A comparison of the Spectral Ewald and Smooth Particle Mesh Ewald methods in GROMACS

Davood Saffar Shamshirgar[a], Berk Hess[b], Anna-Karin Tornberg[a]

[a] *KTH Mathematics, Swedish e-Science Research Centre, 100 44 Stockholm, Sweden.*
[b] *KTH Computational Physics, Swedish e-Science Research Centre, 106 91 Stockholm, Sweden.*

arXiv:1712.04718v1 [math.NA] 13 Dec 2017

Abstract

The smooth particle mesh Ewald (SPME) method is an FFT based method for the fast evaluation of electrostatic interactions under periodic boundary conditions. A highly optimized implementation of this method is available in GROMACS, a widely used software for molecular dynamics simulations. In this article, we compare a more recent method from the same family of methods, the spectral Ewald (SE) method, to the SPME method in terms of performance and efficiency. We consider serial and parallel implementations of both methods for single and multiple core computations on a desktop machine as well as the Beskow supercomputer at KTH Royal Institute of Technology. The implementation of the SE method has been well optimized, however not yet comparable to the level of the SPME implementation that has been improved upon for many years. We show that the SE method is very efficient whenever used to achieve high accuracy and that it already at this level of optimization can be competitive for low accuracy demands.

Keywords: Fast Ewald summation, Fast Fourier transform, Coulomb potentials, Smooth particle mesh Ewald, Spectral Ewald, GROMACS.

1. Introduction

In molecular dynamics simulations (MD), the most time consuming task is often the computation of the electrostatic interactions. This remains true despite the fact that many algorithms exist that can reduce the computational cost from $\mathcal{O}(N^2)$ to $\mathcal{O}(N \log N)$ for evaluating the interaction of N particles under the common assumption of periodic boundary conditions. Therefore, this can not be considered a solved problem, and improvements are highly desirable.

The problem can be stated as follows. Suppose N charged particles with charges $q_{\mathtt{n}}$ are located in positions $\mathbf{x}_{\mathtt{n}}$, $\mathtt{n} = 1, \ldots, N$, and interacting in a simulation box $\Omega = [0, L)^3$. Moreover, assume the so-called *charge neutrality* condition, $\sum_{\mathtt{n}=1}^{N} q_{\mathtt{n}} = 0$. This condition is added to avoid the divergence of the electric potential energy in the presence of periodic boundary conditions. The objective is to compute the electrostatic potential,

$$\varphi(\mathbf{x}_{\mathtt{m}}) = {\sum_{\mathbf{p}}}' \sum_{\mathtt{n}=1}^{N} \frac{q_{\mathtt{n}}}{|\mathbf{x}_{\mathtt{mn},\mathbf{p}}|}, \quad \mathbf{x}_{\mathtt{mn},\mathbf{p}} = \mathbf{x}_{\mathtt{mn}} + \mathbf{p}L = \mathbf{x}_{\mathtt{m}} - \mathbf{x}_{\mathtt{n}} + \mathbf{p}L, \quad \mathbf{p} \in \mathbb{Z}^3, \tag{1.1}$$

Email addresses: davoudss@kth.se (Davood Saffar Shamshirgar), hess@kth.se (Berk Hess), akto@kth.se (Anna-Karin Tornberg)

or two related quantities, the electrostatic potential energy,

$$E = \frac{1}{2}\sum_{\mathtt{m}=1}^{N} q_{\mathtt{m}}\varphi(\mathbf{x}_{\mathtt{m}}), \tag{1.2}$$

and the electrostatic force,

$$\mathbf{F}(\mathbf{x}_{\mathtt{m}}) = -\frac{\partial E}{\partial \mathbf{x}_{\mathtt{m}}} = -\frac{1}{2} q_{\mathtt{m}} \frac{\partial \varphi(\mathbf{x}_{\mathtt{m}})}{\partial \mathbf{x}_{\mathtt{m}}}, \tag{1.3}$$

$\mathtt{m} = 1, \ldots, N$. The sum over $\mathbf{p}$ is over all periodic images of particles and the prime ($'$) denotes that the term with $\{\mathtt{m} = \mathtt{n}, \mathbf{p} = 0\}$ is excluded from the sum. The electrostatic potential (1.1) decays as $1/r$ (inverse of the distance between particles) and is very slow to compute in this form. Moreover, the sum is only conditionally convergent and the result depends on the order of summation. If the sum is evaluated in a spherical order, it yields the same result as the Ewald potential [21]. The *Ewald summation* technique proposed by Ewald [8] in 1921 can be used to compute the conditionally convergent sums of (1.1)-(1.3) by decomposition of $1/r$ into two parts as

$$\frac{1}{r} = \frac{\operatorname{erfc}(\xi r)}{r} + \frac{\operatorname{erf}(\xi r)}{r}, \tag{1.4}$$

where $\operatorname{erf}(\cdot)$ and $\operatorname{erfc}(\cdot)$ are the *error function* and the *complementary error function* respectively and $\xi > 0$ is called the *Ewald parameter*. The first term in (1.4) tends to zero quickly as $r \to \infty$. The second term, on the other hand, is a smoothly varying function with a rapidly converging representation in Fourier space. Using the decomposition (1.4) we can rewrite (1.1) and (1.2) as

$$\begin{aligned}
\varphi(\mathbf{x}_{\mathtt{m}}) =& \varphi^{\mathrm{S}}_{\mathtt{m}} + \varphi^{\mathrm{L}}_{\mathtt{m}} + \varphi^{\mathrm{self}}_{\mathtt{m}} \\
=& \sum_{\mathtt{n}} \sum_{\mathbf{p}}{}' q_{\mathtt{n}} \frac{\operatorname{erfc}(\xi|\mathbf{x}_{\mathtt{mn},\mathbf{p}}|)}{|\mathbf{x}_{\mathtt{mn},\mathbf{p}}|} + \frac{4\pi}{L^3} \sum_{\mathbf{k}\neq 0} \frac{e^{-k^2/4\xi^2}}{k^2} \sum_{\mathtt{n}} q_{\mathtt{n}} e^{\mathrm{i}\mathbf{k}\cdot\mathbf{x}_{\mathtt{mn}}} - \frac{2\xi}{\sqrt{\pi}} q_{\mathtt{m}}, \qquad (1.5) \\
E =& E^{\mathrm{S}} + E^{\mathrm{L}} + E^{\mathrm{self}} \\
=& \frac{1}{2} \sum_{\mathtt{m}} \sum_{\mathtt{n}} \sum_{\mathbf{p}}{}' q_{\mathtt{m}} q_{\mathtt{n}} \frac{\operatorname{erfc}(\xi|\mathbf{x}_{\mathtt{mn},\mathbf{p}}|)}{|\mathbf{x}_{\mathtt{mn},\mathbf{p}}|} \\
&+ \frac{4\pi}{2L^3} \sum_{\mathbf{k}\neq 0} \frac{e^{-k^2/4\xi^2}}{k^2} \sum_{\mathtt{m}} \sum_{\mathtt{n}} q_{\mathtt{m}} q_{\mathtt{n}} e^{\mathrm{i}\mathbf{k}\cdot\mathbf{x}_{\mathtt{mn}}} - \frac{\xi}{\sqrt{\pi}} \sum_{\mathtt{m}} q_{\mathtt{m}}^2. \qquad (1.6)
\end{aligned}$$

The electrostatic force can also be obtained by differentiating the energy equation (1.6) with respect to $\mathbf{x}_{\mathtt{m}}$,

$$\begin{aligned}
\mathbf{F}(\mathbf{x}_{\mathtt{m}}) =& F^{\mathrm{S}} + F^{\mathrm{L}} \\
=& q_{\mathtt{m}} \sum_{\mathtt{n}} q_{\mathtt{n}} \sum_{\mathbf{p}}{}' \left(\frac{2\xi}{\sqrt{\pi}} e^{-\xi^2|\mathbf{x}_{\mathtt{mn},\mathbf{p}}|^2} + \frac{\operatorname{erfc}(\xi|\mathbf{x}_{\mathtt{mn},\mathbf{p}}|)}{|\mathbf{x}_{\mathtt{mn},\mathbf{p}}|} \right) \frac{\mathbf{x}_{\mathtt{mn},\mathbf{p}}}{|\mathbf{x}_{\mathtt{mn},\mathbf{p}}|^2} \\
&- \frac{4\pi \mathrm{i} q_{\mathtt{m}}}{2L^3} \sum_{\mathbf{k}\neq 0} \frac{\mathbf{k} e^{-k^2/4\xi^2}}{k^2} \sum_{\mathtt{n}} q_{\mathtt{n}} e^{\mathrm{i}\mathbf{k}\cdot\mathbf{x}_{\mathtt{mn}}}, \qquad (1.7)
\end{aligned}$$

where $\mathbf{k} \in \{2\pi\mathbf{n}/L, \mathbf{n} \in \mathbb{Z}^3\}$ and $k = |\mathbf{k}|$. In (1.5)-(1.7), superscripts *S*, *L* and *self* refer to *Short range* (computed in real space), *Long range* (computed in Fourier space) and *Self interaction*

contribution respectively. The decomposition parameter ξ, controls the convergence of the two sums in a way that larger ξ results in a faster convergence of the real space sum in expense of slow convergence of the Fourier space sum. Furthermore, equations (1.5)-(1.7) are exact in the current form and their results are invariant to the Ewald parameter. In a charge neutral system, the mean-value of the potential will be zero and, therefore, the zero mode of the Fourier transformed potential is set to zero, i.e. $\widehat{\varphi}(\mathbf{k}=0) = \langle \varphi^L(\mathbf{x}) \rangle = 0$.

In the formulas above, the real space sum can be evaluated ignoring interactions between particles further apart than a given distance such that the $\operatorname{erfc}(\cdot)$ function is well enough decayed. The Fourier space sum can be evaluated using various methods that differ in structure. Among all methods proposed for the Fourier space sum, are the fast Fourier transform (FFT) based methods [5, 6, 7, 12, 15], non-FFT based or Hierarchical methods [9] and PDE-based methods [22]. These methods are all widely used in major software packages such as GROMACS [14], AMBER [2], NAMD [18], and SCAFACOS [20].

The aim of FFT-based methods is to use the FFT in order to accelerate the calculation of the Fourier space sum. This allows for a different choice of the Ewald parameter to shift more work to the Fourier part, and decrease it in the real space calculations. With a proper scaling, these methods reduce the total complexity of the summation to $\mathcal{O}(N \log N)$.

Among the most important methods which benefit from the FFT, are methods within the Particle-Mesh Ewald (PME) family, including the original PME method by Darden *et al.* [5], the SPME method by Essmann *et al.* [7], Particle-Particle-Particle-Mesh Ewald ($\mathrm{P^3M}$) originally introduced by Hockney and Eastwood [12], its variant Interlaced $\mathrm{P^3M}$ by Neelov and Holm [16], and Particle-Mesh NFFT by Pippig and Potts [19]. In this paper, we review another method of the PME family, the Spectral Ewald (SE) method introduced by Lindbo & Tornberg [15], and compare it to the SPME method.

We provide a comparison between the spectral Ewald (SE) and smooth particle mesh Ewald (SPME) methods in GROMACS for different systems of particles under three-dimensional periodic boundary conditions. We discuss similarities and differences, advantages and disadvantages for the two methods, including the approximation errors that they introduce and the computational cost that is incurred for different parts of the algorithms.

The evaluation of the real space sum is the same in the two methods. Therefore, we only study the runtime estimate of this sum and do not discuss the details of the algorithm used for this calculation. In the evaluation of the k-space sum, a regular FFT grid is introduced. The difference between the two methods lies in the choice of window functions that are used to interpolate the point charges onto this grid in the first step of the algorithm, as well as to interpolate back to the irregular point locations from grid values at the end of the calculations - the other differences follow as a consequence of this choice.

In the SE method, the window function is a suitably scaled and truncated Gaussian. Given the number of points in the support of the Gaussian and the grid size, parameters are adjusted to scale it optimally. This way, approximation errors introduced in the SE method are essentially independent of the grid size, and the size of the FFT grid can be chosen considering only the truncation of the original Ewald k-space sum.

In contrast, the SPME method works with a pth order cardinal B-spline as the window function, which introduces an approximation error of order h^p, where h is the grid size of the FFT grid. Hence, the size of the FFT grid must be in general larger than the truncation of the k-space sum (more points, smaller h) to reduce the approximation errors.

In terms of computational cost to obtain a certain accuracy, the cost from computing FFTs will hence be smaller for the SE method as compared to the SPME method, since smaller grids

will be used. On the other hand, the number of points in the support of the Gaussian is larger than that of the support of the cardinal B-splines, and hence the cost of interpolating to and from the grid will be larger for the SE method. In terms of which method that offers the lowest computational cost, this will depend on systems under study and parameters, as well as the quality of the implementation. Generally speaking, the SE method will outperform the SPME method if high accuracy is required, whereas the SPME method will be more efficient if only a few digits of accuracy is asked for.

We have implemented the SE method in GROMACS ver 5.1 [14], using the same routines for the real space sum and for evaluating FFTs as the SPME method. The remaining implementation of the SE method has been well optimized, even if not as fine tuned as the implementation for the SPME method that has been available in GROMACS since late 1999. Hence, the computational runtimes for SE relative to SPME can be expected to be reduced by additional optimization of the implementation.

This article is organized as follows. In section 2.1 we review truncation error estimates in Ewald sums. The structure of the FFT-based methods is reviewed in section 3. Then, in section 4, we review the SE method and consider errors in the SE and SPME methods [7] in section 5 and 6 respectively. The reader that is familiar with the SPME method can get a quick overview of the differences in implementation between the two methods by glancing at table 2 in section 6.1. In section 7 we discuss runtime estimates of the SE method in GROMACS. The numerical results for the serial and parallel implementations are presented in sections 8 and 9.

2. Ewald summation method

2.1. Truncation error estimates

The sums in (1.5) are all infinite and need to be truncated in practice. The error committed due to this truncation is discussed in this section. To consider one measure of accuracy, we use the root mean square (rms) error,

$$e_{\text{rms}}^2 := \frac{1}{N}\sum_{\mathtt{n}=1}^{N}(\varphi(\mathbf{x}_{\mathtt{n}}) - \varphi^*(\mathbf{x}_{\mathtt{n}}))^2,$$

where φ^* denotes an exact or a well converged reference solution. Choosing a cut-off radius $r_{\text{c}} < L/2$, the real space sum in (1.5)-(1.7) is truncated such that the sum only includes the terms for which $|\mathbf{x}_{\mathtt{mn},\mathtt{p}}| \leq r_{\text{c}}$. Using the well established Kolafa & Perram error estimates [13], the real space truncation errors are estimated as

$$e_{\text{P,rms}}^{\text{S}} \approx (Qr_{\text{c}}/2L^3)^{1/2}(\xi r_{\text{c}})^{-2}e^{-r_{\text{c}}^2\xi^2}, \tag{2.1a}$$

$$e_{\text{E,rms}}^{\text{S}} \approx Q(r_{\text{c}}/2L^3)^{1/2}(\xi r_{\text{c}})^{-2}e^{-r_{\text{c}}^2\xi^2}, \tag{2.1b}$$

$$e_{\mathbf{F},\text{rms}}^{\text{S}} \approx 2Q(1/r_{\text{c}}L^3)^{1/2}e^{-r_{\text{c}}^2\xi^2}, \tag{2.1c}$$

for the potential, energy and force respectively and with $Q := \sum_{\mathtt{n}=1}^{N} q_{\mathtt{n}}^2$. The Fourier space sum includes an exponential term of the form $e^{-k^2/4\xi^2}$ which decays exponentially as k increases. This sum can also be truncated at some maximum wave number $k_\infty \in \mathbb{Z}^+$ such that it only includes the terms for which $|\mathbf{k}| \leq 2\pi k_\infty/L$. The Fourier space truncation errors are also

estimated as [13],

$$e^{\mathrm{L}}_{\mathrm{P,rms}} \approx \xi\pi^{-2}k_\infty^{-3/2}\sqrt{Q}e^{-(\pi k_\infty/\xi L)^2}, \tag{2.2a}$$

$$e^{\mathrm{L}}_{\mathrm{E,rms}} \approx \xi\pi^{-2}k_\infty^{-3/2}Qe^{-(\pi k_\infty/\xi L)^2}, \tag{2.2b}$$

$$e^{\mathrm{L}}_{\mathbf{F},\mathrm{rms}} \approx \xi(L\pi)^{-1}(8/k_\infty)^{1/2}Qe^{-(\pi k_\infty/\xi L)^2}. \tag{2.2c}$$

Given a certain error tolerance $\mathcal{E}$ and a cut-off radius r_c, the Ewald parameter ξ can be obtained using (2.1)

$$\xi_\mathrm{P} = \frac{1}{r_\mathrm{c}}\left[W\left(\frac{1}{\mathcal{E}}\sqrt{\frac{Qr_\mathrm{c}}{2L^3}}\right)\right]^{1/2}, \tag{2.3a}$$

$$\xi_\mathrm{E} = \frac{1}{r_\mathrm{c}}\left[W\left(\frac{1}{\mathcal{E}}Q\sqrt{\frac{r_\mathrm{c}}{2L^3}}\right)\right]^{1/2}, \tag{2.3b}$$

$$\xi_\mathbf{F} = \frac{1}{r_\mathrm{c}}\left[\log\left(\frac{2Q}{\mathcal{E}}\sqrt{\frac{1}{r_\mathrm{c}L^3}}\right)\right]^{1/2}. \tag{2.3c}$$

Here, $W(\cdot)$ denotes the *Lambertw* function and is defined as the inverse of $f(x) = xe^{-x}$. Having ξ in hand we find k_∞ through

$$k_{\mathrm{P},\infty} = \frac{\sqrt{3}L\xi}{2\pi}\left[W\left(\frac{4}{3L^2}(\frac{Q}{\pi\xi\mathcal{E}^2})^{2/3}\right)\right]^{1/2}, \tag{2.4a}$$

$$k_{\mathrm{E},\infty} = \frac{\sqrt{3}L\xi}{2\pi}\left[W\left(\frac{4}{3L^2}(\frac{Q^2}{\pi\xi\mathcal{E}^2})^{2/3}\right)\right]^{1/2}, \tag{2.4b}$$

$$k_{\mathbf{F},\infty} = \frac{L\xi}{2\pi}\left[W\left(\frac{2^8\xi^2Q^4}{\mathcal{E}^4L^6\pi^2}\right)\right]^{1/2}, \tag{2.4c}$$

for the potential, energy and force. For a fixed cut-off r_c and an error tolerance $\mathcal{E}$, we have $\xi_\mathrm{P} \leq \xi_\mathrm{E} \leq \xi_\mathbf{F}$ as given by (2.3a-c). Based on the relations in (2.4a-c) and for given values of ξ and $\mathcal{E}$, we have that, $k_{\mathrm{P},\infty}$ is the smallest among all three and $k_{\mathrm{E},\infty}$ and $k_{\mathbf{F},\infty}$ are comparable.

As was mentioned before, the Ewald parameter controls the relative decay of the real space and Fourier space sums. For a fixed error tolerance $\mathcal{E}$, allowing for a larger r_c will yield a smaller ξ, and in turn, a smaller k_∞, decreasing the computational cost of the Fourier space sum. The optimal ξ is chosen such that the computational cost of evaluating the real and Fourier space sums are balanced. This however depends on the algorithm, implementation and hardware used for simulations.

2.2. Computational complexity

The direct evaluation of the Ewald sums (1.5)-(1.7), has a complexity of $\mathcal{O}(N^2)$, where N is the number of particles in the system. With a proper scaling of ξ with N, and assuming a uniform distribution of charges, the computational cost of the Ewald sum can be reduced to $\mathcal{O}(N^{3/2})$ [17], though, with a large constant. In the next section we will consider methods that reduce the complexity to $\mathcal{O}(N\log N)$, also with a much smaller constant.

3. FFT-based methods

As mentioned earlier, the aim of FFT-based methods is to use the FFT in order to accelerate the calculation of the Fourier space sum. These methods reduce the total complexity of the summation from $\mathcal{O}(N^2)$ (or $\mathcal{O}(N^{3/2})$ by optimizing the involved parameters) to $\mathcal{O}(N \log N)$ with a proper scaling of the Ewald parameter. The FFT acceleration allows for a larger ξ, corresponding to a smaller r_c, and thereby reduces the cost of evaluating the real space sum. We shall get back to this discussion later in section 4.2 when we present the computational complexity of the SE method.

The principles of all FFT-based methods within the Particle-Mesh Ewald (PME) family, including the SPME and SE methods, are the same. All methods are inherited from the original P^3M and PME methods. In all these methods, the evaluation is conducted in Fourier space by some modification to make possible the use of the FFT efficiently. Considering the evaluation of the Fourier space part of the potential ($\varphi_{\mathtt{m}}^{\mathrm{L}}$ in (1.5)), the resulting algorithms have the following general steps,

Algorithm 1 The structure of FFT-based Ewald summation methods

1: A uniform grid is introduced.
2: Spreading step: Using a suitable window function, the charges of the non-uniform particles are distributed on the uniform grid.
3: A 3D FFT is performed on the grid function.
4: The grid function is scaled by a multiplication with a modified Fourier transformed Green's function in Fourier space.
5: A 3D inverse FFT is performed.
6: Gathering step: The potential (force or energy) is evaluated at targets with the same window function as in the spreading step.

The spreading and gathering steps can be performed using different window functions, resulting in different flavors of methods with different accuracies and computational costs. The modified Green's function in step 4 depends on the choice of the window function in steps 2 and 6. In all FFT-based methods, the Fourier space truncation error is determined by the number of grid points as given by the Kolafa & Perram error estimates, but the approximation errors depend on the details of the methods. We discuss more on the approximation errors of the SE and SPME methods in sections 5.1 and 6.2.

3.1. Force calculation

Referring back to (1.3), the calculation of the electrostatic force involves differentiation of the electrostatic energy. The differentiation can be performed in three different ways which differ in accuracy and computational cost. In the first approach, used in the original PME method, differentiation is done with i$\mathbf{k}$ multiplication in Fourier space. This is the most accurate approach to evaluate the force but it is computationally very expensive as it needs two extra 3D IFFTs. The second way of obtaining the force, used in the original P^3M method, is to perform a discrete differentiation on the mesh in real space. This method, however, is only useful for low accuracy demands. The last approach is to analytically differentiate the charge assignment function in real space. This method is employed in the SPME method and since it only requires one 3D IFFT, it is cheaper to perform as compared to the i$\mathbf{k}$ differentiation. Another important advantage of the latter approach is that, as a result of analytic differentiation of the energy, and independent of the truncation errors, the resulting scheme conserves energy. This, in addition to the performance, is an important reason for us to choose this scheme for the SE method as well (see section 4.1).

4. The Spectral Ewald method

Here, we review the Spectral Ewald (SE) method [15] for electrostatic calculations under fully periodic boundary conditions. As mentioned before, the method has the structure of all PME family methods but unlike other methods in the family, it is spectrally accurate. Consider the evaluation of the long range part of the electrostatic potential in (1.5),

$$\varphi^{\mathrm{L}}(\mathbf{x}_{\mathtt{m}}) = \frac{4\pi}{L^3} \sum_{\mathbf{k}\neq 0} \frac{e^{-k^2/4\xi^2}}{k^2} \sum_{\mathtt{n}} q_{\mathtt{n}} e^{\mathrm{i}\mathbf{k}\cdot\mathbf{x}_{\mathtt{mn}}}, \tag{4.1}$$

under the assumptions made in section 1. In the SE method, suitably scaled (and truncated) Gaussian functions are used in the spreading and gathering steps of Algorithm 1. To derive the relevant formulas, a free parameter $\eta \in (0,1)$ is introduced to split the Gaussian factor $e^{-k^2/4\xi^2}$ as,

$$e^{-k^2/4\xi^2} = \underbrace{e^{-\eta k^2/8\xi^2}}_{a} \underbrace{e^{-(1-\eta)k^2/4\xi^2}}_{b} \underbrace{e^{-\eta k^2/8\xi^2}}_{c}.$$

The term (a) will be used in the spreading step, (b) for multiplication in Fourier space and (c) in the gathering step. We start by writing (4.1) as,

$$\begin{aligned}
\varphi^{\mathrm{L}}(\mathbf{x}_{\mathtt{m}}) &= \frac{4\pi}{L^3} \sum_{\mathbf{k}\neq 0} \frac{e^{-k^2/4\xi^2}}{k^2} e^{\mathrm{i}\mathbf{k}\cdot\mathbf{x}_{\mathtt{m}}} \sum_{\mathtt{n}} q_{\mathtt{n}} e^{-\mathrm{i}\mathbf{k}\cdot\mathbf{x}_{\mathtt{n}}} \\
&= \frac{4\pi}{V} \sum_{\mathbf{k}\neq 0} \frac{e^{-(1-\eta)k^2/4\xi^2}}{k^2} e^{\mathrm{i}\mathbf{k}\cdot\mathbf{x}_{\mathtt{m}}} e^{-\eta k^2/8\xi^2} \sum_{\mathtt{n}} q_{\mathtt{n}} e^{-\eta k^2/8\xi^2} e^{-\mathrm{i}\mathbf{k}\cdot\mathbf{x}_{\mathtt{n}}}.
\end{aligned} \tag{4.2}$$

Now, (4.2) can be written as

$$\varphi^{\mathrm{L}}(\mathbf{x}_{\mathtt{m}}) = \frac{4\pi}{L^3} \sum_{\mathbf{k}\neq 0} \frac{e^{-(1-\eta)k^2/4\xi^2}}{k^2} e^{\mathrm{i}\mathbf{k}\cdot\mathbf{x}_{\mathtt{m}}} e^{-\eta k^2/8\xi^2} \widehat{H}_{-k}, \tag{4.3}$$

where

$$\widehat{H}_k := \sum_{\mathtt{n}} q_{\mathtt{n}} e^{-\eta k^2/8\xi^2} e^{\mathrm{i}\mathbf{k}\cdot\mathbf{x}_{\mathtt{n}}}. \tag{4.4}$$

The inverse Fourier transform of (4.4) is given by

$$H(\mathbf{x}) = \sum_{\mathtt{n}} q_{\mathtt{n}} \left(\frac{2\xi^2}{\pi\eta}\right)^{3/2} e^{-2\xi^2 |\mathbf{x}-\mathbf{x}_{\mathtt{n}}|_*^2/\eta}, \tag{4.5}$$

where * denotes that periodicity is applied in all directions. The function $H(\mathbf{x})$ is a smooth and periodic function, evaluated on a uniform grid. Algorithmically, evaluating $H(\mathbf{x})$ corresponds to the spreading step, step 2 in Algorithm 1. Since the charges are distributed on a uniform grid, an FFT can be applied to compute $\widehat{H}_k$ efficiently. The Fourier transformed smeared particle function $\widehat{H}_k$, are then scaled by $k^{-2}e^{-(1-\eta)k^2/4\xi^2}$ (step 4 in Algorithm 1). We define

$$\widehat{\widehat{H}}_k := \frac{e^{-(1-\eta)k^2/4\xi^2}}{k^2} \widehat{H}_k, \tag{4.6}$$

and with this, (4.3) is simplified to

$$\varphi^{\mathrm{L}}(\mathbf{x}_{\mathtt{m}}) = \frac{4\pi}{L^3} \sum_{\mathbf{k}\neq 0} \widehat{\widetilde{H}}_{-k} e^{-\eta k^2/8\xi^2} e^{\mathrm{i}\mathbf{k}\cdot\mathbf{x}_{\mathtt{m}}}. \tag{4.7}$$

Then, as a result of exerting the Parseval theorem on (4.7), one obtains

$$\varphi^{\mathrm{L}}(\mathbf{x}_{\mathtt{m}}) = 4\pi \left(\frac{2\xi^2}{\pi\eta}\right)^{3/2} \int_{\Omega} \widetilde{H}(\mathbf{x}) e^{-2\xi^2|\mathbf{x}-\mathbf{x}_{\mathtt{m}}|_*^2/\eta} \mathrm{d}\mathbf{x}. \tag{4.8}$$

This interpolation, corresponds to the gathering step, step 6 in Algorithm 1. Finally, the integral in (4.8) can be evaluated by the trapezoidal rule,

$$\varphi^{\mathrm{L}}_{\mathrm{SE}}(\mathbf{x}_{\mathtt{m}}) = 4\pi h^3 \left(\frac{2\xi^2}{\pi\eta}\right)^{3/2} \sum_{\mathtt{n}} \widetilde{H}(\mathbf{x}_{\mathtt{n}}) e^{-2\xi^2|\mathbf{x}_{\mathtt{n}}-\mathbf{x}_{\mathtt{m}}|_*^2/\eta}. \tag{4.9}$$

Due to the fact that $\widetilde{H}(\mathbf{x})$ is smooth and periodic (as a result of periodicity and smoothness of H), the trapezoidal rule yields spectral accuracy. Moreover, $\widetilde{H}(\mathbf{x})$ in (4.9) can be evaluated by applying an inverse FFT on $\widehat{\widetilde{H}}$. Note here that, the same interpolation function used in the spreading (4.5) and gathering (4.8) steps. The reader may consult [15] for more details on the derivation of the SE method. The formulas that will be used in the different steps of Algorithm 1 are summarized in table 2 of section 6.1, side by side with the formulas needed for the SPME method.

Consider a uniform grid with M points in each direction, where M is related to the maximum Fourier mode k_∞ by $M = 2k_\infty$. The Gaussians do not have compact support, and hence they have to be truncated. Let P, where $P \leq M$, be the number of points in the support of Gaussians in each direction. Also we denote by $w = Ph/2 = PL/2M$ the half width of the Gaussians. Earlier in this section, we introduced the parameter η which controls the width of Gaussians. At the point of truncating the Gaussians, $|\mathbf{x} - \mathbf{x}_{\mathtt{n}}|_* = w$, assume that the term $e^{-2\xi^2|\mathbf{x}-\mathbf{x}_{\mathtt{n}}|_*^2/\eta}$ has decayed to $e^{-m^2/2}$. Hence, the choice of m will completely determine η and control the level at which the Gaussians are truncated. Then,

$$\eta = \left(\frac{2w\xi}{m}\right)^2.$$

In section 5.1 we will explain how the choice of m is tied to P in order to find the optimal η. As a consequence, P remains the only parameter to select.

Truncating the Gaussians decreases the computational complexity of the spreading and gathering steps significantly, but the computations in these steps are still expensive as they involve evaluation of P^3 exponentials for each particle. Using the Fast Gaussian Gridding (FGG) approach suggested by Greengard & Lee [10], these computations are accelerated at the expense of using more memory space to store parts of the exponentials and reuse them. Here, we briefly describe the idea of the FGG and refer the reader to [10] or [15] for more details. Let $\mathbf{x}_{\mathtt{n}}$ be a particle position and $\mathbf{x} = (i, j, k)h$, $(i, j, k) \in \mathbb{Z}^3$ a grid point. Then, the evaluation of

$$e^{-c(\mathbf{x}-\mathbf{x}_{\mathtt{n}})^2} = e^{-c(ih-x_{\mathtt{n}})^2} e^{-c(jh-y_{\mathtt{n}})^2} e^{-c(kh-z_{\mathtt{n}})^2}, \tag{4.10}$$

involves M^3N exponential evaluations. Using the FGG approach we can rewrite (4.10) e.g. in the x-direction as,

$$e^{-c(ih-x_{\mathrm{n}})^2} = \underbrace{e^{-c(ih)^2}}_{a}\underbrace{(e^{2chx_{\mathrm{n}}})^i e^{-cx_{\mathrm{n}}^2}}_{b}. \tag{4.11}$$

Since the term (a) is independent of particle positions, it can be precomputed and stored. This computation on the whole grid involves evaluation of M^3 exponentials and requires M^3 storage. The term (b), on the other hand, can be precomputed once for each particle and requires evaluation of $4N$ exponentials, $3(2N + MN)$ multiplications and $3NM$ storage in three dimensions. Note that the first term in (b) is computed by successive multiplications with itself. The FGG approach is much more effective as compared to the direct approach when the Gaussians are truncated with P points in the support in each direction. Table 1 shows that using the FGG, the cost of evaluating (4.5) and (4.9) reduces from $\mathcal{O}(P^3N)$ to $\mathcal{O}(PN)$ with the expense of $\mathcal{O}(PN)$ multiplication and storage.

	Direct approach	**FGG approach**
exp(·) evaluation	P^3N	$P^3 + 3PN + 4N$
multiplication	-	$3(2N + PN)$
storage	-	$P^3 + 3PN$

Table 1: The cost of evaluating exponentials in (4.5) and (4.9).

4.1. Calculation of the electrostatic force

Based on the conclusion made in section 3.1, to compute the force in the SE method, we differentiate the spreading function. The charge spreading function in the SE method has the form of $e^{-2\xi^2|\mathbf{x}-\mathbf{x}_{\mathrm{n}}|_*^2/\eta}$. By analytic differentiation of this function and using (1.3), (4.8) is then reformulated as,

$$\mathbf{F}(\mathbf{x}_{\mathtt{m}}) = 2\pi\frac{4\xi^2}{\eta}\left(\frac{2\xi^2}{\pi\eta}\right)^{3/2} q_{\mathtt{m}}\int_{\Omega}(\mathbf{x}-\mathbf{x}_{\mathtt{m}})_*\widetilde{H}(\mathbf{x})e^{-2\xi^2|\mathbf{x}-\mathbf{x}_{\mathtt{m}}|_*^2/\eta}\mathrm{d}\mathbf{x}. \tag{4.12}$$

As in (4.9), this integral can be approximated by the trapezoidal rule to spectral accuracy. We remark here that since the force and potential differ only in the term $\mathbf{x}-\mathbf{x}_{\mathtt{m}}$, it is possible to compute these quantities simultaneously at the interpolation step.

4.2. Computational complexity of the SE method

The computational cost of the SE method can be divided into three parts; the spreading and gathering cost, the scaling cost, and the FFT cost. Consider a system of N particles located in a box of length L and assume that the system particle density $\rho = N/L^3$ is constant as N grows, i.e. $L^3 \propto N$ or equivalently $L \propto N^{1/3}$. We assume again that the domain $[0, L)^3$ is discretized with a total of M^3 points and truncated Gaussians have P^3 points in their supports. Therefore, the computational complexity of the spreading and gathering steps are $\mathcal{O}(NP^3)$. The scaling step involves M^3 multiplications and, therefore, the cost of this step scales as M^3. To evaluate the full complexity of the algorithm, and tie N to M, we need to account also for the computational complexity of the real space part of the Ewald sum. Let us assume that

the cut-off radius r_c is fixed. As N grows, the real space sum scales as N with a somewhat large constant that depends on the average number of nearest neighbors. Let τ be the level of accuracy at which we truncate the real and Fourier space sums. Since r_c is fixed, the error estimate (2.1) suggests that $e^{-\xi^2 r_c^2} \approx \tau$. This implies that ξ is fixed. Then, from the error estimate (2.2), $e^{-\pi k_\infty^2/(\xi L)^2} \approx \tau$ and therefore, $M \propto k_\infty \propto L$, i.e. $M^3 \propto N$. Moreover, if $h = L/M$ is fixed, there exists an integer P such that the approximation error also satisfies the requirement on the level of accuracy. Hence, the cost of the Fourier space scaling and FFTs are of the order N and $N \log N$ respectively, and the full algorithm has a complexity of $\mathcal{O}(N \log N)$.

5. Errors in the SE method

In addition to the truncation errors that stem from truncating the sums with infinite number of terms, there are also approximation errors introduced in the SE method. These errors occur due to truncation of the Gaussians and using quadrature rules to compute the integrals in (4.8) and (4.12), for the potential and force, respectively. Here, we aim to present experimental upper bounds for the approximation errors in the evaluation of the potential, energy and force and provide numerical results to assess the sharpness of the error bounds.

5.1. Approximation error in the potential calculation

The derivation of (4.8) (or (4.12)) involves no approximation, but we commit errors when we evaluate the integral with the trapezoidal rule and also truncate the Gaussians. The following theorem adapted from [15] provides an upper bound for the approximation error in evaluating the potential and sets a foundation for obtaining corresponding bounds for the energy and force.

Theorem 1. *(Error estimate in the potential calculation using the SE method). Given ξ, $h > 0$ and an integer $P > 0$, define $w = hP/2$ and $\eta = (2w\xi/m)^2$. The error committed in evaluating the Fourier space part of the potential* (4.1) *by truncating the Gaussians at $|\mathbf{x} - \mathbf{x}_\mathtt{m}| = w$ and applying the trapezoidal rule as in* (4.9) *can be estimated by*

$$|\varphi^{\mathrm{L}}_{\mathrm{SE}} - \varphi^{\mathrm{L}}|_{\mathrm{rms}} \leq |\varphi^{\mathrm{L}}_{\mathrm{SE}} - \varphi^{\mathrm{L}}|_\infty \leq C(e^{-\pi^2 P^2/(2m^2)} + \mathrm{erfc}(m/\sqrt{2})), \tag{5.1}$$

and C is a constant independent of m and P.

Proof. See [15]. □

Clearly, the error estimate is related to m, the shape controller, and P. It is natural to balance two terms such that they contribute equally in the error. This gives, $m = c\sqrt{\pi P}$ and we experimentally find $c = 0.95$ to set a good choice and hence

$$\eta = \left(\frac{2w\xi}{m}\right)^2 = \frac{PL^2\xi^2}{c^2\pi M^2}. \tag{5.2}$$

Therefore, we can write

$$|\varphi^{\mathrm{L}}_{\mathrm{SE}} - \varphi^{\mathrm{L}}|_{\mathrm{rms}} \lessapprox A_{\mathrm{P}} e^{-\pi P c^2/2}. \tag{5.3}$$

We emphasize here that, now, P is the only parameter to control the approximation error and in contrast to the other PME-like methods, the approximation error is independent of the grid size, M. To find A_{P}, we choose the grid size large enough for the truncation error to be negligible in the Fourier space sum. Through numerical experiments, we find that $A_{\mathrm{P}} \approx \sqrt{Q\xi L}/L$. To

confirm this, the error in the Fourier space part of the potential has been measured for 100 different cases and for each P including all possible combinations of $L \in \{1, 5, 10, 20, 40\}$, $N \in \{200, 400, 800, 1200\}$, and $\xi \in \{5, 15, 25, 30, 35\}/L$. In figure 1 (left), we plot absolute rms error in the potential calculation as scaled by A_{P}. We observe that the error for 100 different cases and all P collapse rather well under this scaling. It also turns out that if P is more than 24, independent of the grid size, machine accuracy is attained, [15].

5.2. Approximation error in the energy calculation

Following a similar approach as in Theorem 1, the approximation error in the energy can be written as,

$$|E^{\mathrm{L}}_{\mathrm{SE}} - E^{\mathrm{L}}|_{\mathrm{rms}} \lessapprox A_{\mathrm{E}} e^{-\pi P c^2/2}, \tag{5.4}$$

and considering (1.6), the natural choice would be to choose $A_{\mathrm{E}} \approx \frac{Q\sqrt{\xi L}}{L}$.

5.3. Approximation error in the force calculation

It is straightforward to do the similar procedure as in [15] to derive the approximation error in the force calculation. In (1.7) we have seen that differentiation of the potential yields an extra term $\frac{1}{2} q_{\mathtt{m}} (\frac{4\xi^2}{\eta})(\mathbf{x}_{\mathtt{m}} - \mathbf{x}_{\mathtt{n}})$. The term $(\mathbf{x}_{\mathtt{m}} - \mathbf{x}_{\mathtt{n}})$ can be bounded by w, the half width of the Gaussians. Hence,

$$|\mathbf{F}^{\mathrm{L}}_{\mathrm{SE}} - \mathbf{F}^{\mathrm{L}}|_{\mathrm{rms}} \lessapprox \frac{1}{2}\sqrt{Q}\left(\frac{4\xi^2 w}{\eta}\right) A_{\mathrm{P}} e^{-\pi P c^2/2}, \tag{5.5}$$

where $A_{\mathrm{P}} \approx \sqrt{Q\xi L}/L$. On the other hand, using the definition of η we have

$$\frac{4\xi^2 w}{\eta} = \frac{4\xi^2 w}{4\xi^2 w^2/m^2} = \frac{m^2}{w}.$$

To eliminate the truncation error, the term $e^{-(\pi k_\infty/\xi L)^2}$ in the error estimate (2.2) suggests that $M = 2k_\infty \approx 4\xi L$. Therefore, $w = Ph/2 = PL/2M \approx P/8\xi$. Also, since $m^2 \approx \pi P$, (5.5) can be simplified as

$$|\mathbf{F}^{\mathrm{L}}_{\mathrm{SE}} - \mathbf{F}^{\mathrm{L}}|_{\mathrm{rms}} \lessapprox A_{\mathbf{F}} e^{-\pi P c^2/2}, \quad A_{\mathbf{F}} \approx 4\pi Q\sqrt{\frac{\xi^3}{L}}. \tag{5.6}$$

Figure 1 (right) shows the experimental error bound and measured error in calculating the force as scaled with $A_{\mathbf{F}}$ using the SE method. The error is computed with the same 100 cases as in figure 1 (left) for the potential. For a specific error tolerance, the resulting P value obtained from the error bound (5.6) has to be rounded up to the next even integer to be able to utilize faster spreading/gathering routines available for even values of P. Hence, the fact that the approximation error bound is not as sharp as the bound for the potential, has no significant effect on the choice of P.

Remark 1. *The error bounds* (5.1) *and* (5.6) *suggest that the relative approximation error (obtained by scaling with* A_{P} *and* $A_{\mathbf{F}}$ *respectively) has the form of* $e^{-\pi P c^2/2}$*. Therefore with* $P \approx 10, 20, 24$*, independent of the other parameters, the relative approximation error is less than* 10^{-6}*,* 10^{-12} *and* 10^{-15} *respectively.*

Remark 2. *The force defined throughout this article is computed based on* (4.12) *but whenever we measure the absolute rms error in the SE and SPME methods in GROMACS, the force is scaled with the electrostatic conversion factor* $1/4\pi\epsilon_0 \approx 138.935$*.*

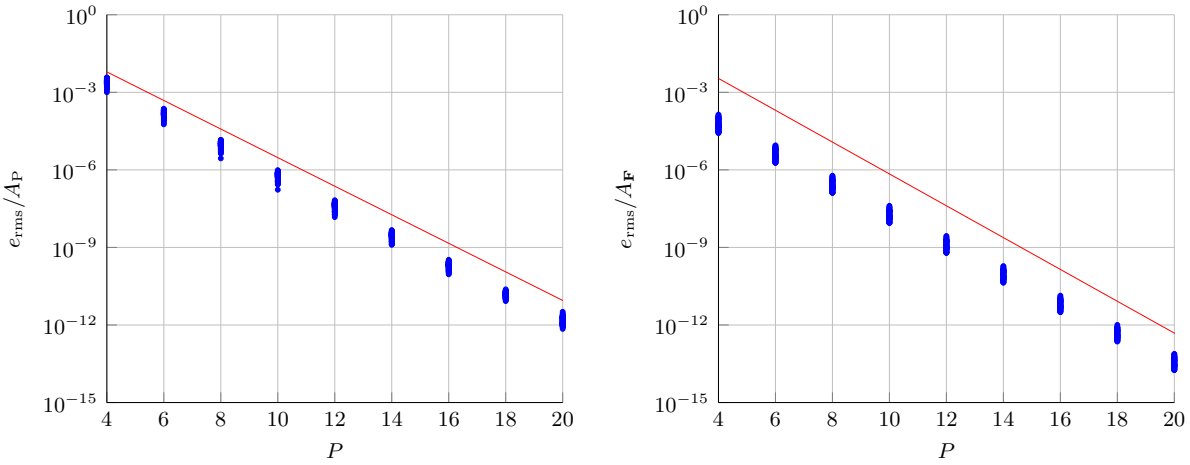

Figure 1: Scaled rms error in evaluating (left) the potential and (right) force. Scaled rms error confirms the dependency to P. The error is computed for each P and for 100 different cases with $L \in \{1, 5, 10, 20, 40\}$, $N \in \{200, 400, 800, 1200\}$, and $\xi \in \{5, 15, 25, 30, 35\}/L$. The scaled results satisfy the error bound $e^{-\pi P c^2/2}$ (red line).

5.4. Total error

There are two types of errors involved in calculation of the electrostatic potential (energy or force) using the SE method, namely the truncation and approximation errors. The truncation errors are estimated by (2.1) and (2.2) and the approximation error is estimated by (5.1), (5.4) and (5.6).

To study the behavior of the truncation errors, we increase P such that the approximation error is down to machine precision. We generate a uniformly distributed system of $N = 800$ particles in a box of length $L = 20$ and measure the absolute rms error as a function of the Ewald parameter for different grid sizes and cut-offs. Figure 2 (left) shows how the total error in evaluating the potential for $\{r_{\mathrm{c}} = 4, M = 32\}$ and $\{r_{\mathrm{c}} = 5, M = 64\}$ (solid curves) follow the truncation error estimates (2.1) and (2.2) (dashed curves). Next, we investigate the effect of the approximation error in the total error. We choose a cloud-wall system (figure 5-middle) with $N = 1200$ and $L = 10$ and compute the total error for $\{r_{\mathrm{c}} = 5, M = 64\}$ and different P values. Figure 2 (right) shows how a smaller P adds approximation error to the total error. Again, the dashed curves represent the error levels obtained from theoretical estimates of truncation errors and the solid curves are the total error in evaluating the potential. We demonstrate in figure 3, the total error in computing the force as a function of ξ using the SE method. As we observe in both figure 2 (left) and 3, there is an excellent agreement between the actual and estimated errors.

5.5. Momentum conservation

The conservation of momentum and energy are important criteria in the simulation of particles. Narrowing many events which affect the conservation to the methods that we use, some methods conserve momentum while others conserve energy. The conservation of energy seems more demanding than the conservation of momentum in molecular dynamics simulations. In classical Newtonian dynamics, the energy of an isolated system is conserved. Therefore, the molecular dynamics as a numerical solution of classical dynamics, has to conserve energy as well.

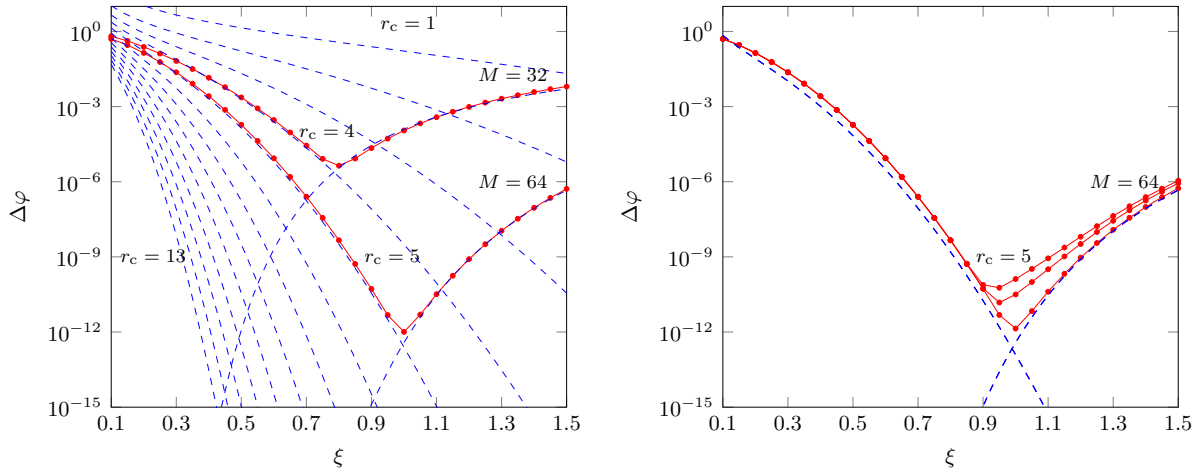

Figure 2: Measured and estimated absolute rms errors in computing the potential vs. ξ using the SE method. The dashed curves show truncation error levels for different r_c and grid size M based on the error estimates (2.1a) and (2.2a) and the solid curves represent the total measured error. A system of $N = 800$ uniformly distributed particles in a box of length $L = 20$ is used. (Left) $\{r_c = 4,\ M = 32\}$ and $\{r_c = 5,\ M = 64\}$. The parameter P is chosen such that the approximation error is negligible. (Right) $r_c = 5$ and $M = 64$. From top to bottom $P = 18, 20, 24$.

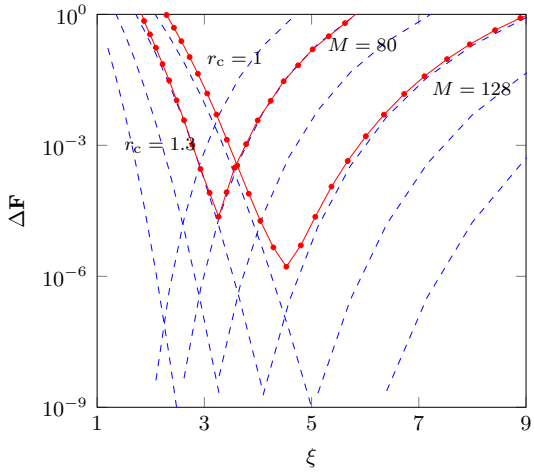

Figure 3: Measured and estimated absolute rms errors in computing the force vs. ξ using the SE method. The dashed curves show truncation error levels for different r_c and grid size M based on the error estimates (2.1c) and (2.2c) and the solid curves represent the total measured error with $\{r_c = 1,\ M = 128\}$ and $\{r_c = 1.3,\ M = 80\}$. The parameter P is chosen such that the approximation error is negligible. A cloud-wall system with $N = 1200$ particles in a box of length $L = 10$ is used (figure 5-middle).

In Darden *et al.* version of the PME method [5], the forces are interpolated symmetrically in space and the same function is used to distribute charges and to interpolate forces. Therefore, the PME method conserves momentum and not energy. On the contrary, the SPME method conserves energy and not momentum. Essmann *et al.* showed [7], as an example, that the sum of electrostatic forces is not zero and that the simulation leads to a small random particle drift. This, however, can be removed by subtracting the drift from all particle before doing another time iteration [7]. The momentum conservation can be examined by computing the sum of forces in a charge neutral system. Due to the fact that in the SE method the force is obtained through analytic differentiation of the energy, the method conserves energy, however, it does not conserve momentum. To have an understanding of the magnitude of the sum of forces, we compute the electrostatic force for a pair of oppositely charged particles. The experiment shows that the net force is not zero but is in the same order as the rms error. This is in agreement with the observation in [7].

6. Errors in the SPME method

In this section, we briefly review the SPME method and give all the formulas needed for its implementation, before we summarize the formulas for the SE and SPME method side by side in Table 2. We discuss the character of the approximation errors in the SPME method in comparison to those of the SE method in order to understand the different requirements on the FFT grid size in the two methods.

6.1. Methodology

The SPME method [7] is a reformulation of the PME method [5] which itself is inspired by the P^3M method [6], and follows the structure of the algorithm described in section 3. The specific choice of window functions is cardinal B-splines of different orders, as will be described below. The electrostatic force is obtained by analytic differentiation of the energy as generically described in section 3.1 and detailed for the SE method in section 4.1.

Suppose N particles with charges $q_{\mathtt{n}}$ at positions $\mathbf{x}_{\mathtt{n}}$, $\mathtt{n} = 1, 2, \ldots, N$, are interacting with each other in a unit box and let $\mathbf{x} = (x, y, z) = (a_1 h, a_2 h, a_3 h)$, $a_i \in \{0, 1, 2, \ldots\}$ for $i = 1, 2, 3$, with $h = 1/M$ for a positive integer M. The cardinal B-spline of order $p = 2$ is defined as

$$\mathcal{M}_2(u) = \begin{cases} 1 - |u - 1|, & 0 \le u \le 2, \\ 0, & \text{otherwise}, \end{cases}$$

and for $p > 2$ is defined recursively as

$$\mathcal{M}_p(u) = \frac{u}{p-1}\mathcal{M}_{p-1}(u) + \frac{p-u}{p-1}\mathcal{M}_{p-1}(u-1).$$

The charge distribution function $\sum_{\mathtt{n}} q_{\mathtt{n}} \delta(\mathbf{x} - \mathbf{x}_{\mathtt{n}})$ is smeared by a convolution with an order p cardinal B-spline on a uniform grid of size M^3 as

$$Q(\mathbf{x}) = \sum_{\mathtt{n}=1}^{N} q_{\mathtt{n}} \boldsymbol{\mathcal{M}}_p(\mathbf{x}_{\mathtt{n}} - \mathbf{x})_* = \sum_{\mathtt{n}=1}^{N} q_{\mathtt{n}} \mathcal{M}_p(x_{\mathtt{n}} - x)_* \mathcal{M}_p(y_{\mathtt{n}} - y)_* \mathcal{M}_p(z_{\mathtt{n}} - z)_*. \tag{6.1}$$

Here, as before, * denotes that periodicity is applied in all direction. The gridded distribution function $Q(\mathbf{x})$ acts in the same way as $H(\mathbf{x})$, cf. (4.5), introduced in the SE method. The

discrete Fourier transform of Q, denoted by $\widehat{Q}$, can then be computed via a 3D FFT. The scaled gridded distribution function now is given by a multiplication in Fourier space,

$$\widehat{\widetilde{Q}}(\mathbf{k}) = B(\mathbf{k})\frac{e^{-k^2/4\xi^2}}{k^2}\widehat{Q}(\mathbf{k}), \tag{6.2}$$

where

$$B(k_1,k_2,k_3) = |b_1(k_1)|^2|b_2(k_2)|^2|b_2(k_3)|^2, \quad b_j(k_j) = \frac{e^{(p-1)k_j h}}{\sum_{\ell=0}^{p-2}\mathcal{M}_p((\ell+1)h)e^{\mathrm{i}k_j\ell h}}.$$

An inverse 3D FFT transforms $\widehat{\widetilde{Q}}$ to the real space (denoted by $\widetilde{Q}$) and the result is then interpolated using another cardinal B-spline to recover the potential at target points, i.e.

$$\varphi^{\mathrm{L}}_{\mathrm{SPME}}(\mathbf{x}_{\mathtt{m}}) = \frac{4\pi}{L^3}\sum_{\mathtt{n}}\widetilde{Q}(\mathbf{x}_{\mathtt{n}})\boldsymbol{\mathcal{M}}_p(\mathbf{x}_{\mathtt{n}}-\mathbf{x}_{\mathtt{m}})_*. \tag{6.3}$$

As mentioned previously, the force field is derived with analytic differentiation when needed [7]. In table 2, we give the formulas for the SE and SPME methods, where the steps correspond to those introduced in Algorithm 1.

steps	**SE**	**Eq**	**SPME**	**Eq.**
1	regular grid of size M^3		regular grid of size M^3	
2	$\left(\frac{2\xi^2}{\pi\eta}\right)^{3/2} e^{-2\xi^2\|\mathbf{x}_{\mathtt{n}}-\mathbf{x}\|_*^2/\eta}$	4.5	$\boldsymbol{\mathcal{M}}_p(\mathbf{x}_{\mathtt{n}}-\mathbf{x})_*$	6.1
3	Apply 3D FFT		Apply 3D FFT	
4	$k^{-2}e^{-(1-\eta)k^2/4\xi^2}$	4.6	$B(\mathbf{k})k^{-2}e^{-k^2/4\xi^2}$	6.2
5	Apply 3D IFFT		Apply 3D IFFT	
6	$\left(\frac{2\xi^2}{\pi\eta}\right)^{3/2} e^{-2\xi^2\|\mathbf{x}_{\mathtt{n}}-\mathbf{x}\|_*^2/\eta}$	4.9	$\boldsymbol{\mathcal{M}}_p(\mathbf{x}_{\mathtt{n}}-\mathbf{x})_*$	6.3

Table 2: A stepwise comparison of the SE and SPME methods with the steps as given in Algorithm 1 and corresponding equation numbers. The number of points in the support of the truncated Gaussian (P, which sets η from equation (5.2)) must be chosen for the SE method and the degree of the B-splines (p) for the SPME method.

6.2. Approximation error

In the SPME method, as a result of using pth order B-splines in the spreading/gathering steps, an approximation error of order h^p, $h = L/M$ is added. Hence, to reduce this type of error, one needs to increase the grid size. Consequently, the cost of scaling and specially FFT steps increases. Essmann *et al.* [7] used B-splines of order $p = 3$ for low accuracies ($< 10^{-4}$), $p = 5$ for moderate accuracies ($< 10^{-6}$) and $p = 7$ for high accuracies ($< 10^{-8}$) but they did not present a systematic way of choosing B-spline orders since they had no estimate of the approximation error available. In 2012, Ballenegger *et al.* [4] showed that the P^{3}M method can be converted into the SPME method with a modification in the influence function. Therefore, the error estimates of the P^{3}M method can be used for the SPME method as well. Later, Wang *et al.* [24] presented an error estimate for the approximation error, but their estimate was not simple enough to be used in practice.

The approximation error in the SE method is independent of the grid size and is only a function of the number of points in the support of suitably scaled and truncated Gaussians. Hence, if the grid is coarser, the Gaussians are made wider. As a result of this difference, a much smaller FFT grid size is used in the SE method compared to the SPME method. This is an important advantage of the SE method over the SPME method specially on parallel architectures. Parallel FFT algorithms require all-to-all communications between processes and the number of communications increases as the square number of processes. Therefore, there have been several attempts to replace FFTs by iterative solvers or to employ a Fast multipole method (FMM) [11]. We shall return to this discussion in section 9. On the other hand, the spreading/gathering steps are more expensive in the SE method compared to the SPME method, even though their costs are substantially reduced by employing the Fast Gaussian Gridding. In section 8.2 we will see how changes in parameters will influence the behavior of both methods.

To visualize the dependency of the SE and SPME methods to grid size, in figure 4 we plot the relative rms error in evaluation of the force as a function of total number of grid points in three directions, M^3, for both methods and with two different Ewald parameters. The figure shows the spectral accuracy in the computation of the force in the SE method and polynomial accuracy in the SPME method. Note that if a smaller value of P were chosen for the SE method, the error curve would follow the current curve down to a certain accuracy level, and then levels out.

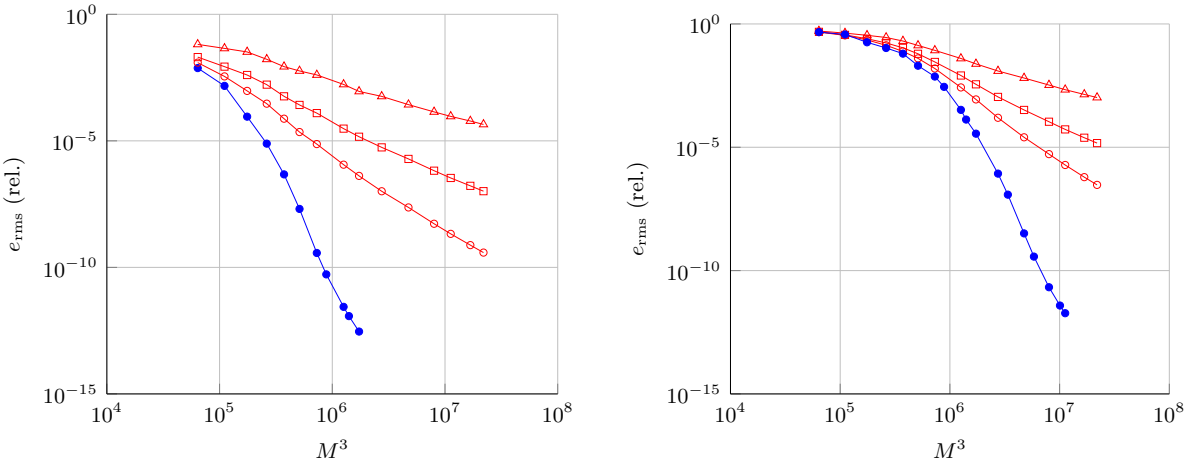

Figure 4: Relative rms error in evaluation of the force as a function of M^3 for the SE and SPME methods with (left) $\xi = 3$ and (right) $\xi = 6$. A uniform system of $N = 81\,000$ particles in a box of length $L = 9.3222$ is used. The red curves represent the SPME method with B-splines of order $p = 3, 5, 7$ from top to bottom and the blue curve represents the SE method with $P = 24$ points in the support of the Gaussians in each direction.

7. Runtime estimates

So far, we have explained the fundamentals of the SE method and its error. In this section, we estimate the runtime of different parts of the electrostatic calculations to later aid us in the choice of the optimal Ewald parameter for our implementation of the SE method. For this, we implement the algorithm in GROMACS ver. 5.1 [14]. GROMACS is one of the fastest Molecular Dynamics simulation packages available which serves many users. It is efficient, highly optimized for serial use as well as parallel using MPI and/or multi-threading. Three different

systems are examined; A uniformly distributed system with $N \in \{1200, 3000, 81\,000, 1\,029\,000\}$ particles, a cloud-wall system with $N \in \{1200, 1\,012\,500\}$ particles used in [3] and an isolated dense-cloud system with $N = 1200$ particles (see figure 5). The cloud-wall and isolated systems are generated artificially to allow for strong long range interactions.

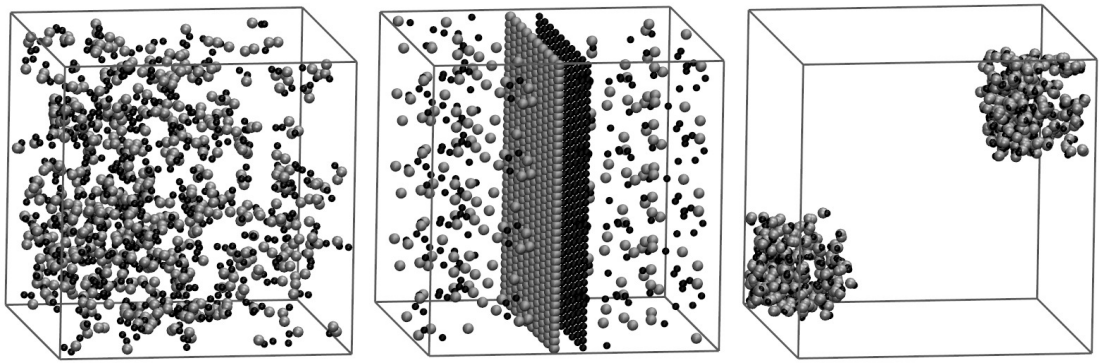

Figure 5: Three systems with $N = 1200$ particles. (Left) A homogeneous system of particles. (Middle) a cloud-wall system of particle (borrowed from [3]). (Right) an isolated dense-cloud system of particles.

Finding the optimal Ewald parameter is always a concern in Ewald methods. The optimal parameter balances the runtime of the real and Fourier space parts of the sum. This optimal value depends not only on the algorithms that are used but also on the implementation of the methods. Therefore, what we present here might require a modification with a different implementation of the SE method and also with the use of other MD packages. Here, we strive to estimate the runtime of different parts of the Electrostatic calculations. The results obtained here are useful to estimate the Ewald parameter while using the SE method. We start by generating a charge neutral system of $N = 3000$ uniformly distributed particles placed in a box of length $L = 3.1074$. To generate other uniformly distributed systems of particles, this system is scaled up such that the charge density remains constant, i.e. $N/L^3 \approx 3000/30 = 100$.

All runs (except those related to figure 17) are performed on a desktop machine with 8 GB of memory and an Intel Core i7-3770 3.40 GHz CPU with four cores and eight threads. GROMACS was built with the Gnu C Compiler version 4.8.3.

7.1. Timing the real space sum

Here we denote by n, the average number of nearest neighbors within a cut-off radius r_c around each particle. As the points are uniformly distributed, the maximum and minimum number of nearest neighbors are approximately equal and therefore we have

$$\mathsf{n} = \frac{\frac{4}{3}\pi r_c^3}{L^3} N.$$

The computation of the real space sum in GROMACS consists of two parts; neighbor searching and force calculation. The neighbor searching is the process of finding the neighboring particles which lie in a specified cut-off radius for each particle. In GROMACS, this step has a complexity of $\mathcal{O}(N)$ (though with a large prefactor of 125), cf. [23]. The experiment (performed in GROMACS) suggests the following for the runtime in the neighbor searching step,

$$t_{\text{ns}} = c_{\text{ns}} \mathsf{n} N,$$

where $c_{\text{ns}} \approx 1.4 \cdot 10^{-8}$.

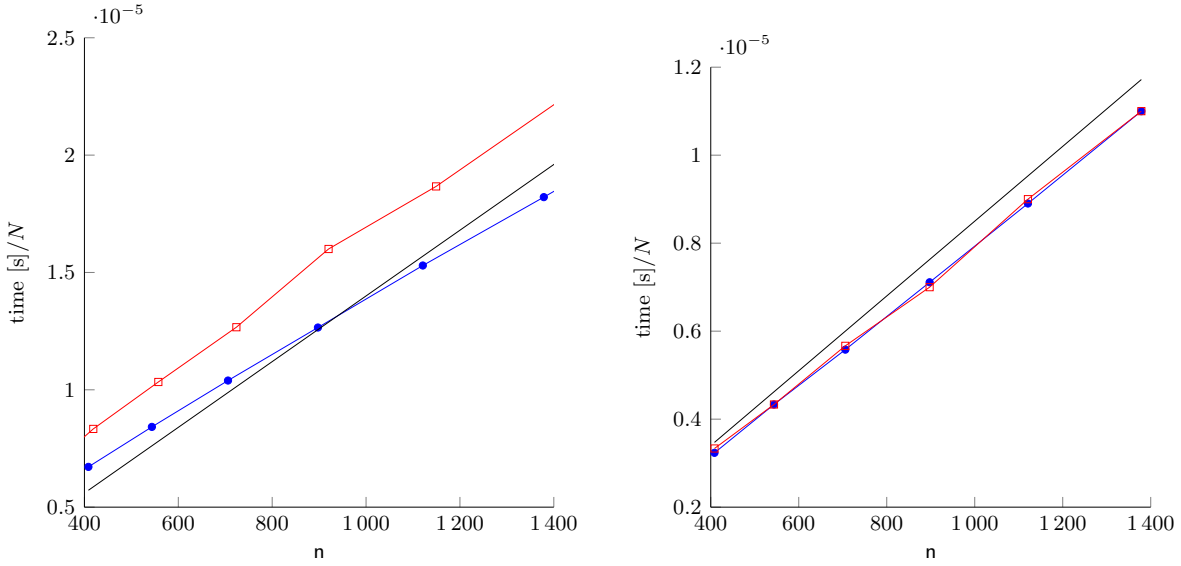

Figure 6: The experiment suggests $\mathcal{O}(N) = c\mathsf{n}N$ for (left) neighbor searching and (right) real space force calculation. The black line represents the estimated time. The red and blue dotted lines are the measured runtimes scaled with the number of particles $N = 3000$ and $N = 81\,000$ respectively.

The calculation of the so called short-range or real space part of the Ewald sum has $\mathcal{O}(N)$ complexity as well, and the experiment suggests that the runtime follows

$$t_{\text{force}} = c_{\text{force}}\mathsf{n}N,$$

where $c_{\text{force}} \approx 10^{-8}$. Figure 6 (left) shows the scaled runtime of the neighbor searching and figure 6 (right) shows the dependency of the scaled runtime of the real space sum calculation as a function of number of nearest neighbors. The computational time of the real space sum can then be estimated as

$$t_{\text{real}} = t_{\text{ns}} + t_{\text{force}} = c_{\text{real}}\mathsf{n}N, \tag{7.1}$$

with $c_{\text{real}} \approx 2.4 \cdot 10^{-8}$. Again, the constants c_{ns}, c_{force} and c_{real} are found by numerical experiments and depend on the implementation and the machine that is used.

We should also mention that the runtime of the real space sum presented here includes the runtime of other computations which do not actually belong to the real space sum but they cannot be excluded in GROMACS. Yet GROMACS internal parameters are chosen in a way that these calculations are small and fixed for all different runs.

7.2. Timing the Fourier space sum

The Fourier space computation consists of three different routines; FFT/IFFT, spreading/gathering and solving (scaling) routines. The computational costs of the FFT and IFFT routines are similar and are of the order $M^3 \log M^3$. Since GROMACS performs a real to complex FFT and a complex to real IFFT, we can write

$$t_{\text{fft-ifft}} = c_{\text{fft-ifft}} \frac{M^3}{2} \log(\frac{M^3}{2}),$$

with $c_{\text{fft-ifft}} \approx 4.5 \cdot 10^{-9}$.

Another step in the evaluation of the Fourier space sum using the SE method is the scaling step which has $\mathcal{O}(M^3)$ computational complexity. The runtime estimate in this step is in the form of

$$t_{\text{solve}} = c_{\text{solve}} M^3,$$

and $c_{\text{solve}} \approx 2 \cdot 10^{-9}$.

In the spreading (and gathering) step, the number of operations required to spread to (or interpolate from) the P^3 neighboring points of a source point are equal and scale as $\mathcal{O}(NP^3)$. For some P values, it is possible to accelerate the routines with the aid of single instruction, multiple data (SIMD) instructions. Employing SIMD intrinsics will decrease the computational time for those P values. Considering this timing complicates our analysis but for simplicity we suggest the following form

$$t_{\text{sp-ga}} = c_{\text{sp-ga}} N P^3,$$

and $c_{\text{sp-ga}} \approx 5 \cdot 10^{-9}$. Note that in the timing of the spreading and gathering steps, the runtime of the pre-computation step is also included. The pre-computation step includes the computation of the exponentials in the FGG routine. The total time spent in the Fourier space part of the sum can then be estimated as

$$\begin{aligned} t_{\text{fourier}} &= t_{\text{fft-ifft}} + t_{\text{sp-ga}} + t_{\text{solve}} \\ &= c_{\text{fft-ifft}} \frac{M^3}{2} \log(\frac{M^3}{2}) + c_{\text{sp-ga}} N P^3 + c_{\text{solve}} M^3, \end{aligned} \tag{7.2}$$

where $c_{\text{fft-ifft}} \approx 4.5 \cdot 10^{-9}$, $c_{\text{sp-ga}} \approx 2 \cdot 10^{-9}$, and $c_{\text{solve}} \approx 5 \cdot 10^{-9}$.

To assess reliability of the runtime estimates for the real and Fourier space sum presented above, a system of $N = 81\,000$ uniformly distributed charged particles in a box of length $L = 9.3222$ is considered. We measure the runtime in evaluating the force for different Ewald parameters and to achieve a relative rms error of $\approx 10^{-5}$. Figure 7 (left) shows that the estimates can predict actual runtimes rather well. It also suggests an optimal value of $\xi \approx 6.75$ (corresponding to $r_{\text{c}} = 0.60$).

We repeat the same experiment but now to achieve a relative rms error of $\approx 10^{-12}$. Figure 7 (right) suggests another optimal value, $\xi \approx 6.31$ (corresponding to $r_{\text{c}} = 0.9$). Comparing the plots in figure 7 we observe that for the SE method, the Ewald parameter has a weak dependency to the error tolerance. This is again in contrast with the SPME method, see figure 11 where the dependence is much stronger.

To convey the effect of system setting on the optimal Ewald parameter we find the optimal ξ for three different systems in figure 5 with the same number of particles and box size. We use systems with $N = 1200$ oppositely charged particles and boxes of length $L = 10$. Figure 8 shows that, as expected, the optimal ξ is only slightly different for each system. The density of the clouds in the isolated-cloud system is much higher than the density of other systems. Therefore, to keep the number of nearest neighbors fixed, the radius cut-off r_{c} is smaller for the isolated-cloud system and as a result, ξ is larger. On the other hand, the cloud-wall system is designed in the way that the number of particles located on the walls are 4 times more than the number of particles in each cloud. Due to the 2D structure of the walls, r_{c} has to be increased to have the same average number of near neighbors. This gives rise to a smaller Ewald parameter. Figure 8 shows the optimal Ewald parameter obtained for three different cases.

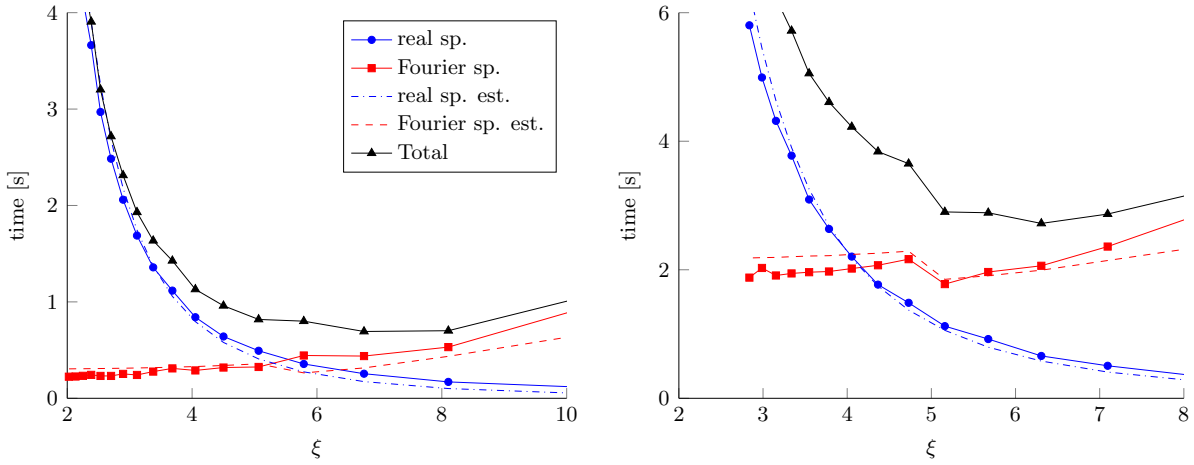

Figure 7: Real space, Fourier space, and total runtime as a function ξ for a system of $N = 81\,000$ uniformly distributed particles in a box of length $L = 9.3222$. (Left) To achieve a relative rms error of $\approx 10^{-5}$. $M = 50, \ldots, 250$ and $P \in \{8, 10\}$. The optimal Ewald parameter is $\xi \approx 6.75$ which corresponds to $r_c = 0.60$. (Right) To achieve a relative rms error of $\approx 10^{-12}$, $M = 100, \ldots, 280$ and $P \in \{20, 22\}$. The optimal Ewald parameter is $\xi \approx 6.31$ which corresponds to $r_c = 0.90$.

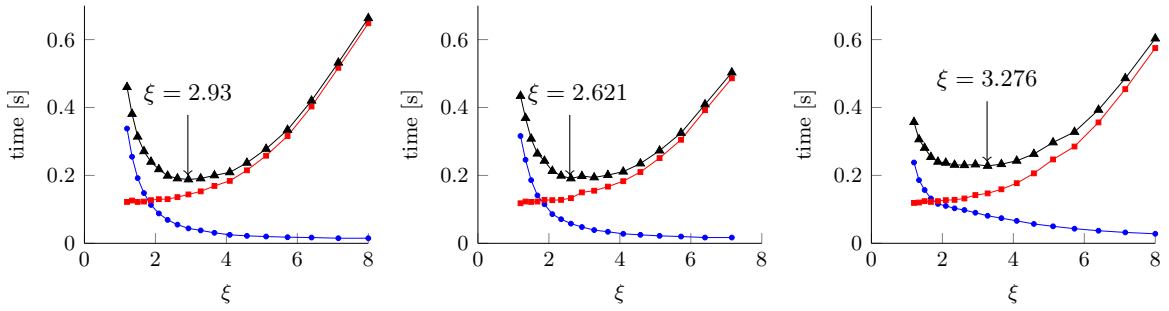

Figure 8: Real space (●), Fourier space (■) and total runtime (▲) as a function of ξ (left) for the uniform system, (middle) cloud-wall system and (right) two isolated-cloud system of particles to achieve a relative rms error of $\approx 10^{-8}$. $N = 1200$, $L = 10$, $P = 18$, $M = 40, \ldots, 220$.

8. Comparison of the SE and SPME methods - serial implementation

There are many different FFT-based methods available to calculate the long-range electrostatics in periodic systems of particles. However, one of the most efficient methods available is still the SPME method. Moreover, one of the fastest and highly optimized implementations of the SPME method is available in GROMACS. Considering the man-hour time spend in optimizing the method, this is a difficult challenge for the SE method. In this section, we consider the serial implementation of the SE and SPME methods. Also, due to the fact that in our comparisons, the real space part of the sum is the same in the computation of the electrostatic interactions, from now on, we ignore the runtime of the real space part and only present results of the comparison in the Fourier space part of the Ewald sum.

In sections 2.1 and 5.1 the estimates are presented for the absolute errors though, in this and the following sections we measure relative rms errors. To obtain estimates for relative errors one can, however, use the magnitude of the potential and force provided in section 5.1.

8.1. SE-SPME stepwise comparison

We use a uniform system of $N = 81000$ particles located in a box of length $L = 9.3222$. This system is an equilibrated system of 1000 water molecules in a box of length 31.074 Å $= 3.1074\ nm$ with charges -0.8476 and 0.4238 for Oxygen and Hydrogen atoms respectively. Figure 9 shows the runtime in evaluating the force as a function of relative rms error with $P = 16$ in the SE method and $p = 3, 5$ in the SPME method. The figure includes only the runtime of the FFT and *solve* steps and not the spreading/gathering step. Also we present only the error committed in the Fourier space part of the sum.

FFT/IFFT

Since both methods are implemented in GROMACS, the same FFT routine is used for both methods. Figure 9 (left) represents the total time spent in the FFT and IFFT routines. Since the SE method needs smaller mesh size compared to the SPME method, the runtime difference to achieve a certain level of accuracy is large. Note that the grid size is chosen to suit the FFT routines and to minimize runtime fluctuations.

Solve

In this context, *solve* refers to the scaling of the Fourier transformed distributed charges on a uniform mesh. Again, the only difference in runtime of both methods in this step is due to the difference in the grid size. Referring to section 7, this step has a complexity of $\mathcal{O}(M^3)$ and, therefore, the scaling step is almost on the same order of computational complexity as the FFT/IFFT step.

Spreading and Gathering

Theoretically in the SPME method, it is possible to achieve any desired accuracy with all B-spline orders (even though it might not be practically feasible) but due to the low accuracy demand in MD, B-splines of order 3 are frequently used. In the SE method, on the other hand, increasing the grid size while keeping P fixed, reduces the error up to a certain level of accuracy. When the approximation error dominates truncation errors, to decrease the total error, P has to be increased.

The runtime for the SE method is dominated by the spreading and gathering steps. Hence, when the total computational time is considered, the comparison of the two methods is much more complex, as will be discussed next.

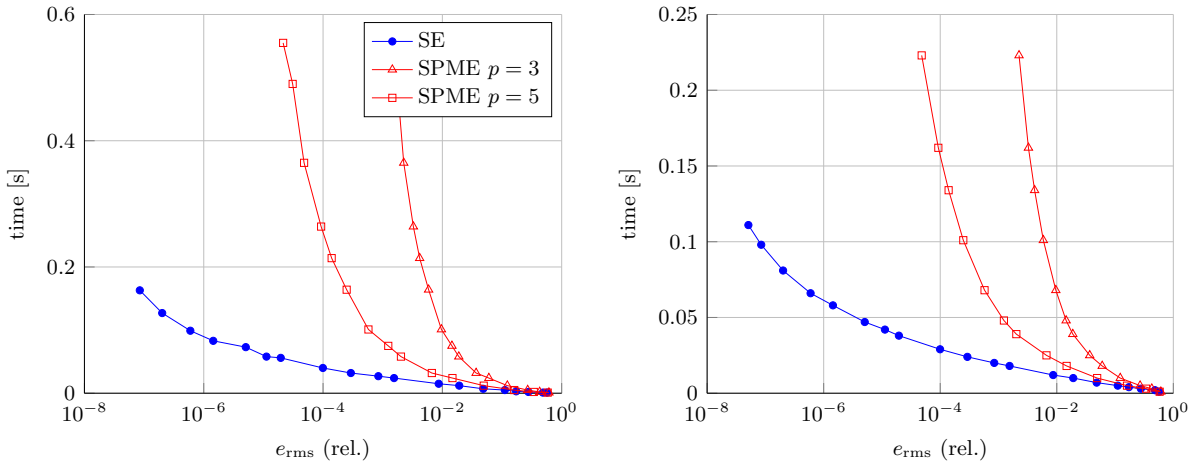

Figure 9: Runtime vs. relative rms error of (left) the FFT and (right) *solve* routines. A uniform system with $N = 81000$, $L = 9.3222$ and $\xi = 6.5$ is used. In the SE method, $P = 16$ and $M = 40, \ldots, 192$ and in the SPME method $p = 3$ and $p = 5$ (from top to bottom) and $M = 40, \ldots, 350$. The red curves show the runtime of the SPME method and the blue curves represent the runtime of the SE method.

8.2. SE-SPME runtime comparison

Our aim here is to examine sensitivity and behavior of both methods for different Ewald parameters. We use a system of $N = 81\,000$ uniformly distributed particles in a box of length $L = 9.3222$ and plot Fourier space runtime as a function of relative rms error, see figure 10. For the SE method, P is obtained using the approximation error bounds calculated in section 5.3. For the SPME method we use $p = 3$ and 5. We observe that the runtime of the SE method is less sensitive to the Ewald parameter than the SPME method and therefore the optimal Ewald parameter can be found with less effort. Also, we note that for the case of $\xi = 6.5$, error tolerance of 10^{-5} cannot be achieved with $p = 3$ in a reasonable wall-clock time.

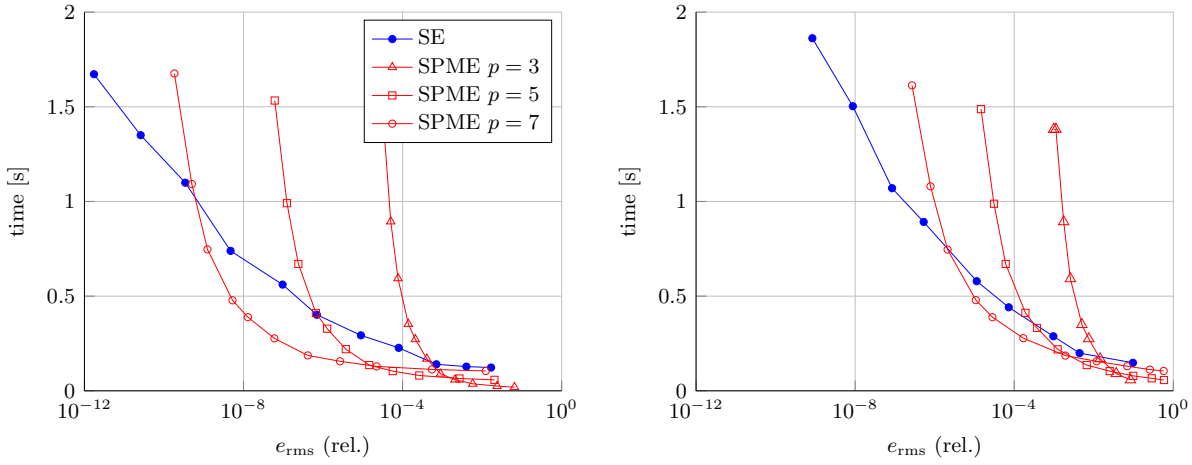

Figure 10: Fourier space runtime comparison of the SE and SPME methods using a single CPU core. (left) $\xi = 3$, $P = 8, \ldots, 18$ and $M = 40, \ldots, 112$ in the SE method and $M = 40, \ldots, 400$ in the SPME method.(right) $\xi = 6.5$, $P = 8, \ldots, 18$ and $M = 80, \ldots, 240$ in the SE method and $M = 40, \ldots, 400$ in the SPME method. A uniform system with $N = 81\,000$ particles and $L = 9.3222$ is used.

In the next experiment, in figure 11 we compare Fourier space runtime as a function of ξ for a relative rms error of $\approx 10^{-5}$. For the SPME method, B-splines of order 3, 5 and 7 were

used. For the SE method, as before, we obtain P from the approximation error bound. The figure confirms that for the SPME method to be efficient, ξ has to be kept small enough or p has to be increased otherwise. For example, increasing ξ from 3 to 6.5, the runtime of the SPME method with $p = 7$ increases with a factor of 5 while this is less than a factor of 2 for the SE method. The cost of spreading/gathering and pre-computation steps depend only on P (since N is fixed). Also P is uniquely determined by the error tolerance and, therefore, is constant for all ξ in this test. Increasing ξ, however, increases the grid size and consequently the cost of the FFT and *solve* steps.

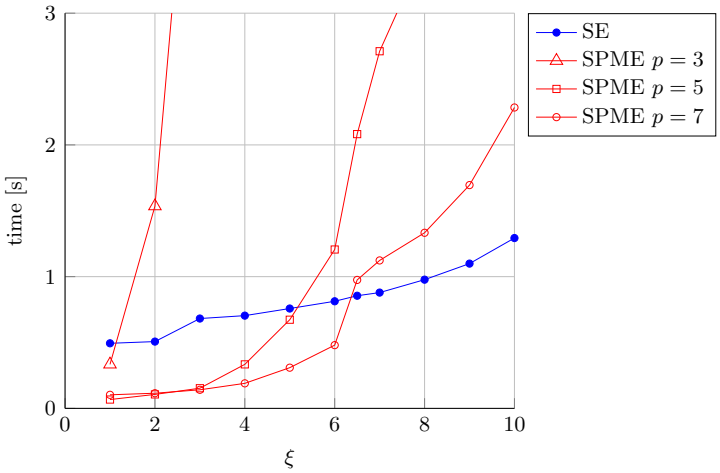

Figure 11: Fourier space runtime comparison of the SE and SPME methods using a single CPU core to achieve a relative rms error of $\approx 10^{-5}$. The red lines are the runtime of the SPME method with B-splines of orders $p = 3, 5, 7$. A uniform system with $N = 81000$ particles in a box of length $L = 9.3222$ is used.

9. Comparison of the SE and SPME methods - parallel implementation

On parallel computers or while using heterogeneous architectures, a larger grid size in the 3D FFT algorithm will increase the cost of communications between processes. This cost can degrade the performance of the FFT-based methods and specially those with larger FFT grid requirement. Unlike FFTs, spreading/gathering and solve steps benefit more from parallel processing. We use the GROMACS parallel paradigm [11] and implement a parallel version of the SE method. As has already been explained, two 3D FFTs are needed in the implementation of the SE and SPME methods. The parallel implementation of the 3D FFT in GROMACS requires also two redistributions (or four, considering also the 3D IFFT) of 3D grids that is clearly expensive. In figure 12 we use a uniform system of $N = 81\,000$ particles with $L = 9.3222$ and plot Fourier space runtime of the SE method with $P = 8, \ldots, 18$ and SPME method with $p = 3, 5, 7$ for $\xi = 3$ (left) and $\xi = 6.5$ (right) using 4 CPU cores. Comparing figure 10 and 12 discloses that for $\xi = 6.5$, the intersection of the runtime curve for the SE and SPME ($p = 7$) methods changes from an error level around 10^{-6} in the serial implementation to 10^{-5} in the parallel implementation with 4 CPU cores.

Next, we generate a uniform system of $N = 1\,029\,000$ particles within a box of size $L = 22.2698$. This system is generated by replicating a uniform system of $N = 3000$ particles in a box of length $L = 9.3222$. As in the previous section and similar to the case with $N = 81\,000$ uniformly distributed system of particles, we choose $\xi = 6.5$. If we also consider the real

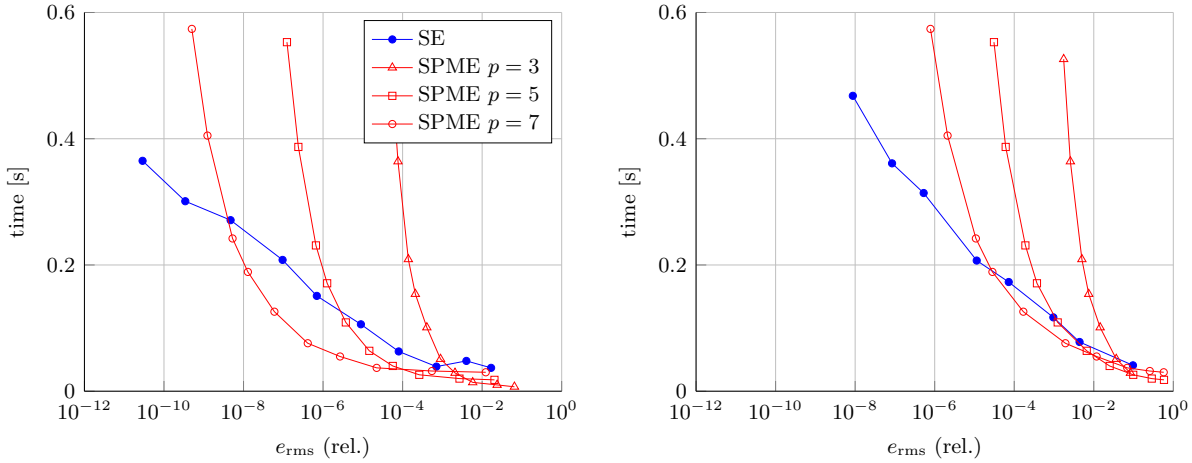

Figure 12: Fourier space runtime comparison of the SPME method of order $p = 3, 5, 7$ and the SE method using 4 CPU cores. (left) $\xi = 3$, $P = 8, \ldots, 18$ and $M = 40, \ldots, 112$ in the SE method and $M = 40, \ldots, 400$ in the SPME method. (right) $\xi = 6.5$, $P = 8, \ldots, 18$ and $M = 80, \ldots, 240$ in the SE method and $M = 40, \ldots, 400$ in the SPME method. A uniform system with $N = 81\,000$ particles and $L = 9.3222$ is used.

space sum cost, runtime estimates in section 7 can be used to obtain an almost optimal Ewald parameter, $\xi = 4.184$. The result using 8 CPU threads is shown in figure 13. We observe that the SE method is still competitive for low error tolerances. In fact, the cross overs are decreased and occur at $\approx 10^{-5}$ for $\xi = 6.5$ and at $\approx 10^{-6}$ for $\xi = 4.184$ instead. In the right plot of figure 13, due to insufficient memory, the simulation can not be performed for $M > 540$ (red curves) on our desktop machine.

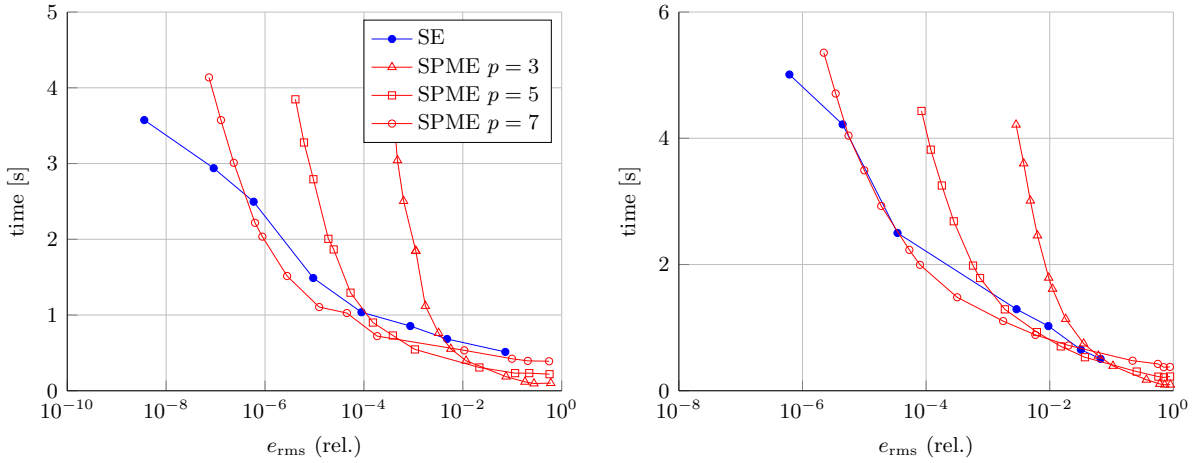

Figure 13: Fourier space runtime comparison of the SE and SPME methods using 8 CPU threads. A uniform system of $N = 1\,029\,000$ particles in a box of length $L = 22.2698$ is used. For the SE method $P = 6, ...18$ and $M = 220, \ldots, 360$, and for the SPME method, $p = 3, 5, 7$ and $M = 50, \ldots, 540$. (left) $\xi = 4.184$, (right) $\xi = 6.5$.

Figure 14 shows a similar experiment using a cloud-wall system of $N = 1\,012\,500$ particles in a box of length $L = 150$. As in [3], we set $\xi = 0.7002$ for the left and $\xi = 0.8267$ for the right plot of figure 14. We should note that based on the Kolafa & Perram truncation error estimates described in section 2.1, Fourier space truncation error is controlled by ξL and not only ξ. Hence, the Ewald parameters are one order of magnitude smaller than the parameters

in the previous simulations. Based on these plots we can conclude that, for systems with strong long-range interactions, the SE method can be competitive even for error tolerances below 10^{-5}.

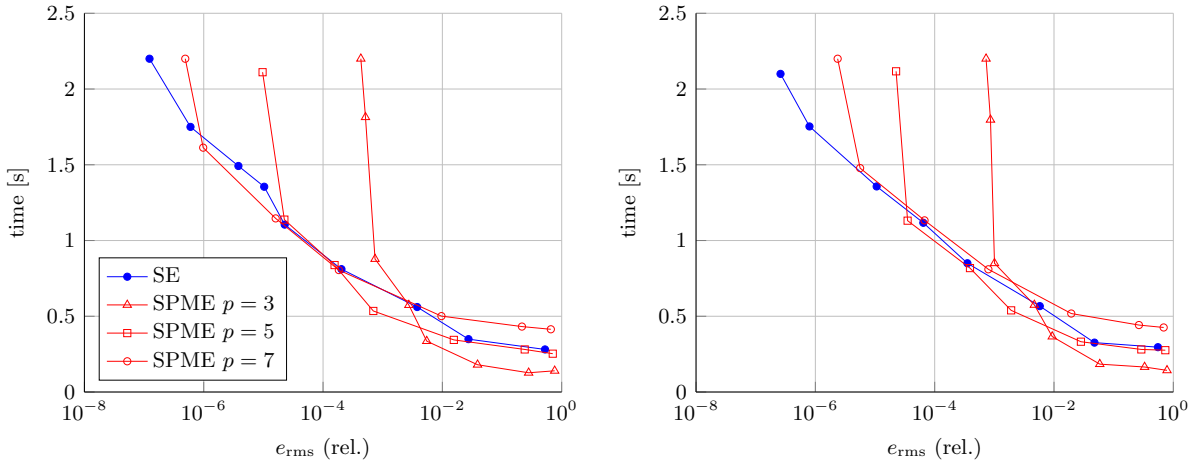

Figure 14: Fourier space runtime comparison of the SE and SPME methods using 8 CPU threads. A cloud-wall system with $N = 1\,012\,500$ particles in a box of length $L = 150$ is used. For the SE method $P = 6, \ldots, 16$ and for the SPME method $p = 3, 5, 7$. Also $M = 32, \ldots, 448$ for both methods. (left) $\xi = 0.7002$, (right) $\xi = 0.8267$.

10. Further optimization

GROMACS can be accelerated with SIMD extensions, e.g. AVX & AVX2. In the SE method implementation, the spreading/gathering routines for $P = 8, 16$ and 24 can be accelerated using AVX and AVX2 instructions if the necessary conditions are met. Our desktop processor supports AVX instructions, but for a benchmark with AVX2, we will perform an experiment on Beskow supercomputer.

We start by a cloud-wall system of $N = 1\,012\,500$ particles with box size $L = 150$ and set $\xi = 0.8267$. First, we choose $P = 8, 16$ and plot Fourier space runtime for the SE method as a function of total grid size M^3. Figure 15 (left) shows that using AVX instructions, the total runtime is smaller and this reduction is more pronounced when $P = 16$.

So far we have performed all experiments in double precision, but as a realistic comparison for low accuracy demands, single precision has to be used. We consider the same system as in the previous example with the same Ewald parameter and set $P = 8$. The routine for $P = 8$ can be accelerated using AVX and SSE instructions. Figure 15 (right) shows a runtime comparison of the SE method in single and double precision as a function of the total grid size M^3. We clearly gain a factor of 2 using single precision, however, the difference between using AVX and SSE is minor in this case.

Now, we can compare the runtime of both SE and SPME methods using AVX and in single precision. We choose the same system and Ewald parameter as before and use $p = 3$ and 5 for the SPME method and $P = 8$ for the SE method. Figure 16 shows the relative rms error as a function of runtime. Note that, in contrast to the right plot in figure 14, the result here is obtained with a fixed P. Therefore, the blue curve (SE method) levels out for large M and to achieve a higher accuracy, P has to be increased.

To complete our benchmarks and investigate the performance, we compare both methods on Beskow supercomputer at KTH Royal Institute of Technology, Sweden. Beskow is a Cray XC40

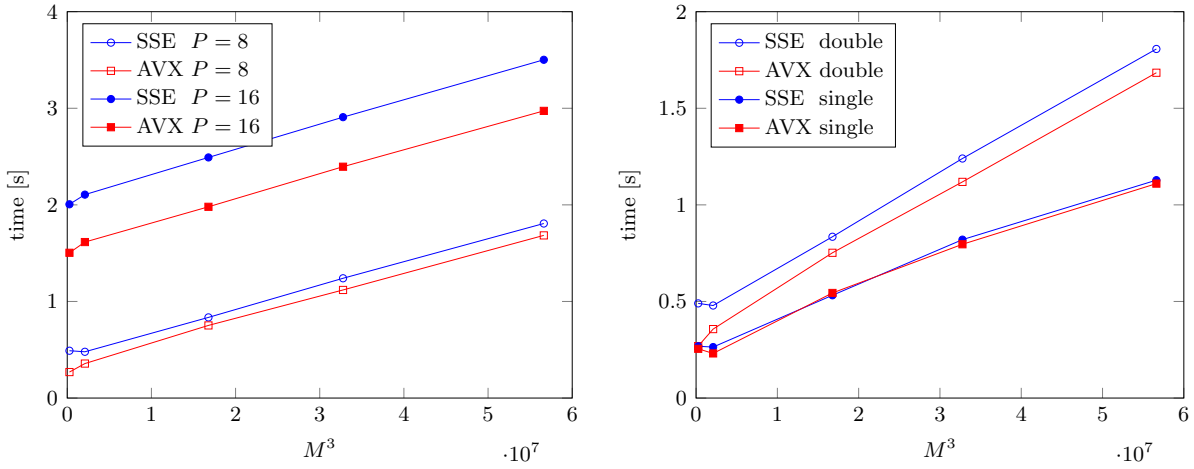

Figure 15: Fourier space runtime of the SE method as a function of the total grid size M^3. A cloud-wall system of $N = 1\,012\,500$ particles in a box of length $L = 150$ with $\xi = 0.8267$ is used. Also $M = 64, \ldots, 400$. (left) SSE-AVX comparison with $P = 8$ and 16 in double precision. (right) Single and double precision comparison using 8 CPU threads with $P = 8$.

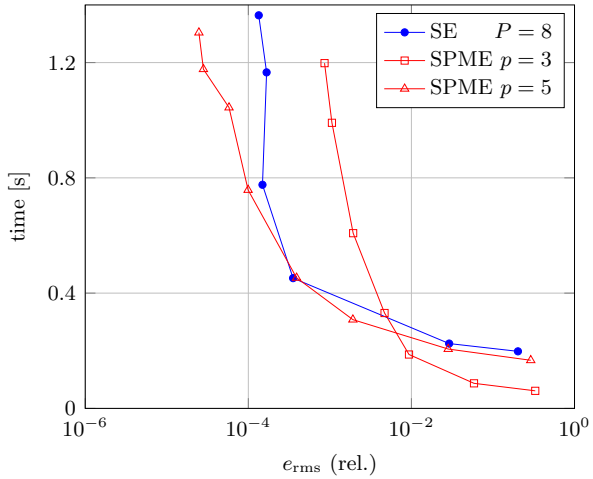

Figure 16: Fourier space runtime comparison of the SE ($P = 8$) and SPME ($p = 3, 5$) methods with AVX intrinsics enabled in single precision using 8 CPU threads. A cloud-wall system of $N = 1\,012\,500$ particles in a box of length $L = 150$ with $\xi = 0.8267$ is used. Also $M = 64, \ldots, 400$.

system, based on Intel Haswell processors and consists of 53632 compute cores. Haswell processors support AVX2 as well as *Fused Multiply Add* (FMA) instructions. This is more favorable for the SE method since there are more multiply-add operations in the spreading/gathering step compared to the SPME method.

We use a uniform system of $N = 1\,029\,000$ particles with $L = 22.2698$ and set $\xi = 6.5$. We, furthermore, set an error tolerance of 10^{-5}. To achieve a relative rms error less than the set tolerance, we choose $P = 16$ and $M = 320$ for the SE method and $p = 7$ and $M = 500$ for the SPME method. Note that with $P = 10$, the same error tolerance can be reached, but the performance is degraded.

GROMACS is also compiled and accelerated with AVX2 SIMD instructions. Figure 17 shows the runtime comparison of both methods for up to 256 MPI ranks. Even though the SE method works only marginally better than the SPME method this difference is more significant if the time iteration is to be considered. In the right plot of figure 17, we break down the cost of FFT and spreading/gathering steps for both methods. It is worth mentioning that, for a fixed grid size, using a very large number of cores can reduce the performance since the cost of communications in the FFTs might dominate the computational time. This is more evident in the SE method (blue circle-dotted line) as the grid is smaller.

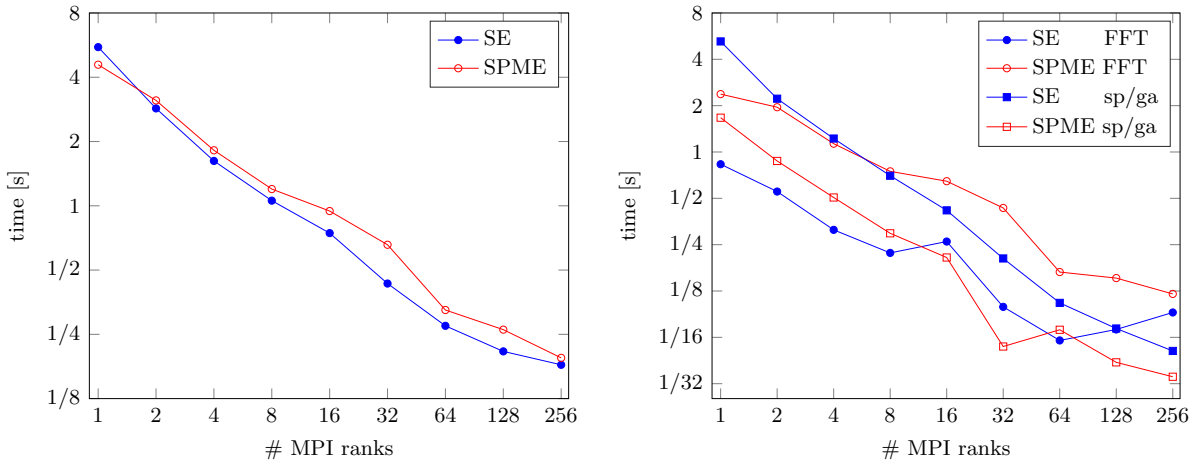

Figure 17: (left) Runtime and (right) breakdown runtime of the SE and SPME methods in computing the electrostatic force to achieve a relative rms error of $\approx 10^{-5}$. Computation is done on Beskow supercomputer in single precision and with up to 256 MPI ranks. Also $P = 16$, $M = 320$ for the SE method and $p = 7$, $M = 500$ for the SPME method and $\xi = 6.5$. A uniform system of $N = 1\,029\,000$ particles with $L = 22.2698$ is used.

11. Conclusions

In this work we compared our implementation of the SE method introduced by Lindbo & Tornberg in [15] and a highly optimized version of the SPME method [7] in GROMACS [14] version 5.1. The two methods belong to the family of Particle Mesh Ewald methods, and are similar in their structure. However, by the different choice of window functions - suitably scaled and truncated Gaussians vs pth order B-splines ($p = 3, 5, 7$), significant differences do arise.

Approximation errors in the SE method are essentially independent of the FFT grid size and can be controlled by increasing the number of points of support in the Gaussian while simultaneously rescaling it. Whereas in the SPME method approximation errors decay algebraically as h^p as the grid size h is decreased. This means that smaller FFT grids can be used in the

SE method as compared to the SPME method. On the other hand, the gridding and spreading steps that use the window functions for interpolation, are more costly for the SE method.

Both methods can use the same routine for the evaluation of the real space sum, so the comparison really entails the evaluation of the k-space sum. However, it is desirable to balance the work done for the two parts, and hence the implementation for the real space sum will effect the choice of the Ewald parameter ξ. As we have seen in the results section (figure 11), the k-space runtime for the SPME method is much more sensitive to the choice of ξ than the SE method.

The comparison between the two methods will depend on the specific problem, the value of ξ, the accuracy requirement, and also on the implementation. Generally speaking, the SE method will be most efficient if high accuracy is required, whereas the SPME method will be most efficient for low accuracy demands. The choice of $p = 3$ for the B-splines is essentially only competitive with the other choices if a very low accuracy is sufficient, and it also shows the highest sensitivity in the choice of ξ, which makes it less attractive to work with.

Our benchmarks include system of particles with periodic boundary conditions and different number of particles, box sizes and topologies. We ran GROMACS on a regular desktop machine with an Intel Core i7-3770 CPU as well as up to 256 MPI ranks on a supercomputers with AVX2 support.

Both algorithms were accelerated using SIMD instructions in single and double precision. The spreading/interpolating steps can be done by either computing the effect of all particles on a specific grid point or by computing all the grid values for one particle at the time. In our implementation, we choose the latter case. From the computer science point of view, this case has the benefit of writing on a consecutive memory locations. It is also advantageous for the distributed memory architectures though it should be revised again while using Graphics processing Units (GPU).

Our benchmarks show that the typical relative rms error below which the SE method becomes more efficient than the SPME method is around 10^{-6}. We however observed that for some cases where the system is large enough and long range interactions are strong, it can be competitive for accuracies around 10^{-5}, and whenever SIMD instructions are used in single precision, both methods are comparable even for low accuracy demands. On the Beskow computer, a runtime comparison for a system of around 10^6 particles at a relative error level of 10^{-5} showed the SE method to be somewhat faster than the SPME method already from 2 MPI ranks, scaling up to 256.

Already at this level of optimization, we can see that the SE method can be competitive with the SPME method, and we expect to be able to make further enhancements and accelerations of the method. As of yet, the implementation of the SE method is not available as an integrated part of GROMACS, but is publicly available at [1].

Acknowledgement

This work has been supported by the the Swedish Research Council under grant no. 2011-3178 and by the Swedish e-Science Research Center. The authors gratefully acknowledge this support.

References

[1] Spectral Ewald method in GROMACS URL `https://github.com/davoudss/gromacs`.

[2] Amber. - molecular dynamics package. URL `http://ambermd.org`.

[3] A. Arnold, F. Fahrenberger, C. Holm, O. Lenz, M. Bolten, H. Dachsel, R. Halver, I. Kabadshow, F. Gähler, F. Heber, J. Iseringhausen, M. Hofmann, M. Pippig, D. Potts, and G. Sutmann. Comparison of scalable fast methods for long-range interactions. *Phys. Rev. E - Stat. Nonlinear, Soft Matter Phys.*, 88(6), 2013.

[4] V. Ballenegger, J. J. Cerdá, and C. Holm. How to Convert SPME to P3M: Influence Functions and Error Estimates. *J. Chem. Theory Comput.*, 8(3):936–947, 2012.

[5] T. Darden, D. York, and L. Pedersen. Particle mesh Ewald: An Nlog(N) method for Ewald sums in large systems. *J. Chem. Phys.*, 98(12):10089, 1993.

[6] M. Deserno and C. Holm. How to mesh up Ewald sums . I . A theoretical and numerical comparison of various particle mesh routines How to mesh up Ewald sums . I . A theoretical and numerical comparison of various particle mesh routines. *J. Chem. Phys.*, 7678(1998): 7678–7693, 2010.

[7] U. Essmann, L. Perera, M. L. Berkowitz, T. Darden, H. Lee, and L. G. Pedersen. A smooth particle mesh Ewald method. *J Chem Phys*, 103(1995):8577–8593, 1995.

[8] P. P. Ewald. Die berechnung optischer und eletrostatischer gutterpotentiale. *Ann. Phys.*, 64:253–287, 1921.

[9] L. Greengard and J.-Y. Lee. Accelerating the Nonuniform Fast Fourier Transform. *SIAM Rev.*, 46(3):443–454, 2004.

[10] L. Greengard, V. Rokhlin, L. G. Rokhlin, and V. A fast algorithm for particle simulations. *J. Comput. Phys.*, 73(2):325–348, 1987.

[11] B. Hess, C. Kutzner, D. Van Der Spoel, and E. Lindahl. GROMACS 4: Algorithms for highly efficient, load-balanced, and scalable molecular simulation. *J. Chem. Theory Comput.*, 4(3):435–447, 2008.

[12] R. W. Hockney and J. W. Eastwood. *Computer simulation using particles*. McGraw- Hill, New York, 1981.

[13] J. Kolafa and J. W. Perram. Cutoff Errors in the Ewald Summation Formulae for Point Charge Systems. *Mol. Simul.*, 9(5):351–368, 1992.

[14] E. Lindahl, B. Hess, and D. van der Spoel. GROMACS 3.0: a package for molecular simulation and trajectory analysis. *J. Mol. Mod.*, 7:306–317, 2001.

[15] D. Lindbo and A.-K. Tornberg. Spectral accuracy in fast Ewald-based methods for particle simulations. *J. Comput. Phys.*, 230(24):8744–8761, 2011.

[16] A. Neelov and C. Holm. Interlaced P3M algorithm with analytical and ik-differentiation. *J. Chem. Phys.*, 132(23), 2010.

[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. *Mol. Phys.*, 65(4):875–893, 1988.

[18] J. C. Phillips, R. Braun, W. Wang, J. Gumbart, E. Tajkhorshid, E. Villa, C. Chipot, R. D. Skeel, L. Kalé, and K. Schulten. Scalable molecular dynamics with NAMD. *J. Comput. Chem.*, 26(16):1781–1802, 2005.

[19] M. Pippig and D. Potts. Particle simulation based on nonequispaced fast Fourier transforms. *Fast Methods for Long-Range Interactions in Complex Systems, Forschungszentrum Jülich, Jülich*, IAS Series 6, G. Sutmann, P. Gibbon, and T. Lippert, eds.:131–158, 2011.

[20] SCAFACOS. - Scalable Fast Coulomb Solvers. URL `http://www.scafacos.de`.

[21] E. R. Smith. Electrostatic potentials in systems periodic in one, two, and three dimensions. *J. Chem. Phys.*, 128(17), 2008.

[22] G. Sutmann and B. Steffen. A particle-particle particle-multigrid method for long-range interactions in molecular simulations. *Comput. Phys. Commun.*, 169(1-3):343–346, 2005.

[23] D. van der Spoel, E. Lindahl, B. Hess, and the GROMACS Development Team. GROMACS User Manual version 4.6.7, 2014. URL `www.gromacs.org`.

[24] H. Wang, P. Zhang, and C. Schütte. On the numerical accuracy of Ewald, smooth particle mesh ewald, and staggered mesh Ewald methods for correlated molecular systems. *J. Chem. Theory Comput.*, 8(9):3243–3256, 2012.