}|/\beta)$ and $g(\beta|\mathbf{r}_{jk}+\mathbf{n}|)$ are given by

$$f(x) = 1/3\left[(1-2x^2)\exp(-x^2)+2x^3\sqrt{\pi}\,\text{erfc}(x)\right] \tag{5}$$

$$g(x) = \exp(-x^2)(1+x^2+\frac{x^4}{2}). \tag{6}$$

Assuming the interaction coefficients $C(j,k)$ can be factorized, $C(j,k)=\sqrt{C(j,j)}\cdot\sqrt{C(k,k)}$, we can remove the double sum in the second term of 4 by defining the so-called structure factors, $S(\mathbf{m})$, as

$$S(\mathbf{m}) = \sum_{j}^{N} C(j,j)\,\exp\left[2\pi i\mathbf{m}\cdot\mathbf{r}_j\right], \tag{7}$$

which translates this second term into

$$\frac{\pi^{3/2}\beta^3}{2\mathbf{V}}\sum_{\mathbf{m}} f(\pi|\mathbf{m}|/\beta)\times S(\mathbf{m})S(-\mathbf{m}) \tag{8}$$

on which we can apply the PME methodology in the same way as it is done for electrostatic interactions. This makes it relatively painless to use LJ-PME together with force fields that utilize geometric combination rules ($C(j,k)=\sqrt{\sigma_j\sigma_k}$). But when the force field combines the coefficients according to the LB rules, the coefficient $C(j,k)$ instead splits into seven terms

$$C(j,k) = (\sigma_j+\sigma_k)^6 = \sum_{n=0}^{6} P_n\sigma_j^n\sigma_k^{(6-n)}, \tag{9}$$

which thereby also forces the splitting of the reciprocal dispersion energy into seven terms

$$\begin{aligned} S(\mathbf{m})S(-\mathbf{m}) &= \sum_{j,k}\sum_{n=0}^{6} P_n \sigma_j^n \sigma_k^{(6-n)} \exp(2\pi i \mathbf{m}\cdot\mathbf{r}_{jk}) \\ &= \sum_{n=0}^{6} P_n \left[\sum_j \sigma_j^n \exp(2\pi i \mathbf{m}\cdot\mathbf{r}_j)\right]\left[\sum_k \sigma_k^{(6-n)} \exp(-2\pi i \mathbf{m}\cdot\mathbf{r}_k)\right] \\ &= \sum_{n=0}^{6} P_n Z_n(\mathbf{m}) Z_{6-n}(-\mathbf{m}), \end{aligned} \tag{10}$$

where $Z_\alpha(\mathbf{m}) = \sum_j \sigma_j^\alpha \exp(2\pi i \mathbf{m}\cdot\mathbf{r}_j)$, and P_n are the Pascal triangle coefficients.

This effectively makes the PME-method impractical to use together with force fields utilizing LB combination rules, since the reciprocal interactions require the evaluation of several separate fast Fourier transforms (FFTs). As we have shown previously,[18] it is possible to circumvent this problem by approximating the interaction coefficients in reciprocal space using geometric combination rules, and this will not cause any serious effects in normal production runs. Below, we present a way to modify the direct-space interactions in order to reduce the introduced error to a level where it is completely negligible.

2.2 Correcting the reciprocal energy in direct space

From 3 we can see that the dispersion energy in our system is calculated as

$$U_{\text{dispersion}} = \frac{1}{2}{\sum_{\mathbf{n}}}' \sum_j \sum_k \frac{C_6^{jk}}{|\mathbf{r}_{jk}+\mathbf{n}|^6} \tag{11}$$

with C_6^{jk} being the combined dispersion coefficient. Using Ewald summation we split this slowly converging sum into two sums

$$U_{\text{dispersion}} = \underbrace{\frac{1}{2}\sum_{\mathbf{n},j,k} \frac{C_6^{jk} g(\beta|\mathbf{r}_{jk}+\mathbf{n}|)}{|\mathbf{r}_{jk}+\mathbf{n}|^6}}_{\text{Sum in direct space}} + \underbrace{\frac{1}{2}\sum_{\mathbf{n},j,k} \frac{C_6^{jk,\text{recip}}[1-g(\beta|\mathbf{r}_{jk}+\mathbf{n}|)]}{|\mathbf{r}_{jk}+\mathbf{n}|^6}}_{\text{Sum in reciprocal space}}, \tag{12}$$

where the function $g(\beta|\mathbf{r}_{jk}+\mathbf{n}|)$ is given by 6, $C_6^{jk,\text{recip}}$ is the dispersion coefficient used for the reciprocal sum (i.e., not always equal to C_6^{jk}), and we have defined $\sum_{\mathbf{n},j,k} = {\sum_{\mathbf{n}}}' \sum_j^N \sum_k^N$. The first sum in 12 is a fast converging sum in direct space, while the second sum is a slowly varying function of $|\mathbf{r}_{jk}+\mathbf{n}|$, thus making its Fourier transform a fast converging sum in reciprocal space.

For our old implementation, and the case where $C_6^{jk,\text{recip}} \neq C_6^{jk}$, this gave the total interaction energy within the cutoff as

$$U_{\text{dispersion}} = \frac{1}{2}\sum_{\mathbf{n},j,k} \frac{C_6^{jk,\text{recip}}}{|\mathbf{r}_{jk}+\mathbf{n}|^6} + \frac{1}{2}\sum_{\mathbf{n},j,k} \frac{(C_6^{jk}-C_6^{jk,\text{recip}})[1-g(\beta|\mathbf{r}_{jk}+\mathbf{n}|)]}{|\mathbf{r}_{jk}+\mathbf{n}|^6}, \tag{13}$$

which maintains a well-defined Hamiltonian while keeping the associated error at a negligible level (below 0.5% of total dispersion energy) in the production runs we performed. A more correct treatment of this potential would instead be to modify the direct-space calculations so that we subtract the reciprocal contribution from the original dispersion potential in 11

$$\begin{aligned} U_{\text{dispersion}} &= \underbrace{\frac{1}{2}\sum_{\mathbf{n},j,k}\frac{C_6^{jk}}{|\mathbf{r}_{jk}+\mathbf{n}|^6} - \frac{C_6^{jk,\text{recip}}[1-g(\beta|\mathbf{r}_{jk}+\mathbf{n}|)]}{|\mathbf{r}_{jk}+\mathbf{n}|^6}}_{\text{Sum in direct space}} \\ &+ \underbrace{\frac{1}{2}\sum_{\mathbf{n},j,k}\frac{C_6^{jk,\text{recip}}[1-g(\beta|\mathbf{r}_{jk}+\mathbf{n}|)]}{|\mathbf{r}_{jk}+\mathbf{n}|^6}}_{\text{Sum in reciprocal space}} \\ &= \frac{1}{2}\sum_{\mathbf{n},j,k}\frac{C_6^{jk}}{|\mathbf{r}_{jk}+\mathbf{n}|^6}. \end{aligned} \tag{14}$$

This direct-space modification leaves us with an unmodified potential up to the cutoff. At this point the direct-space terms are set to zero, and hence the error we obtain only comes from the reciprocal part outside of the cutoff. In order to account for the abrupt shift in the potential that happens at the cutoff, we also add the constant

$$\frac{-C_6^{jk} + C_6^{jk,\text{recip}}[1-g(\beta r_{\text{cutoff}})]}{r_{\text{cutoff}}^6} \tag{15}$$

to ensure that we maintain a continuous potential. This enables us to use geometric combination rules for the reciprocal calculations regardless of the combination rules in the applied force field, and as we will show in the following sections this reduces the approximation error by an order of magnitude compared to the old implementation.

3 Methods

For the simulations of the palmitoyl-oleoyl-phosphocholine (POPC) bilayers in this work we extended the simulations using LJ-PME with Lorentz-Berthelot combination rules in our previous work.[18] The structures simulated originally come from the work of Hub *et al.*[25] These structures contain 128 POPC molecules and 5,262 SPC[26] water molecules. The systems were simulated for 100 ns at 298 K, using a time step of 2 fs. Bond lengths were constrained with the LINCS[27] algorithm. Temperature was coupled with the Nosé-Hoover thermostat, and pressure was coupled semi-isotropically at 1 bar with the Parrinello-Rahman barostat. The POPC parameters were modified from the parameters of Berger[28] by using parameters for the double bond and atom types from the OPLS united atom force field.[29] Electrostatic interactions were evaluated every step using PME,[2] and the LJ-interactions were evaluated using three different setups of LJ-PME: (I) full treatment of the non-bonded interactions using LB combination rules, (II) LB combination rules in direct space together with geometric combination rules in reciprocal space, and (III) LB combination rules with

modified direct space interactions (as described in section 2.2) together with geometric combination rules in reciprocal space. A fourth-order spline interpolation was applied for the PME calculations, together with a grid spacing of 1.2 Å.

For the performance benchmark simulations we used three different systems: a box of SPC/E water,[30] a dipalmitoyl-phosphocholine (DPPC) bilayer, and a well-studied ion channel system containing the membrane protein GluCl.[31] The water system was constructed by placing 666 molecules in a cubic box and then condensing the system for 50000 step using a reference pressure of 10 bar and a temperature of 298 K. After this the reference pressure was set to 1 bar, and the system was equilibrated for an additional 100000 steps reaching a final density of approximately 1 kg/L. For the production runs, this system was then replicated in three dimensions to reach the desired system size. The strong scaling simulations were carried out with a system size of 42624 molecules, while the weak scaling was measured using systems with 42624, 340992 and 2727936 molecules maintaining an average number of particles per core at approximately 250. The system was simulated using geometric combination rules in LJ-PME, and for the strong scaling a reference simulation with an LJ-cutoff of 10 Å and dispersion correction was also performed. The simulations were performed using the new Verlet cutoff scheme (later referenced as Verlet scheme) in Gromacs.[32]

The DPPC bilayer, containing 400 DPPC molecules and 12,820 SPC/E water molecules, was constructed replicating our previous DPPC-system[18] first along the x- and then along the y-axis making it four times larger. The system was first equilibrated using a conjugate gradient method with the tolerance 1 kJ mol^{-1} nm^{-1}, followed by an MD simulation using a timestep of 0.2 fs for 30,000 steps after which the time step was increased to 2 fs and an additional 30,000 steps were performed. Electrostatic interactions were evaluated using PME, LJ-interactions by using a twin-range 10/16 Å cutoff updated every 10th step. DPPC parameters were taken from Chiu *et al.*[33] The performance benchmark simulations were carried out using a direct-space cutoff of 10 Å for both electrostatic and Lennard-Jones PME, and since the DPPC-coefficients are derived using geometric combination rules we treated both direct and reciprocal space in LJ-PME using these rules. Reference benchmarks were obtained from three simulations using different choices for the LJ-setup: 10 Å cutoff, twin-range 10/14 Å or twin-range 10/16 Å both with an update frequency of 10 steps, and dispersion correction applied for all three setups. Other than the choice of treatment for the LJ-interactions (and temperature) all of the DPPC-systems were simulated using the same parameters as the POPC system above. The DPPC simulations were perfomed using the charge group based cutoff scheme in Gromacs (later referenced as group scheme).

The ion channel system, consisting of GluCl embedded in a solvated bilayer (a total of 142 000 atoms), represents a typical membrane protein, which proved to be quite challenging to scale due to its relatively small size and the inhomogeneous computational workload. The benchmarks were set up with the Gromacs verlet cutoff scheme, a time step of 2 fs and h-bond constraints. To compare the performance of the GPU accelerated LJ-PME implementation to the performance when treating Lennard-Jones interactions with cutoffs, we carried out benchmarks with a 10 Å cutoff and geometric combination rules in reciprocal space for LJ-PME as well as runs with LJ interactions truncated at a slightly longer distance, 14 Å. The performance was measured over the second half of three minute long runs ensuring that the influence of initialization and initial load-balancing on performance is minimized. Extreme strong scaling is shown to 1100 atoms per node or 138 atoms per CPU core.

All simulations were carried out with a modified development version of the Gromacs simulation package (5.0). The implementation of the modified direct space interactions (section 2.2) is part of the publicly available Gromacs-5.0 release, while the performance modification in section 4.3 is in the last stages of development and will be included in a future patch level.

4 Results

4.1 Modifiying the direct-space interactions increases accuracy

The implementation of the modified direct-space interactions (section 2.2) in Gromacs was tested on the same POPC bilayer as we used in previous work.[18] The bilayer simulation using LB combination rules was extended for an additional 100 ns using three setups: (I) Full LB combination rules in both direct and reciprocal space, (II) approximation of reciprocal space with geometric combination rules, and (III) modified direct-space interactions together with geometric combination rules in reciprocal space. Since the force field parameters for these lipids originates from the Berger lipids,[28] and hence uses geometric combination rules, the errors introduced by the application of LB combination rules will have a small effect on the calculated properties. But considering that we here are investigating the relative difference between the simulation setups, which all use the same combination rules, we deem that this will not have a significant effect on our results.

From these simulations we calculated the average area per lipid during the last 20 ns of simulation time, and the results can be seen in 1. In agreement with our previous results[18] the values obtained from simulation II are very similar to the ones obtained from simulation I, displaying the small effect from the reciprocal approximation on simulated properties. However, the addition of the direct-space modification present in simulation III reduces this difference even further, making their average value identical at 66.9 Å^2 and within reported standard errors of each other. Furthermore, the high correlation between area per lipid at equilibrium and lipid chain ordering makes it possible to also study this effect through the deuterium order parameters. In 1 we show the order parameters S_{cd} for both the sn1- and sn2-chain from simulations I-III, calculated as an average over the last 20 ns of simulation using

$$S_{\mathrm{cd}} = \frac{2}{3}S_{xx} + \frac{1}{3}S_{yy} \tag{16}$$

with S_{xx} and S_{yy} defined according to

$$S_{ij} = \frac{1}{2}\langle 3\cos\theta_i \cos\theta_j - \delta_{ij}\rangle, \tag{17}$$

where θ_i represents the angle between the ith molecular axis and the bilayer normal.[34]

Here, the results from 1 are reinforced by the calculated order parameters for the sn1-chain of POPC. A discrepancy can be seen between the values calculated from simulation II and those from simulations I and III, with II having slightly higher order parameters for all atoms. For the sn2-chain, the existence of the double bond reduces the differences due to the conformational restrictions imposed on the atoms, but there is still a clearly consistent

trend of simulation III being closer than II to the values obtained in simulation I.

In addition to the structural analysis we also performed a recalculation of all the energies and forces at every frame in the trajectory from simulation I, using the other two methodologies, and the results for this are shown in 2. The difference in total energy throughout the simulation when only using the approximation in reciprocal space (simulation II) gives an error of approximately 470-480 kJ/mol, or 0.1% of the total energy. When using the modified potential in direct space the difference is reduced by an order of magnitude, to approximately 50 kJ/mol, thereby almost eliminating the statistical difference compared to the original simulation which has an error estimate of 29 kJ/mol. The remaining difference might very well be due to the summation differences with single precision, rather than the algorithm. This difference also holds true for the RMSD of the total force in the simulations. Here the setup using parameters from simulation II obtains a RMSD of 0.168 kJ mol^{-1} nm^{-1} while setup III displayed a RMSD reduced to 0.0142 kJ mol^{-1} nm^{-1}.

4.2 Accuracy

In order to investigate the accuracy of the LJ-PME implementation we performed a series of single frame calculations for different truncations of the direct-space potentials (using the DPPC bilayer from our previous work[18]). The reciprocal grid spacing was also modified with the same proportional change as the direct-space cutoff starting at 1.2 Å for a cutoff of 10 Å. These were compared to a reference setup in which we used a direct-space cutoff of 10 Å and a reciprocal grid with a spacing of 0.2 Å together with a sixth-order spline interpolation. The LJ-PME calculations were performed using geometrical combination rules and we also performed the same calculations using a 10 Å cutoff together with dispersion correction for long-range LJ-interactions. The calculations were performed using both the Verlet and group cutoff schemes available in Gromacs, and the results for these calculations can be seen in 3.

For the cutoff calculations the most accurate results are non-surprisingly found when using a cutoff of 10 Å, which is the value specified for the applied force field. Increasing the cutoff has a small effect on the force-RMSD, and a quite substantial effect on the relative error in the LJ-energy. For LJ-PME the RMSD of the force seems to be proportional to the cutoff (an increase of the cutoff beyond 13 Å further decreased the RMSD), while the most accurate result for the LJ-energy is obtained at 9 Å. A further increase of the cutoff beyond 10 Å seem to have no further effect on the LJ-energy, with the error staying constant at approximately 9%.

A recent suggestion by Isele-Holder *et al.*,[17] regarding the applicable cutoffs when using long-range dispersion solvers, noted that cutoffs shorter than twice the largest LJ-diameter would impact the system properties since effects from the truncation of the repulsion term becomes non-negligible. For the DPPC-system studied here, the LJ-diameter of the methylene groups in the united atom tails is approximately 4.1 Å, which suggests that results obtained with cutoffs lower than ˜8.2 Å would be innacurate to use. This is confirmed by the results from our LJ-PME calculations, which have a steep increase of the error in LJ-energy going from a direct space cutoff of 9 Å down to 8 Å.

From these calculations we can see that the accuracy of LJ-PME and cutoffs seems to be comparable when we apply the same direct space cutoff as specified for the utilized force field. Diverting from this cutoff quickly reduces the accuracy of the energies obtained from

the cutoff calculations, while the errors from the LJ-PME calculations stay reasonably close to the reference. The most interesting result here would be that regardless of the applied cutoff LJ-PME seems to maintain an accuracy that is much closer to the standard 10 Å cutoff calculation with dispersion correction than those obtained in the corresponding cutoff-calculation. Furthermore, the relatively small differences in the obtained force-RMSDs when using LJ-PME indicates that LJ-PME might be usable with a much smaller direct-space cutoff than what the applied force field dictates. It also gives a good indication that future force fields, developed with dispersion solvers for all the non-bonded interactions, hopefully can remove the cutoff restriction commonly imposed by force fields utilized today.

4.3 Performance

In our previous work[18] we showed that when using LJ-PME with LB combination rules we only obtain 25% of the performance we get by using cutoffs, which makes it hard to motivate its use. However, we also noted that by using geometric combination rules for the reciprocal part of the potential, the performance could be increased up to almost 80% of the performance obtained using a twin-range 10/16 Å cutoff. Thus, here we will only show numbers concerning geometric rules in reciprocal space since the newly implemented direct-space modification (section 2.2) should make it unnecessary to use LB combination rules in standard production runs.

The previously reported implementation of LJ-PME had challenges with scalability when the number of particles per core dropped below 500, which was coupled to the extremely prohibitive communication costs when performing all-to-all communication during the calculation of the three-dimensional FFTs. At the scaling limit the number of messages, not their size becomes the bottleneck, We have now implemented a combined all-to-all communication for the electrostatic and Lennard-Jones PME in order to effectively reduce this cost by a factor of two. Together with new analytical interaction kernels that also use SIMD (single-instruction, multiple-data) instructions (e.g. AVX on x86) for LJ-PME this now enables LJ-PME to scale as well as the cutoff simulations down to about 30 particles per core. The performance was measured on three separate systems: a box of SPC/E-water simulated at the Blue Gene/Q system at the Jülich Supercomputing Centre, a DPPC bilayer simulated on the Triolith supercomputer (Intel E5-2660 CPUs) at the National Supercomputer Centre in Linköping, and the ion channel system GluCl[31] simulated on the Piz Daint supercomputer (a hybrid Cray XC30 machine with Intel Xeon E5-2670 CPUs, NVIDIA Tesla K20X GPUs and an Aries Dragonfly network) at the Swiss National Supercomputer Centre in Zürich.

On the Blue Gene/Q system, we investigated the weak and strong scaling limits of the LJ-PME implementation, and the results are displayed in 2. In the strong scaling simulations, LJ-PME maintains a relative performance of 70% compared to the one obtained using a single range 10 Å cutoff, a large improvement compared to the old implementation due to the fact that we here go all the way down to 32 atoms per core in the last simulation. We now observe similar performance characteristics for the cutoff and LJ-PME performance simulations, and the largest difference in absolute performance are coupled to the interpolation of the interaction coefficients onto the PME-grid.

In regards to weak, scaling LJ-PME performs slightly worse than linear going from 512 to 4096 cores, and the small performance loss is coupled to increased communication costs in

the domain decomposition, neighbour search, and coordinate communication. Increasing the number of cores by an additional factor of 8 incurs a slightly larger performance degradation, which is coupled to an increase in the cost of the redistribution of particle positions and forces within the PME-ranks, along with the interpolation of the interaction coefficients onto the PME-grid. Since these regions were found to be the limiting factor in the measurments of both strong and weak scaling they are now under consideration for future improvements to the algorithm.

From the DPPC simulations (4) we see that when simulating a more heterogeneous system the current implementation has an absolute performance which is slightly faster than a cutoff simulation using a twin-range 10/16 Å cutoff, and which stays at around 90-100% of what is achieved with the 10/14 Å cutoff (16 cores being an outlier with a relative performance of approximately 80%). Compared to the simulations with a cutoff of 10 Å LJ-PME still performs relatively well with a relative performace of approximately 75-80% (85% at 512 cores as the highest and 72% at 16 cores the lowest), which can be compared to 10/14 Å which stays at a relative performance of 80-90% (compared to 10 Å) and 10/16 Å at 70-80%.

In order to assess the performance obtained with LJ-PME on more complex systems commonly simulated today we also performed benchmarks on the well-studied ion channel GluCl, and the results is presented in 5. On heterogeneous hardware, such as the setup on Piz Daint, Gromacs offloads non-bonded interactions to GPUs and therefore the performance is very sensitive to the CPU-GPU balance. The Cray XC30 hybrid machine represents a well-balanced setup were the GPUs are fast enough to reduce the cost of computing short-range non-bonded interactions to the point where it is possible to use cutoffs that are longer than default, up to 12-14 Å, without any additional loss of performance. Because of this the overall performance, as shown in 5, is bound by the PME performance from around 4 nodes and above. This have the effect of making the difference between the simulations using cutoffs almost negligible, and hence the comparison with the LJ-PME implementation even more challenging. Still, the GPU accelerated LJ-PME implementation manages to achieve, on a single node, 70% and 92% of the performance of the simulations utilizing 10 and 14 Å cutoff. It also exhibits equally well strong scaling as the cutoff simulations, achieving 81% and 96% of the 10 and 14 Å cut-off simulations on 128 XC30 nodes. It is worth noting that the LJ-PME non-bonded force kernels, with geometric combination rules, are only about 20% slower than the LJ cutoff kernels and therefore on other hybrid machines with relatively less GPU computational power LJ-PME will have even more advantage than a longer cutoff.

Finally, we would also like to note that these comparisons are against the performance of normal simulation setups in Gromacs, which is much faster than other reference codes available and more or less state-of-the-art for membrane simulations today.

5 Conclusions

This work introduces a large improvement of the direct-space interactions in the implementation of the LJ-PME method in Gromacs. We demonstrate that for interaction distances within the cutoff our implemented direct-space modification completely removes the error introduced by the previously suggested approximation in reciprocal space. This preserves a correct description of the potential up to the cutoff, the obtained error is contained to

the reciprocal part outside of the cutoff, and it retains the huge computational performance advantage of avoiding the Lorentz-Berthelot combination rules in reciprocal space. Furthermore, we have made optimizations to the communication bottlenecks, and together with new accelerated non-bonded kernels for the direct-space interactions in LJ-PME we now attain a performance with LJ-PME which is on par with what is achieved by a standard simulation setup in Gromacs. As the last few optimizations too become integrated in Gromacs we should finally be able to lay aside the issue with performance degradation from mesh-based dispersion solvers. With the first implementations of electrostatic PME we were looking at a performance degration far worse than this, and with a possible implementation of a multiple time step algorithm in a future version of Gromacs, there is even potential to improve performance above the previous cutoff case in Gromacs.

LJ-PME now enables the correct treatment of long-range LJ-interactions in simulations of inhomogeneous systems, which until recently have been treated with artificial cutoffs and insufficient long-range corrections that assume isotropic surroundings. For electrostatic interactions, the severe effects coupled to the use of cutoffs[7] were quickly understood by the community, but for dispersion interactions the issue seems to be more subtle. The inclusion of long-range dispersion effects has been shown to have non-negligible effects on structural properties such as area per lipid and surface tension,[14,18,35] and we have previously shown that LJ-PME can be used with existing force fields. However, the force fields we use today have been parametrized using inadequate long-range methods for dispersion interactions, so a reparametrization with correct treatment of the long-range interactions is of interest. The use of mesh-based solvers for all long-range interactions has been shown to transform the cutoff into a possible tuning parameter for performance.[17] Therefore, a force field parametrized with a correct treatment of all long-range interactions might remove the cutoff-restriction present in existing force fields, and thus we would obtain an accurate force field where the cutoff is used as an extra tuning parameter to achieve the highest performance.

Acknowledgement

This work was supported by grants from the Swedish Foundation for Strategic Research, the Swedish Research Council (2010-491, 2010-5107), the CRESTA EU FP7 project (287703) and the Swedish e-Science Research Center. Computer resources on the Triolith supercomputer were provided through the Swedish National Infrastructure for Computing (SNIC 020/11-41). We would also like to acknowledge the Jülich Supercomputing Center (JSC) and the Swiss National Supercomputing Centre (CSCS) for the BlueGene/Q and Piz Daint computing resources.

References

(1) Darden, T.; York, D.; Pedersen, L. *J. Chem. Phys.* **1993**, *98*, 10089–10092.

(2) Essmann, U.; Perera, L.; Berkowitz, M. L.; Darden, T.; Lee, H.; Pedersen, L. G. *J. Chem. Phys.* **1995**, *103*, 8577–8592.

(3) Hess, B. *J. Chem. Phys.* **2002**, *116*, 209–217.

(4) Hünenberger, P. H.; van Gunsteren, W. F. *J. Chem. Phys.* **1998**, *108*, 6117–6134.

(5) Mark, P.; Nilsson, L. *J. Comp. Chem.* **2002**, *23*, 1211–1219.

(6) Anézo, C.; de Vries, A. H.; Höltje, H.-D.; Tieleman, D. P.; Marrink, S.-J. *J. Phys. Chem. B.* **2003**, *107*, 9424–9433.

(7) Patra, M.; Karttunen, M.; Hyvonen, M. T.; Falck, E.; Lindqvist, P.; Vattulainen, I. *Biophys. J.* **2003**, *84*, 3636–3645.

(8) Feller, S. E.; Pastor, R. W.; Rojnuckarin, A.; Bogusz, S.; Brooks, B. R. *J. Phys. Chem.* **1996**, *100*, 17011–17020.

(9) Lindahl, E.; Hess, B. A.; van der Spoel, D. *J. Mol. Mod.* **2001**, *7*, 306–317.

(10) Lagüe, P.; Pastor, R. W.; Brooks, B. R. *J. Phys. Chem. B* **2004**, *108*, 363–368.

(11) Wu, X.; Brooks, B. R. *J. Chem. Phys.* **2005**, *122*, 044107.

(12) Wu, X.; Brooks, B. R. *J. Chem. Phys.* **2008**, *129*, 154115.

(13) Klauda, J. B.; Wu, X.; Pastor, R. W.; Brooks, B. R. *J. Phys. Chem. B* **2007**, *111*, 4393–4400.

(14) Venable, R. M.; Chen, L. E.; Pastor, R. W. *J. Phys. Chem. B.* **2009**, *113*, 5855–5862.

(15) in't Veld, P. J.; Ismail, A. E.; Grest, G. S. *J. Chem. Phys.* **2007**, *127*, 144711.

(16) Isele-Holder, R. E.; Mitchell, W.; Ismail, A. E. *J. Chem. Phys.* **2012**, *137*, 174107.

(17) Isele-Holder, R. E.; Mitchell, W.; Hammond, J. R.; Kohlmeyer, A.; Ismail, A. E. *J. Chem. Theory Comput.* **2013**, *9*, 5412–5420.

(18) Wennberg, C. L.; Murtola, T.; Hess, B.; Lindahl, E. *J. Chem. Theory Comput.* **2013**, *9*, 3527–3537.

(19) Good, R. J.; Hope, C. J. *J. Chem. Phys.* **1971**, *55*, 111–115.

(20) Lorentz, H. A. *Annalen der Physik* **1881**, *248*, 127–136.

(21) Ewald, P. P. *Ann. Phys.* **1921**, *64*, 253–287.

(22) Karasawa, N.; III, W. A. G. *J. Phys. Chem.* **1993**, *93*, 7320–7327.

(23) Toukmaji, A.; Board, J. A. *Comp. Phys. Comm.* **1996**, *95*, 73–92.

(24) Deserno, M.; Holm, C. *J. Chem. Phys.* **1999**, *109*, 7678–7693.

(25) Hub, J. S.; Winkler, F. K.; Merrick, M.; de Groot, B. L. *J. Am. Chem. Soc.* **2010**, *132*, 13251–13263.

(26) Berendsen, H. J. C.; Postma, J. P. M.; van Gunsteren, W. F.; Hermans, J. In *Intermolecular Forces*; Pullman, B., Ed.; D. Reidel Publishing Company: Dordrecht, 1981; pp 331–342.

(27) Hess, B. *J. Chem. Theory Comput.* **2008**, *4*, 116–122.

(28) Berger, O.; Edholm, O.; Jähnig, F. *Biophys. J.* **1997**, *72*, 2002–2013.

(29) Jorgensen, W. L.; Tirado-Rives, J. *J. Am. Chem. Soc.* **1988**, *110*, 1657–1666.

(30) Berendsen, H. J. C.; Grigera, J. R.; Straatsma, T. P. *J. Phys. Chem.* **1987**, *91*, 6269–6271.

(31) Yoluk, O.; Brömstrup, T.; Bertaccini, E.; Trudell, J.; Lindahl, E. *Biophys. J.* **2013**, *105*, 640–647.

(32) Páll, S.; Hess, B. *Comp. Phys. Comm.* **2013**, *184*, 2641–2650.

(33) Chiu, S.; Pandit, S. A.; Scott, H. L.; Jakobsson, E. *J. Phys. Chem. B* **2009**, *113*, 2748–2763.

(34) Egberts, E.; Marrink, S. J.; Berendsen, H. J. C. *Eur. Biophys. J.* **1994**, *22*, 423–436.

(35) Klauda, J. B.; Venable, R. M.; Freites, J. A.; O'Connor, J. W.; Tobias, D. J.; Mondragon-Ramirez, C.; Vorobyov, I.; A. D. MacKerell, J.; Pastor, R. W. *J. Phys. Chem. B* **2010**, *114*, 7830–7843.

Tables

Table 1: Calculated area per lipid from 100 ns simulations of POPC using LJ-PME. Three different setups were utilized: (I) full treatment using LB combination rules, (II) approximation of reciprocal space using geometric combination rules, and (III) geometric combination rules in reciprocal space combined with the modified interactions in direct space. The averages are obtained from the last 20 ns of the simulations.

Area/lipid [Å^2]		
Simulation I	Simulation II	Simulation III
66.9 ±0.4	66.0 ±0.3	66.9 ±0.2

Table 2: Strong and weak scaling measured on a Blue Gene/Q system, for a simulation of SPC/E-water. The strong scaling was measured on a system of 42624 water molecules, while the weak scaling was measured using 999 particles per core. The absolute performance is presented in ns/day, and for the weak scaling the time spent calculating the PME mesh or direct-space non-bonded interactions are shown relative to the total simulation time.

Strong Scaling [ns/day]				
Number of cores	512	1024	2048	4096
LJ-PME	27	42	52	66
Cutoff (10 Å)	37	58	76	92
Weak Scaling				
Number of cores		512	4096	32768
Performance [ns/day]		27	25	19
Non-bonded forces [%]		45	43	33
PME mesh [%]		22	22	23

Figures

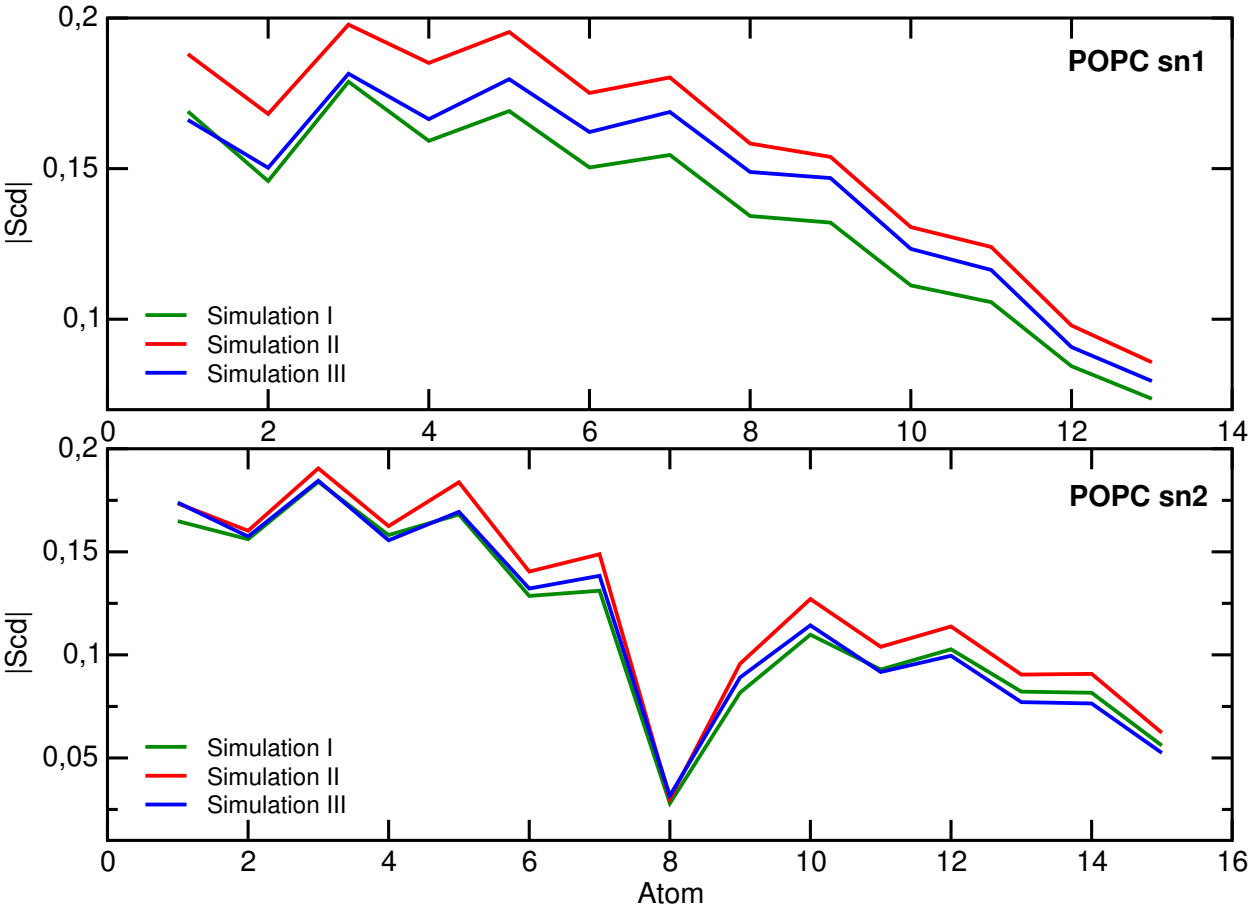

Figure 1: Deuterium order parameters averaged over the last 20 ns of simulation for the simulated LJ-PME systems. (I) full treatment with LB combination rules, (II) approximation of reciprocal space using geometric combination rules, (III) geometric combination rules in reciprocal space and modified interactions in direct space.

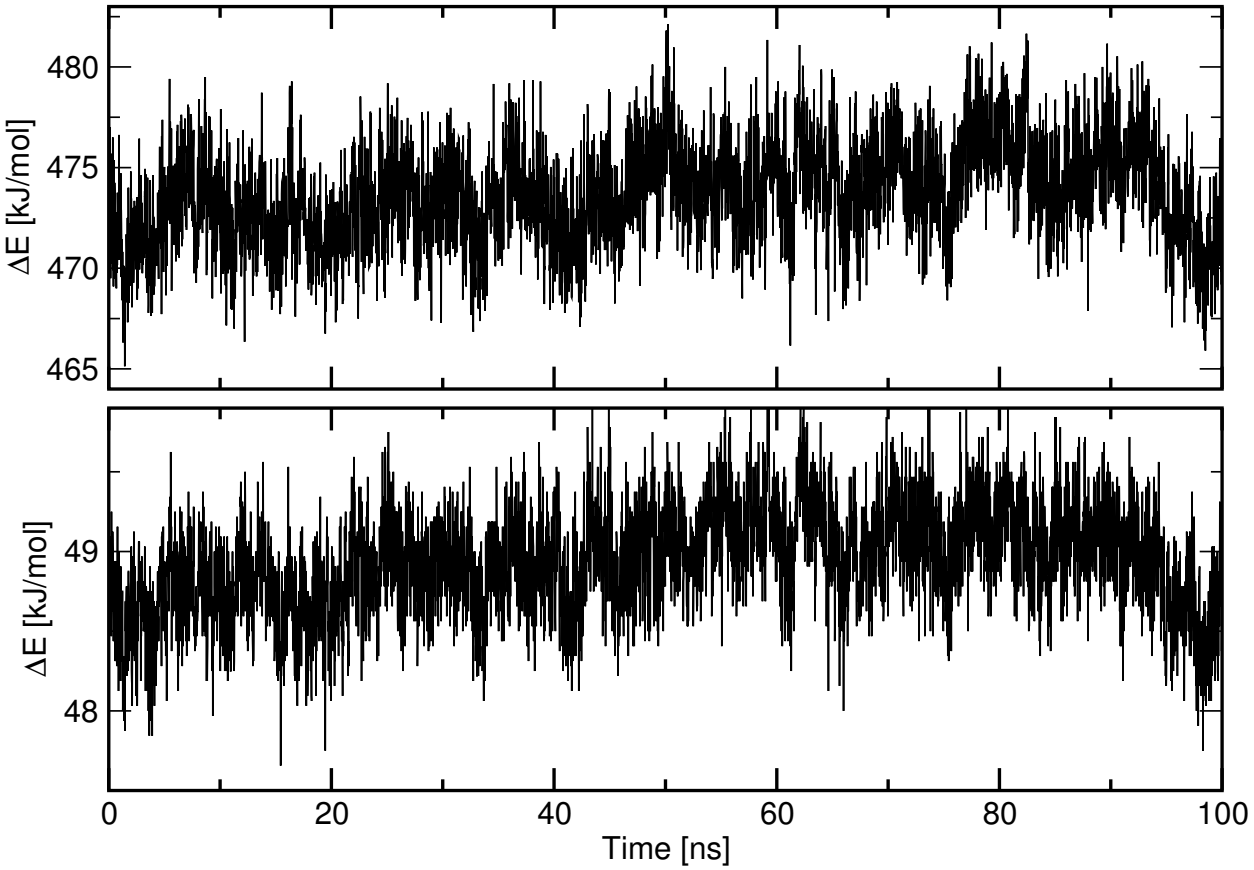

Figure 2: Absolute difference in the total energy from simulation I, compared at every frame with recalculated values using the setup from simulation II (top) or III (bottom). The results are obtained by reevaluating the energy at every frame in simulation I, using either one of the other two methodologies. The difference in the energy is reduced by an order of magnitude in simulation III, compared to simulation II. Total energy of the LB-system is approximately -376,000 kJ/mol.

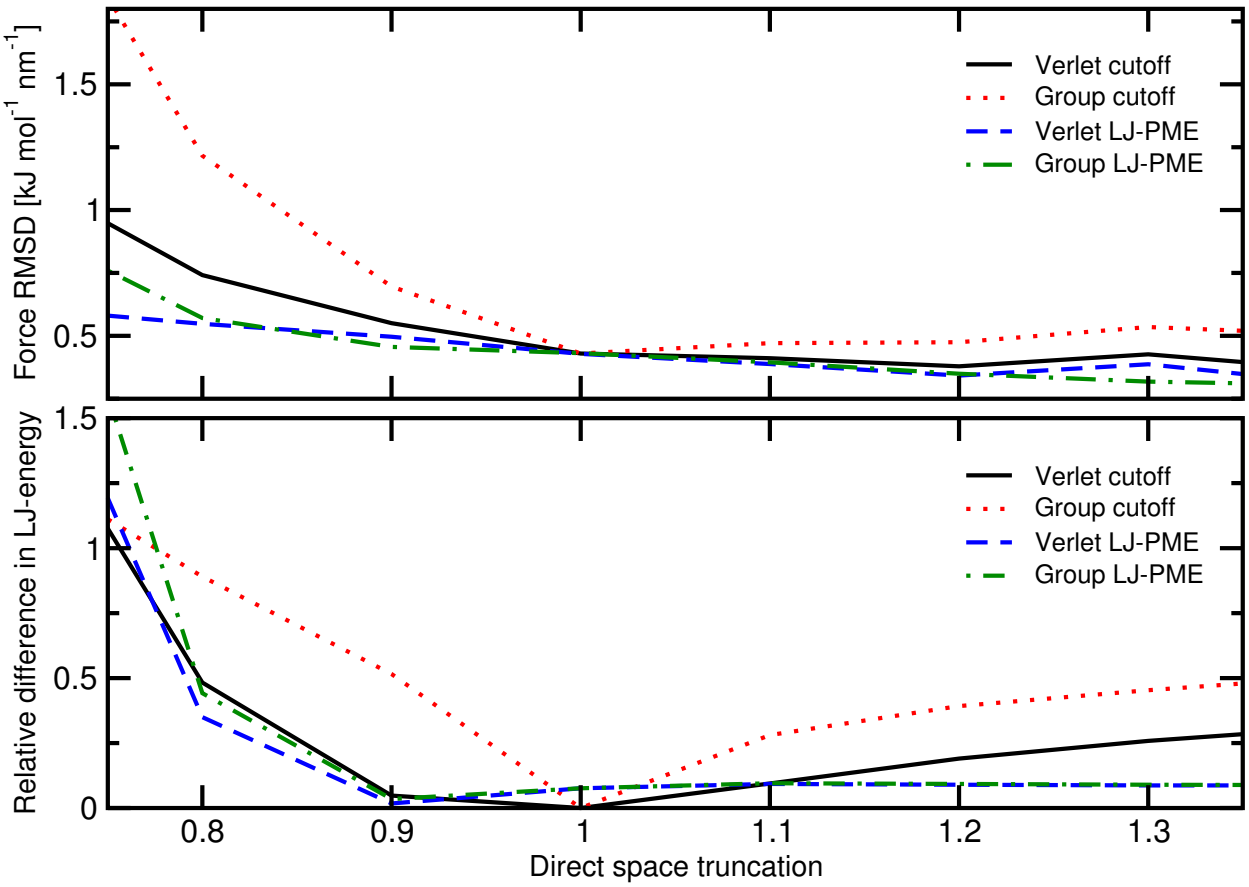

Figure 3: Force-RMSD and relative differences in total Lennard-Jones energy calculated for the POPC-system used in this work, as a function of the truncation in direct space, using either LJ-PME or dispersion correction to correct for long-range LJ-interactions. The reference systems utilized a direct-space cutoff of 10 Å, dispersion correction or LJ-PME for long-range LJ-interactions, and a reciprocal grid spacing of 0.2 Å with a sixth-order spline interpolation. Calculations were performed using either the Verlet[32] or group cutoff scheme in Gromacs.

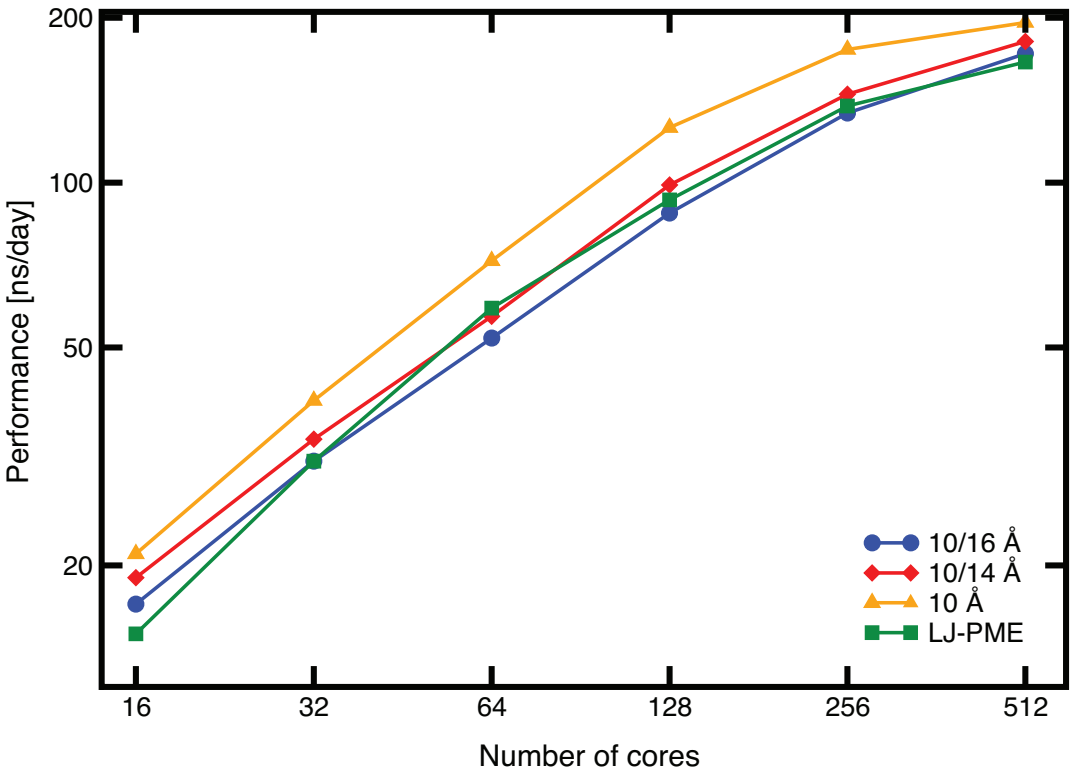

Figure 4: Performance obtained from simulations of a DPPC bilayer (58460 atoms in total) treating Lennard-Jones interactions with either cutoffs (10 Å or twin-range 10/14 Å and 10/16 Å), or using LJ-PME with geometric combination rules in reciprocal space. The simulations were performed on Intel E5-2660 CPUs, and performance was measured over 50,000 simulation steps.

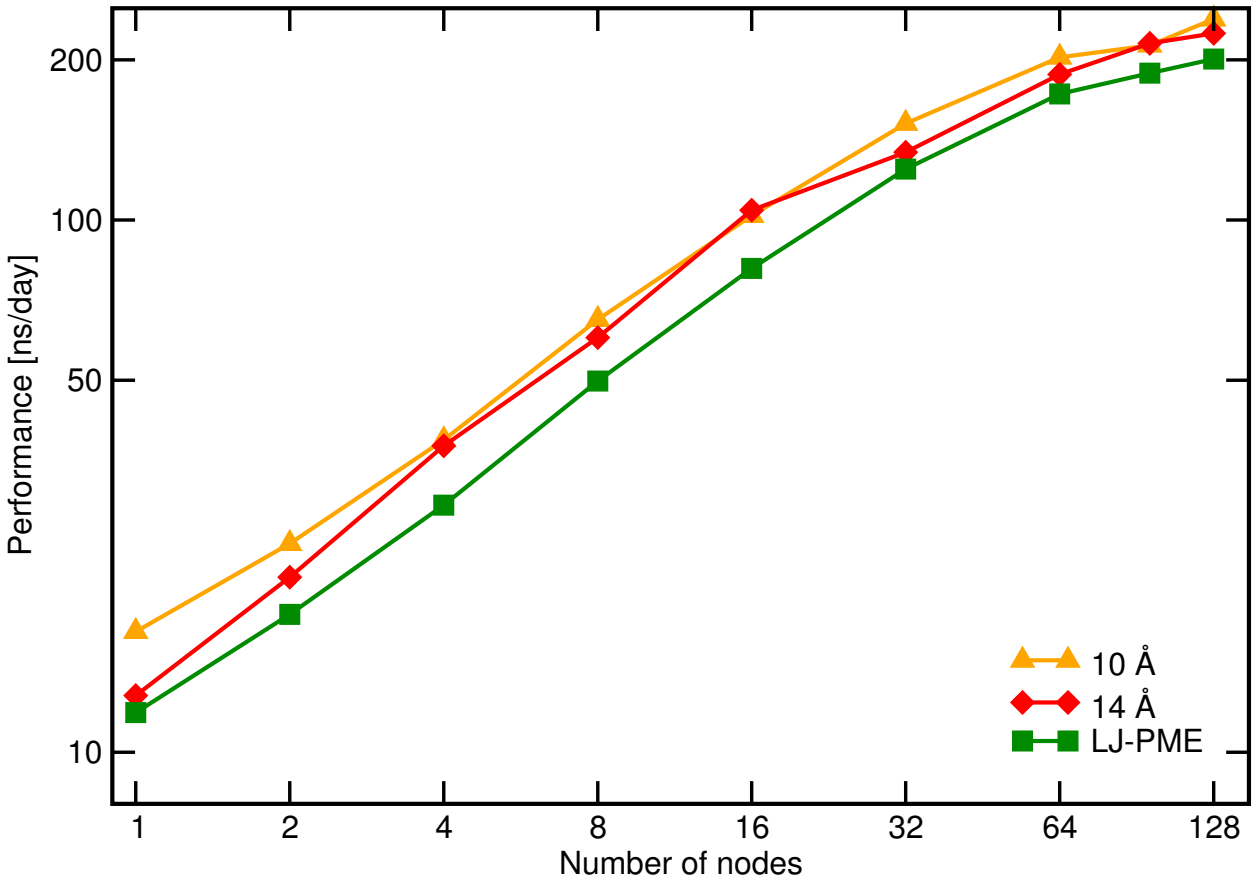

Figure 5: Performance from simulations of GluCl (142 000 atoms) in which the Lennard-Jones interactions were treated using either a 10 or 14 Å cutoff or with LJ-PME using geometric combination rules. The simulations were performed on a heterogeneous hardware setup, each node containing Intel Xeon E5-2670 CPUs and an NVIDIA Tesla K20X GPU. Performance was measured over the second half of three minute long runs.