



# A new algorithm for electrostatic interactions in Monte Carlo simulations of charged particles



William Robert Saunders<sup>a</sup>, James Grant<sup>b</sup>, Eike Hermann Müller<sup>c,\*</sup>

<sup>a</sup> Department of Physics, University of Bath, Bath BA2 7AY, Bath, United Kingdom

<sup>b</sup> Computing Services, University of Bath, Bath BA2 7AY, Bath, United Kingdom

<sup>c</sup> Department of Mathematical Sciences, University of Bath, Bath BA2 7AY, Bath, United Kingdom

## ARTICLE INFO

### Article history:

Available online 4 January 2021

### Keywords:

Monte Carlo  
Electrostatics  
Particle simulations  
Computational complexity  
Fast multipole method

## ABSTRACT

To minimise systematic errors in Monte Carlo simulations of charged particles, long range electrostatic interactions have to be calculated accurately and efficiently. Standard approaches, such as Ewald summation or the naive application of the classical Fast Multipole Method, result in a cost per Metropolis-Hastings step which grows in proportion to some positive power of the number of particles  $N$  in the system. This prohibitively large cost prevents accurate simulations of systems with a sizeable number of particles. Currently, large systems are often simulated by truncating the Coulomb potential which introduces uncontrollable systematic errors. In this paper we present a new multilevel method which reduces the computational complexity to  $\mathcal{O}(\log(N))$  per Metropolis-Hastings step, while maintaining errors which are comparable to direct Ewald summation. We show that compared to related previous work, our approach reduces the overall cost by better balancing time spent in the proposal- and acceptance- stages of each Metropolis-Hastings step. By simulating large systems with up to  $N = 10^5$  particles we demonstrate that our implementation is competitive with state-of-the-art MC packages and allows the simulation of very large systems of charged particles with accurate electrostatics.

© 2020 Elsevier Inc. All rights reserved.

## 1. Introduction

The accurate representation of all pairwise interactions in classical atomistic simulations is important to minimise systematic errors. In this paper we focus on Monte Carlo (MC) simulations of charged particles. Short-range interactions such as the Lennard-Jones potential can be safely truncated at a finite cutoff distance: when calculating energy differences in a proposed MC move only interactions with a fixed number of other particles in close proximity of the moving particle need to be taken into account. For fixed density the cost of one local Metropolis-Hastings (MH) step is constant, independent of the total number  $N$  of particles in the system. However, due to the long-range nature of the Coulomb potential, which decays in proportion to the inverse distance between two charges, including electrostatics is far from trivial since interactions with *all* other particles in the system have to be considered. Worse, interactions with periodic copies or mirror charges have to be taken into account if non-trivial boundary conditions are used.

As the review in [1] shows, a plethora of methods have been developed to address this issue, but care has to be taken to obtain reliable results. In this context the authors of [2] for example find that truncating the Coulomb potential in the

\* Corresponding author.

E-mail address: e.mueller@bath.ac.uk (E.H. Müller).

MC simulation of water in a highly anisotropic geometry leads to significant systematic errors. Other methods which have been proposed to avoid this problem include solving the Poisson equation with a grid-based multigrid method [3], Ewald summation [4] or the naive application of the Fast Multipole Method (FMM) [5–7]. Unfortunately all those approaches result in a prohibitively large computational cost as the number of particles  $N$  in the system grows, typically the cost per MH step increases as  $\mathcal{O}(N)$  or  $\mathcal{O}(\sqrt{N})$ . This renders the simulation of very large systems impossible and limits the predictive power of computer experiments.

This is particularly keenly felt in MC simulations, where in contrast to molecular dynamics, systems evolve through discontinuous moves of individual atoms or small numbers of particles. The current inability to treat electrostatics accurately and efficiently prohibits the application of MC to problems with larger particle counts. On the other hand, in many circumstances it is desirable to simulate a system with MC instead of molecular dynamics. In particular, the grand canonical or semi-grand ensembles where the number of particles can fluctuate are only accessible with MC [8,9]. Exchanging particle position in so-called ‘swap moves’ can be far more efficient at equilibrating systems e.g. when simulating the distribution of impurities near grain boundaries. The lack of efficient electrostatic algorithms also limits the use of methods such as phase- and lattice switch MC [10,11] which allow the accurate determination of free energy differences between phases. Enabling the routine application of MC to larger systems of charged particles will allow researchers to study the physics of a particular system with the computationally most appropriate evolution algorithm.

In this paper we show how the limitations of the MC method for charged particle systems can be overcome by constructing an algorithm which reduces the cost per MH step to  $\mathcal{O}(\log(N))$  without sacrificing accuracy, thereby making much larger simulations with realistic electrostatics feasible. Our multilevel approach is inspired by the Fast Multipole method and similar to the method recently proposed in [12]. However, compared to [12] it leads to an overall reduction of computational cost by balancing the time spent in proposing and accepting MC moves in realistic MC simulations.

The key observation motivating our new method is the following: standard FMM constructs local expansions for evaluating the long range interactions on the finest level of the hierarchical tree. While the evaluation of those expansions (and direct interactions with all close by neighbours) in the proposal stage of a MH step is  $\mathcal{O}(1)$ , re-calculating the local expansions incurs a cost of  $\mathcal{O}(N)$  since all local expansions are re-calculated in the upward and downward pass of the algorithm, resulting in an overall  $\mathcal{O}(N)$  cost per MH step. By storing the local expansions on each level of the multilevel hierarchy instead of accumulating them on the finest level, the relative cost of the proposal- and accept- stage can be balanced since each of those two steps requires a fixed number of operations on each level of the multilevel hierarchy. As the number of levels  $L$  is proportional to  $\log(N)$ , this results in a total computational complexity per MH step which grows logarithmically in the number of particles.

In summary, the new contributions of our work are as follows:

1. We describe a new hierarchical algorithm for Monte Carlo simulations with accurate electrostatics, which has an  $\mathcal{O}(\log(N))$  computational complexity per Metropolis-Hastings step.
2. By comparing to the similar method in [12] we show that our algorithm leads to an overall reduction in cost for realistic Monte Carlo simulations.
3. We describe the efficient implementation of our algorithm in the performance-portable PPMD framework [13] recently developed in our group. Since it has a user-friendly high-level Python interface, PPMD allows the easy implementation of Monte Carlo algorithms with accurate electrostatics.
4. We demonstrate that the runtime of our implementation is competitive with the state-of-the-art MC code DL\_MONTE [14,15], which struggles to simulate systems of the size that are easily accessible with our implementation.

For systems with  $N = 10^5$  particles our implementation is an order of magnitude faster than the DL\_MONTE code. At this system size, we observe that the alternative approach in [12] results in a 30% longer simulation time than our method. By fitting a semi-empirical model to predict the cost of a simulation as a function of the problem size, we show that asymptotically (i.e. for  $N \rightarrow \infty$ ) we expect our algorithm to be twice as fast as the one in [12].

*Structure* This paper is organised as follows: After discussing related work in Section 2, we review the native FMM algorithm and describe our new method in Section 3, where we also compare it to the approach in [12]. Following a description of the Python interface for our implementation of the algorithms introduced in this paper in Section 4, numerical results are presented in Section 5. We conclude and outline directions for further work in Section 6.

## 2. Related work

Methods for including untruncated electrostatic interactions in MC simulations with a computational complexity of  $\mathcal{O}(\log(N))$  per MH step have been developed previously in [12,16]. Compared to our approach, the method in [12] does not construct local expansions, thereby avoiding their recalculation whenever a particle is moved. While this might look like a reasonable simplification, it actually makes evaluating the change in potential for each proposed (but potentially rejected) MC transition more expensive. Since there is typically more than one proposed move per accepted transition, this renders the method in [12] more expensive overall, as our numerical experiments confirm.

A modification of the Barnes-Hut octree algorithm is discussed in [16]. Similar to FMM, the classical Barnes-Hut method constructs a hierarchical mesh structure, and represents the distribution of particles in cells on each level by their multidimensional Taylor expansion coefficients. While the calculation of the total electrostatic energy with the octree algorithm is  $\mathcal{O}(N \log(N))$ , the authors of [16] present a modification of the method which has a cost of  $\mathcal{O}(\log(N))$  per local MC step.

The  $\mathcal{O}(\log(N))$  algorithms presented here and in [12,16] improve on what can be achieved with Ewald summation [4,17], for which the change in electrostatic energy per MC proposal can be calculated at a computational complexity of  $\mathcal{O}(\sqrt{N})$ . This is because the overall  $\mathcal{O}(N^{3/2})$  cost of the Ewald-based energy calculation is made up by an iteration over all  $N$  particles and a sum over  $\mathcal{O}(\sqrt{N})$  reciprocal vectors (long-range contribution) and neighbouring particles (short-range contribution). If only  $\mathcal{O}(1)$  particles move in each proposed move, only a small number of the  $\mathcal{O}(\sqrt{N})$  sums have to be evaluated. A similar approach is currently explored in the DL\_MONTE code [14,15], though the implementation at present is limited by the fact that the short-range cutoff of the Ewald summation has to be identical to the cutoff of all other local interactions. In this paper we present numerical results which show that our new method can be used to simulate systems with up to  $10^5$  charges and accurate electrostatic interactions at a cost of around 1 ms per MH step.

For completeness, it should be pointed out that other methods with a computational complexity of  $\mathcal{O}(N)$  per MH step have also been developed. For example, in [3] a multigrid method is used to solve the Poisson equation to obtain the electrostatic potential generated by the particles. Since the global electrostatic field has to be re-computed for each (local) particle move and the cost of multigrid grows in proportion to the number of grid points [18], which in turn has to increase with the problem size to ensure an accurate representation of the charge distribution on the computational mesh, the cost is inherently  $\mathcal{O}(N)$ . An additional disadvantage of the method is that the mapping of the highly peaked charge distribution to a grid introduces interpolation errors.

While all methods discussed so far focus on the development of fast methods for computing the electrostatic interactions between particles based on Gauss' law and the fact that the electric field is the gradient of a scalar potential, a radically different approach is pursued in [19]. The authors of this paper argue that (statistically) integrating over a fictitious traverse electric field in addition to the different particle configurations produces the same results as a traditional Monte Carlo simulation with standard Coulombic interactions. While this enlarges the configuration space, it crucially only requires local updates for each particle move. Since the traverse electric field relaxes rapidly, this only introduces a small overhead, independent of the problem size. Overall the cost of updating all particle positions and the traverse field grows in proportion to  $N$ , which results in a  $\mathcal{O}(1)$  cost per individual particle move. A non-trivial interpolation scheme is used in [19] to reduce lattice artifacts that are introduced by enforcing Gauss' law on a Cartesian grid.

### 3. Method

We now discuss the new approach introduced in this work. After briefly reviewing key concepts of the classical FMM algorithm we describe our method in detail in Section 3.2 and compare it to the related technique in [12] in Section 3.3.

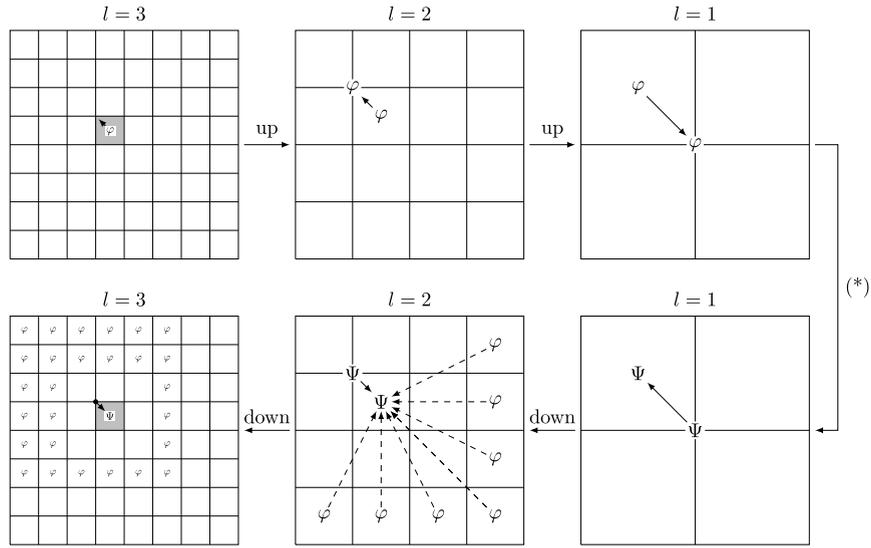
#### 3.1. Fast multipole method

In three dimensions the FMM algorithm [5–7] uses a hierarchical grid with  $L$  levels for the computational domain  $\Omega$  (which is assumed to be a cube of width  $a$ ) such that the number of cells on each level  $\ell$  is  $M_\ell = 8^{\ell-1}$  for  $\ell = 1, \dots, L$ . The number of cells on the finest level is  $M = M_L$ , and (to balance the work between the long range and direct field calculation) typically  $L$  is chosen such that there are  $\mathcal{O}(1)$  particles in each fine level cell. Each cell on level  $\ell = 1, \dots, L-1$  is subdivided into 8 child-cells on the next-finer level; conversely each cell on level  $\ell = 2, \dots, L$  has a unique parent cell. The Fast Multipole Algorithm now computes the electrostatic potential by splitting it into two contributions. First, the long range part is calculated by working out the multipole expansion of all charges in a fine level cell and transforming them into multipole expansions on the coarser levels in the upward pass of the algorithm. In the downward pass the multipole expansions on each level are transformed into local expansions around the centre of a cell. Those are then recursively combined into local expansions in the child cells. By only considering the contribution from multipole expansions in a fixed number of well-separated cells on each level, the potential of distant charges is resolved at the appropriate level of accuracy, while including the contribution from closer charges on finer levels. The method for calculating the long range contribution is shown schematically in Fig. 1; we refer the reader to [5–7] for further details. For the following discussion of our FMM variant for MC simulations the notion of an *interaction list* (IL) of a particular cell is crucial. For a cell  $\alpha$  on level  $\ell$  this interaction list  $IL(\alpha)$  is the set of cells which are the children of the parent cell of  $\alpha$  and its nearest neighbours, but which are well separated from  $\alpha$ , i.e. not direct neighbours of  $\alpha$  on level  $\ell$ . Explicitly, the interaction list is defined as

$$IL(\alpha) = \text{children}(\mathcal{N}_b(\text{parent}(\alpha))) \setminus (\alpha \cup \mathcal{N}_b(\alpha)),$$

where  $\mathcal{N}_b(\alpha)$  is the set of the 26 nearest neighbours of the cell  $\alpha$ ; the functions  $\text{children}(\alpha)$ ,  $\text{parent}(\alpha)$  return the set of child cells or the parent cell of a particular cell  $\alpha$ . An example of an interaction list can be found in the bottom left corner of Fig. 1: all cells labelled with the letter  $\varphi$  are in the interaction list of the grey cell labelled with a  $\Psi$ .

Finally, interactions with charges in neighbouring fine level cells are included by directly evaluating the  $1/r$  potential generated by those charges.



**Fig. 1.** Schematic illustration of the classical Fast Multipole Method in two dimensions first described in [5,6]. The number of levels is  $L = 4$  in this example. The upward pass for constructing the multipole expansions ( $\varphi$ ) is shown in the top row, while the local expansions ( $\Psi$ ) are built in the downward pass at the bottom of the figure. The asterisk (\*) on the right hand side denotes special operations on the coarsest level for incorporating (potentially nontrivial) boundary conditions.

### 3.2. FMM for Monte Carlo simulations

Now consider the following modification of FMM. Let  $\Psi_{\ell,\alpha}^\Delta$  be the  $p$ -term local expansion around the centre of cell  $\alpha$  on level  $\ell$  such that  $\Psi_{\ell,\alpha}^\Delta$  contains contributions from all charges in the interaction list  $IL(\alpha)$  of  $\alpha$ . Note that this is different from standard FMM, where the local expansions  $\Psi_{\ell,\alpha}$  contain the contribution of all charges not contained in the cell  $\alpha$  or its 26 nearest neighbours. However,  $\Psi_{\ell,\alpha}$  can be obtained by summing the  $\Psi_{\ell',\alpha_{\ell'}}^\Delta$  on the current and coarser levels, namely

$$\Psi_{\ell,\alpha} = \sum_{\ell'=1}^{\ell} \Psi_{\ell',\alpha_{\ell'}}^\Delta \quad \text{with } \alpha_\ell = \alpha \text{ and } \alpha_{\ell'} = \text{parent}(\alpha_{\ell'+1}) \text{ for all } \ell' = 1, \dots, \ell - 1. \quad (1)$$

For a cell  $\alpha$  on level  $\ell$  the local expansion  $\Psi_{\ell,\alpha}^\Delta$  can be expressed in terms of the coefficients  $(L_{\ell,\alpha}^\Delta)_n^m$  as

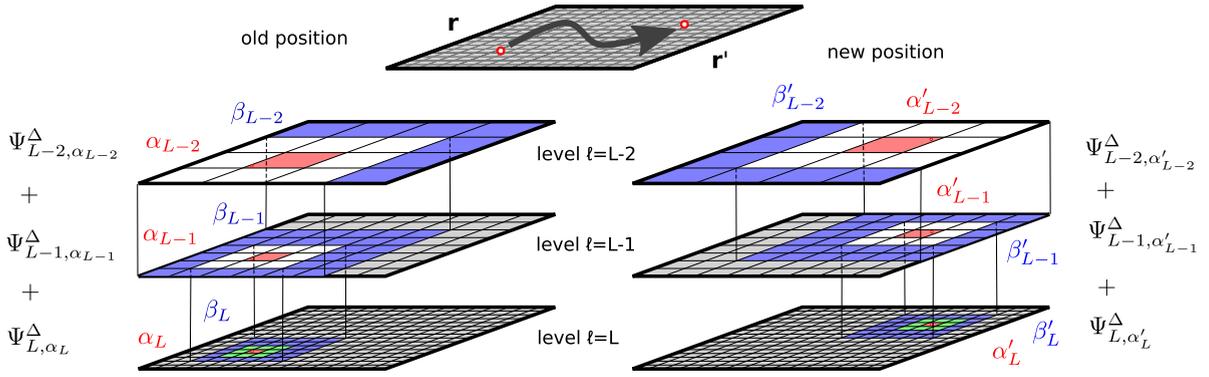
$$\Psi_{\ell,\alpha}^\Delta(\delta\mathbf{r}) = \sum_{n=0}^p \sum_{m=-n}^{+n} (L_{\ell,\alpha}^\Delta)_n^m (\delta r)^n Y_n^m(\delta\theta, \delta\phi) \quad \text{with } (\delta r, \delta\theta, \delta\phi) = \text{spherical}(\delta\mathbf{r}). \quad (2)$$

Here  $\delta\mathbf{r}$  is the position of the particle measured relative to the centre  $\mathbf{R}_\alpha$  of the cell  $\alpha$ . The function  $\text{spherical}(\mathbf{r})$  returns the spherical coordinates  $(r, \theta, \phi)$  of a vector  $\mathbf{r}$ . We further call the set  $\text{ancestors}(\alpha) = \{\alpha_{\ell-1}, \alpha_{\ell-2}, \dots, \alpha_2, \alpha_1\}$  defined recursively in Equation (1) the ancestors of cell  $\alpha$ . Our strategy for evaluating the long range contributions in Monte Carlo simulations is as follows (see Fig. 2):

**Initialisation.** At the beginning of the simulation, calculate the local expansion coefficients  $(L_{\ell,\alpha}^\Delta)_n^m$  for all cells  $\alpha$  and on all levels  $\ell$  by using a slightly modified variant of the upward/downward pass in the Fast Multipole Algorithm.

**Proposal.** Consider a proposed MC move  $\mathbf{r} \rightarrow \mathbf{r}'$  of charge  $q$  such that the original position  $\mathbf{r}$  is contained in the fine-level cell  $\alpha$  and the new position  $\mathbf{r}'$  in the fine-level cell  $\alpha'$  (which could be identical to  $\alpha$ ). To evaluate the change in the long range potential, evaluate and accumulate the  $\Psi_{\ell,\alpha_\ell}^\Delta(\mathbf{r} - \mathbf{R}_{\alpha_\ell})$  and  $\Psi_{\ell,\alpha'_\ell}^\Delta(\mathbf{r}' - \mathbf{R}_{\alpha'_\ell})$  on all levels  $\ell = 1, \dots, L$  of the hierarchy by using the sum in Equation (1), where  $\mathbf{R}_{\alpha_\ell}$  is the centre of cell  $\alpha$ . This gives the change in long-range electrostatic energy  $\Delta U_{\text{lr}} = q(\Psi_{L,\alpha'_\ell}^\Delta(\mathbf{r}' - \mathbf{R}_{\alpha'_\ell}) - \Psi_{L,\alpha_\ell}^\Delta(\mathbf{r} - \mathbf{R}_{\alpha_\ell}))$ . Secondly, add the change in short-range energy  $\Delta U_{\text{direct}}$  from a direct calculation of the interactions with particles in the same and adjacent cells. Finally, remove the spurious self-interaction contribution  $q^2/|\mathbf{r}' - \mathbf{r}|$  which is due to the potential field induced by the charge at the original position. This self-interaction correction is described in detail in [20, Section 3].

**Accept.** Assume we accept a move  $\mathbf{r} \rightarrow \mathbf{r}'$  of charge  $q$  such that  $\mathbf{r}$  lies in cell  $\alpha$  and  $\mathbf{r}'$  in cell  $\alpha'$ . For all cells  $\beta$  in the interaction list of  $\alpha$  and all its ancestors the contribution of a monopole with charge  $q$  is subtracted from  $\Psi_{\ell,\beta_\ell}^\Delta$  on all levels  $\ell = 1, \dots, L$ . This requires updating the expansion coefficients  $L_{\ell,\alpha_\ell}^\Delta$ . Conversely, a monopole of charge  $q$  is added to the local expansions  $\Psi_{\ell,\beta'_\ell}^\Delta$  of for all cells  $\beta'_\ell$  in the interaction list of  $\alpha'$  and its ancestors.



**Fig. 2.** Schematic illustration of the new method for a proposing a MC move  $\mathbf{r} \rightarrow \mathbf{r}'$ , as described in Section 3.2. On each level  $\ell$  the cell  $\alpha_\ell$  containing  $\mathbf{r}$  and the corresponding cell  $\alpha'_\ell$  containing  $\mathbf{r}'$  are marked in red. Cells  $\beta_\ell$  in the interaction list of  $\alpha_\ell$  and cells  $\beta'_\ell$  in the interaction list of  $\alpha'_\ell$  are shown in blue. Interactions with particles in green cells have to be evaluated directly when a move is proposed. (For interpretation of the colours in the figure(s), the reader is referred to the web version of this article.)

The *propose* and *accept* steps are written down explicitly in Algorithms L.1 and L.2; the direct calculation of the interaction with particles in the same fine-level cell or directly adjacent cells to obtain  $\Delta U_{\text{direct}}$  is given in Algorithm 4. In Algorithm 4 we remove the spurious self-interaction that occurs between the charge at the proposed position with itself at the original position.

**Algorithm L.1** Propose a move  $\mathbf{r} \rightarrow \mathbf{r}'$  using local expansions. Input: initial position  $\mathbf{r}$  of particle with charge  $q$ ; target position  $\mathbf{r}'$ . Output: change in total electrostatic energy  $\Delta U$ .

```

1: Find fine-level cells  $\alpha_L, \alpha'_L$  such that  $\mathbf{r} \in \alpha_L$  and  $\mathbf{r}' \in \alpha'_L$ 
2:  $\Delta U \leftarrow 0$ 
3: for  $\ell = L, L-1, \dots, 1$  do
4:   Update  $\Delta U \leftarrow \Delta U + q \left( \Psi_{\ell, \alpha'_\ell}^\Delta(\mathbf{r}' - \mathbf{R}_{\alpha'_\ell}) - \Psi_{\ell, \alpha_\ell}^\Delta(\mathbf{r} - \mathbf{R}_{\alpha_\ell}) \right)$  using local expansion in Equation (2)
5:   if  $\ell > 1$  then
6:     Set  $\alpha_{\ell-1} \leftarrow \text{parent}(\alpha_\ell)$ ;  $\alpha'_{\ell-1} \leftarrow \text{parent}(\alpha'_\ell)$ 
7:   end if
8: end for
9: Calculate change  $\Delta U_{\text{direct}}$  in electrostatic energy from direct interactions with Algorithm 4
10:  $\Delta U \leftarrow \Delta U + \Delta U_{\text{direct}}$ 

```

**Algorithm L.2** Accept a move  $\mathbf{r} \rightarrow \mathbf{r}'$  using local expansions. Input: old position  $\mathbf{r}$ , new position  $\mathbf{r}'$ .

```

1: Find fine-level cells  $\alpha_L, \alpha'_L$  such that  $\mathbf{r} \in \alpha_L$  and  $\mathbf{r}' \in \alpha'_L$ 
2: for  $\ell = L, L-1, \dots, 1$  do
3:   for  $\beta_\ell \in \text{IL}(\alpha_\ell)$  and  $\beta'_\ell \in \text{IL}(\alpha'_\ell)$  do
4:     Set  $(\delta r, \delta \theta, \delta \phi) \leftarrow \text{spherical}(\mathbf{r} - \mathbf{R}_{\beta_\ell})$  and  $(\delta r', \delta \theta', \delta \phi') \leftarrow \text{spherical}(\mathbf{r}' - \mathbf{R}_{\beta'_\ell})$ 
5:     for  $n = 0, \dots, p$  do
6:       for  $m = -n, \dots, +n$  do
7:         Update  $(L_{\ell, \beta_\ell}^\Delta)_n^m \leftarrow (L_{\ell, \beta_\ell}^\Delta)_n^m - q(\delta r)^{-(n+1)} Y_n^{-m}(\delta \theta, \delta \phi)$ 
8:         Update  $(L_{\ell, \beta'_\ell}^\Delta)_n^m \leftarrow (L_{\ell, \beta'_\ell}^\Delta)_n^m + q(\delta r')^{-(n+1)} Y_n^{-m}(\delta \theta', \delta \phi')$ 
9:       end for
10:    end for
11:   end for
12:   if  $\ell > 1$  then
13:     Set  $\alpha_{\ell-1} \leftarrow \text{parent}(\alpha_\ell)$ ;  $\alpha'_{\ell-1} \leftarrow \text{parent}(\alpha'_\ell)$ 
14:   end if
15: end for

```

Since the local expansions with  $\mathcal{O}(p^2)$  terms need to be evaluated in two cells per level, the cost of one proposal is  $\mathcal{O}(p^2 L) = \mathcal{O}(p^2 \log(N))$ . Similarly, when updating the  $\mathcal{O}(p^2)$  expansion coefficients  $L_{\ell, \alpha_\ell}^\Delta$  while accepting a move, the number of cells in the interaction list on each level is constant ( $6^d - 3^d = 189$  in  $d = 3$  dimensions, to be specific). Therefore the computational complexity of the accept stage is also  $\mathcal{O}(p^2 L) = \mathcal{O}(p^2 \log(N))$ . The larger constant (compared to the one in the propose stage) arises due to the fact that  $2 \times 189 = 378$  instead of 2 cells have to be considered on each level in this stage and is partially compensated by two effects:

1. Typically only a fraction of all proposed moves are accepted.

---

**Algorithm L.3** Initialise local expansion coefficients  $(L_{\ell,\alpha}^{\Delta})_n^m$  for electrostatic calculation with Algorithms L.1 and L.2.

---

```

1: for levels  $\ell = 1, \dots, L$  do
2:   for all cells  $\alpha_\ell$  on level  $\ell$  do
3:     for  $n = 0, \dots, p$  and  $m = -n \dots +n$  do
4:       Set  $(L_{\ell,\alpha_\ell}^{\Delta})_n^m = 0$ 
5:     end for
6:   end for
7:   for all cells  $\alpha_\ell$  on level  $\ell$  do
8:     for all particles with charge  $q_i$  and position  $\mathbf{r}_i \in \alpha_\ell$  do
9:       for all cells  $\beta_\ell \in \text{IL}(\alpha_\ell)$  do
10:        Set  $(\delta r_i, \delta \theta_i, \delta \phi_i) \leftarrow \text{spherical}(\mathbf{r}_i - \mathbf{R}_{\beta_\ell})$ 
11:        for  $n = 0, \dots, p$  do
12:          for  $m = -n, \dots, +n$  do
13:            Update  $(L_{\ell,\beta_\ell}^{\Delta})_n^m \leftarrow (L_{\ell,\beta_\ell}^{\Delta})_n^m + q_i (\delta r_i)^{-(n+1)} Y_n^{-m}(\delta \theta_i, \delta \phi_i)$ 
14:          end for
15:        end for
16:      end for
17:    end for
18:  end for
19: end for

```

---

**Algorithm 4** Calculate change in electrostatic energy from direct interactions for a proposed move  $\mathbf{r} \rightarrow \mathbf{r}'$ . Input: initial position  $\mathbf{r} \in \alpha_L$  of particle with charge  $q$ ; target position  $\mathbf{r}' \in \alpha'_L$ . Output: change in direct electrostatic interaction energy  $\Delta U_{\text{direct}}$ .

---

```

1:  $\Delta U_{\text{direct}} \leftarrow 0$ 
2: for all particles with charge  $q_i$  and position  $\mathbf{r}_i \in \alpha_L \cup \mathcal{N}_b(\alpha_L)$ ,  $\mathbf{r}_i \neq \mathbf{r}$  do
3:   Update  $\Delta U_{\text{direct}} \leftarrow \Delta U_{\text{direct}} - \frac{qq_i}{|\mathbf{r} - \mathbf{r}_i|}$ 
4: end for
5: for all particles with charge  $q'_i$  and position  $\mathbf{r}'_i \in \alpha'_L \cup \mathcal{N}_b(\alpha'_L)$ ,  $\mathbf{r}'_i \neq \mathbf{r}'$  do
6:   Update  $\Delta U_{\text{direct}} \leftarrow \Delta U_{\text{direct}} + \frac{qq'_i}{|\mathbf{r}' - \mathbf{r}'_i|}$ 
7: end for
8: Remove self-interaction  $\Delta U_{\text{direct}} \leftarrow \Delta U_{\text{direct}} - \frac{q^2}{|\mathbf{r}' - \mathbf{r}|}$ 

```

---

2. Each MC proposal also requires the calculation of the short-range electrostatic interactions, which is not necessary in the accept stage.

The short-range contribution of the electrostatic potential is evaluated by calculating the contribution of all charges in  $\mathcal{N}_b(\alpha)$  and  $\mathcal{N}_b(\alpha')$  directly in Algorithm 4. As in the standard FMM algorithm the number of charges per fine level cell is constant and independent of the number of levels; each cell typically contains at the order of 1-10 charges. This guarantees that the total cost of the electrostatic calculation in the proposal step is still  $\mathcal{O}(\log(N))$  after the direct, short-range interactions are included using Algorithm 4.

We conclude that the computational complexity of the electrostatics in one MH step, which consists of a proposed move, potentially followed by one accepted transition, is  $\mathcal{O}(p^2 \log(N))$ .

The initialisation of the local expansions  $\Psi_{\ell,\alpha}^{\Delta}$  at the beginning of the simulation could be carried out in  $\mathcal{O}(p^4 N)$  time with a minimally modified variant of the standard FMM algorithm, which is written down explicitly as Algorithm 2 in [20]. Apart from renaming  $L_{\ell,\alpha} \mapsto L_{\ell,\alpha}^{\Delta}$  in the local expansions, the only difference is that line 17 of this algorithm has to be replaced by  $\Psi_{\ell,\alpha} \leftarrow 0$  and the loop over cells to construct  $\bar{\Psi}_{\ell,\alpha}$  in lines 13-15 is no longer necessary. However, since the setup cost is amortised anyway by the large number of MH steps, we chose a slightly more expensive but simpler approach, which is written down in Algorithm L.3 and avoids the calculation of multipole expansions in the upwards pass of the FMM algorithm. For this, the coefficients  $L_{\ell,\alpha}^{\Delta}$  are initialised to zero for all cells  $\alpha$  and levels  $\ell$ . Next, on each level we loop over all cells  $\alpha$  and increment  $\Psi_{\ell,\beta}^{\Delta}$  for all  $\beta \in \text{IL}(\alpha)$  by adding the contribution of all monopoles in  $\alpha$  to the local expansion in  $\beta$ . Since  $N$  monopoles have to be considered on each level, the computational complexity of the setup phase is  $\mathcal{O}(p^2 LN) = \mathcal{O}(p^2 N \log(N))$ .

### 3.3. Alternative algorithm based on multipole expansions

For reference, we now describe the alternative algorithm introduced in [12], which is based entirely on multipole expansions and which we also implemented for reference. In analogy to Equation (2), we define the multipole expansion of all particles contained in box  $\alpha$  on level  $\ell$  around the centre of the box

$$\Phi_{\ell,\alpha}(\delta \mathbf{r}) = \sum_{n=0}^p \sum_{m=-n}^{+n} (M_{\ell,\alpha})_n^m (\delta r)^{-(n+1)} Y_n^m(\delta \theta, \delta \phi) \quad \text{with } (\delta r, \delta \theta, \delta \phi) = \text{spherical}(\delta \mathbf{r}). \quad (3)$$

Assuming that the particles in the box  $\alpha$  have coordinates  $\delta\mathbf{r}_i$  and charges  $q_i$  with  $i \in I_\alpha \subset \{1, \dots, N\}$ , the explicit expression for the multipole expansion coefficients in Equation (3) is

$$(M_{\ell,\alpha})_n^m = \sum_{i \in I_\alpha} q_i (\delta r_i)^n Y_n^{-m}(\delta\theta_i, \delta\phi_i) \quad \text{with } (\delta r_i, \delta\theta_i, \delta\phi_i) = \text{spherical}(\delta\mathbf{r}_i). \quad (4)$$

Algorithms M.1 and M.2 describe how a potential move is proposed and accepted, using only multipole expansions; the two algorithms should be compared to Algorithms L.1 and L.2 above. Both methods have a computational complexity of  $\mathcal{O}(p^2 \log(N))$ , but observe that the expensive loops over the interaction list are now carried out in the proposal step in Algorithm M.1. Again, it would be possible to initialise the multipole expansion coefficients  $M_{\ell,\alpha}$  at the beginning of the simulation in  $\mathcal{O}(p^4 N)$  time with one upward pass of the native FMM method. However, for simplicity we chose to calculate them directly by looping over the cells on all levels and accumulating the multipole coefficients from all particles in a particular cell using Equation (4); this is written down explicitly in Algorithm M.3 which has  $\mathcal{O}(p^2 N \log(N))$  complexity. The resulting coefficients  $(M_{\ell,\alpha})_n^m$  are not identical to those that would have been obtained in the upward pass of the FMM algorithm, where they are calculated by recursively combining expansions on subsequent levels. However, the difference between the two ways of computing the multipole coefficients can be bounded as in the standard FMM error analysis.

---

**Algorithm M.1** Propose a move  $\mathbf{r} \rightarrow \mathbf{r}'$  using multipole expansions. Input: initial position  $\mathbf{r}$  of particle with charge  $q$ ; target position  $\mathbf{r}'$ . Output: change in total electrostatic energy  $\Delta U$ .

---

```

1: Find fine-level cells  $\alpha_L, \alpha'_L$  such that  $\mathbf{r} \in \alpha_L$  and  $\mathbf{r}' \in \alpha'_L$ 
2:  $\Delta U \leftarrow 0$ 
3: for  $\ell = L, L-1, \dots, 1$  do
4:   for  $\beta_\ell \in \text{IL}(\alpha_\ell)$  and  $\beta'_\ell \in \text{IL}(\alpha'_\ell)$  do
5:     Update  $\Delta U \leftarrow \Delta U + q \left( \Phi_{\ell,\beta'_\ell}(\mathbf{r}' - \mathbf{R}_{\beta'_\ell}) - \Phi_{\ell,\beta_\ell}(\mathbf{r} - \mathbf{R}_{\beta_\ell}) \right)$  using multipole expansion in Equation (3)
6:   end for
7:   if  $\ell > 1$  then
8:     Set  $\alpha_{\ell-1} \leftarrow \text{parent}(\alpha_\ell)$ ;  $\alpha'_{\ell-1} \leftarrow \text{parent}(\alpha'_\ell)$ 
9:   end if
10: end for
11: Calculate change  $\Delta U_{\text{direct}}$  in electrostatic energy from direct interactions with Algorithm 4
12:  $\Delta U \leftarrow \Delta U + \Delta U_{\text{direct}}$ 

```

---



---

**Algorithm M.2** Accept a move  $\mathbf{r} \rightarrow \mathbf{r}'$  using multipole expansions. Input: old position  $\mathbf{r}$ , new position  $\mathbf{r}'$ .

---

```

1: Find fine-level cells  $\alpha_L, \alpha'_L$  such that  $\mathbf{r} \in \alpha_L$  and  $\mathbf{r}' \in \alpha'_L$ 
2: for  $\ell = L, L-1, \dots, 1$  do
3:   Set  $(\delta r, \delta\theta, \delta\phi) \leftarrow \text{spherical}(\mathbf{r} - \mathbf{R}_{\alpha_\ell})$  and  $(\delta r', \delta\theta', \delta\phi') \leftarrow \text{spherical}(\mathbf{r}' - \mathbf{R}_{\alpha'_\ell})$ 
4:   for  $n = 0, \dots, p$  do
5:     for  $m = -n, \dots, +n$  do
6:       Update  $(M_{\ell,\alpha_\ell})_n^m \leftarrow (M_{\ell,\alpha_\ell})_n^m - q (\delta r)^n Y_n^{-m}(\delta\theta, \delta\phi)$ 
7:       Update  $(M_{\ell,\alpha'_\ell})_n^m \leftarrow (M_{\ell,\alpha'_\ell})_n^m + q (\delta r')^n Y_n^{-m}(\delta\theta', \delta\phi')$ 
8:     end for
9:   end for
10:  if  $\ell > 1$  then
11:    Set  $\alpha_{\ell-1} \leftarrow \text{parent}(\alpha_\ell)$ ;  $\alpha'_{\ell-1} \leftarrow \text{parent}(\alpha'_\ell)$ 
12:  end if
13: end for

```

---

### 3.4. Boundary conditions

So far we have implicitly assumed that free-space boundary conditions are used for the calculation of the electrostatic energy. In this case the interaction lists on the coarsest two levels of Algorithms L.2, L.3, M.1 and L.3 are empty. This implies that  $\Psi_{1,1}^\Delta = \Psi_{2,\alpha}^\Delta = \Phi_{1,1} = \Phi_{2,\alpha} = 0$  and levels  $\ell = 1, 2$  can be skipped when looping over the hierarchical tree. It is possible to adapt all algorithms in this section for simulations with periodic boundary conditions by making the following modifications:

- In Algorithms L.2 and L.3, extend the domain  $\Omega$  by  $3^3 - 1 = 26$  identical copies of the simulation cell to obtain an extended computational domain  $\bar{\Omega}$ . In the loop over  $\alpha_\ell$  and  $\alpha'_\ell$ , include the copies of those cells in the extended domain  $\bar{\Omega}$ .
- By following the approach described in detail in [20, Section 3.1], extend Algorithm L.3 to initialise the data structures  $K$  and  $E$  required to compute the electrostatic contribution of all charges outside  $\bar{\Omega}$ .

---

**Algorithm M.3** Initialise multipole expansion coefficients  $M_n^m$  for electrostatic calculation with Algorithms M.1 and M.2.

---

```

1: for levels  $\ell = 1, \dots, L$  do
2:   for all cells  $\alpha_\ell$  on level  $\ell$  do
3:     for  $n = 0, \dots, p$  and  $m = -n \dots +n$  do
4:       Set  $(M_{\ell, \alpha_\ell})_n^m = 0$ 
5:     end for
6:   end for
7:   for all cells  $\alpha_\ell$  do
8:     for all particles with charge  $q_i$  and position  $\mathbf{r}_i \in \alpha_\ell$  do
9:       Set  $(\delta r_i, \delta \theta_i, \delta \phi_i) \leftarrow \text{spherical}(\mathbf{r}_i - \mathbf{R}_{\alpha_\ell})$ 
10:      for  $n = 0, \dots, p$  do
11:        for  $m = -n, \dots, +n$  do
12:          Update  $(M_{\ell, \alpha_\ell})_n^m \leftarrow (M_{\ell, \alpha_\ell})_n^m + q_i (\delta r_i)^n Y_n^{-m}(\delta \theta_i, \delta \phi_i)$ 
13:        end for
14:      end for
15:    end for
16:  end for
17: end for

```

---

- ```

1: Set  $K_n^m \leftarrow 0, E_n^m \leftarrow 0 \forall m, n$ 
2: for all particles with charge  $q_i$  and position  $\mathbf{r}_i \in \alpha_\ell$  do
3:   Set  $(\delta r_i, \delta \theta_i, \delta \phi_i) \leftarrow \text{spherical}(\mathbf{r}_i - \mathbf{R}_1)$ 
4:   for  $n = 0, \dots, p$  do
5:     for  $m = -n, \dots, +n$  do
6:       Set  $K_n^m \leftarrow K_n^m + q_i (\delta r_i)^n Y_n^{-m}(\delta \theta_i, \delta \phi_i)$ 
7:       Set  $E_n^m \leftarrow E_n^m + q_i (\delta r_i)^n Y_n^m(\delta \theta_i, \delta \phi_i)$ 
8:     end for
9:   end for
10: end for
11: Store  $K$  and  $E$ .

```
- Extend Algorithm L.2 to update  $K$  and  $E$  when a move is accepted.

```

1: Set  $(\delta r, \delta \theta, \delta \phi) \leftarrow \text{spherical}(\mathbf{r} - \mathbf{R}_1)$  and  $(\delta r', \delta \theta', \delta \phi') \leftarrow \text{spherical}(\mathbf{r}' - \mathbf{R}_1)$ 
2: for  $n = 0, \dots, p$  do
3:   for  $m = -n, \dots, +n$  do
4:     Set  $K_n^m \leftarrow K_n^m + q \left( (\delta r')^n Y_n^{-m}(\delta \theta', \delta \phi') - (\delta r)^n Y_n^{-m}(\delta \theta, \delta \phi) \right)$ 
5:     Set  $E_n^m \leftarrow E_n^m + q \left( (\delta r')^n Y_n^m(\delta \theta', \delta \phi') - (\delta r)^n Y_n^m(\delta \theta, \delta \phi) \right)$ 
6:   end for
7: end for

```
  - Extend Algorithm L.1 to include the contributions to the energy differences from the proposed move in periodic images outside  $\bar{\Omega}$  by computing the proposed differences to  $K$  and  $E$  to determine the change in energy (the linear operator  $\mathcal{R}$  is introduced in [20, Section 2.2.1]).

```

1: Set  $H \leftarrow \mathcal{R}(K)$ 
2: Set  $U^\infty \leftarrow \sum_{n=0}^p \sum_{m=-n}^n E_n^m H_n^m$ 
3: Set  $(\delta r, \delta \theta, \delta \phi) \leftarrow \text{spherical}(\mathbf{r} - \mathbf{R}_1)$  and  $(\delta r', \delta \theta', \delta \phi') \leftarrow \text{spherical}(\mathbf{r}' - \mathbf{R}_1)$ 
4: for  $n = 0, \dots, p$  do
5:   for  $m = -n, \dots, +n$  do
6:     Set  $\delta K_n^m \leftarrow K_n^m + q \left( (\delta r')^n Y_n^{-m}(\delta \theta', \delta \phi') - (\delta r)^n Y_n^{-m}(\delta \theta, \delta \phi) \right)$ 
7:     Set  $\delta E_n^m \leftarrow E_n^m + q \left( (\delta r')^n Y_n^m(\delta \theta', \delta \phi') - (\delta r)^n Y_n^m(\delta \theta, \delta \phi) \right)$ 
8:   end for
9: end for
10: Set  $\delta H \leftarrow \mathcal{R}(\delta K)$ 
11: Set  $\Delta U^\infty \leftarrow \sum_{n=0}^p \sum_{m=-n}^n \delta E_n^m \delta H_n^m - U^\infty$ 
12: Set  $\Delta U \leftarrow \Delta U + \Delta U^\infty$ 

```
  - In Algorithm 4, extend  $\mathcal{N}_b(\alpha_L)$  and  $\mathcal{N}_b(\alpha'_L)$  to include all cells in the extended domain  $\bar{\Omega}$ . Further, the self-interaction term has to be modified to take into account spurious interactions with the additional copies of the original particle. As discussed in detail in [20, Section 3.1] this can be done by replacing the update in line 8 of Algorithm 4 by

$$\Delta U_{\text{direct}} \leftarrow \Delta U_{\text{direct}} - \sum_{\mathbf{v} \in [-1, 0, +1]^3} \frac{q^2}{|\mathbf{r}' - (\mathbf{r} + a\mathbf{v})|} + \frac{q^2}{a} \left( 6 + \frac{8}{\sqrt{3}} + \frac{12}{\sqrt{2}} \right).$$

Similar modifications have to be made to Algorithms M.1 and M.3. In practice, iteration over the 26 additional copies of the simulation cells can be implemented by modifying data structures such as neighbour lists, see [20] for a more detailed discussion.

#### 4. Implementation

The algorithms described in this paper have been implemented as an extension to the performance portable framework for molecular dynamics (PPMD) described in [13], which is freely available at

<https://github.com/ppmd/ppmd>

PPMD provides a high-level Python interface for particle-based simulations which require the efficient execution of user-defined operations over all particles or pairs of particles in a system. An obvious example of the latter is the calculation of inter-particle forces in molecular dynamics simulations, but the interface is sufficiently abstract to support more general operations such as the structure analysis algorithms discussed in Section 4 of [13]. PPMD automatically generates efficient code for executing short user-defined C-kernels for all particles or particle-pairs on different parallel architectures; both distributed- and shared- memory parallelism are supported and the code can run on non-standard architectures such as GPUs. Particle-specific data (such as e.g. charge, mass and velocity) is stored as instances of the Python `ParticleDat` class. Particle positions, which contain information that is relevant for parallel domain decompositions, are stored as instances of the specialised `PositionDat` class. In addition to electrostatic potential- and force- calculation with classical Ewald summation [4,21], PPMD also contains an implementation of the standard FMM algorithm in three dimensions given in [7]. The PPMD framework therefore provides all necessary data structures for storing information (such as local- and multipole-expansion coefficients) on a nested hierarchy of grids which is required to implement the algorithms discussed in this paper. Algorithms L.1 - L.3, M.1 - M.3 and 4 are implemented in the separate `coulomb_mc` Python package which is based on PPMD and can be downloaded from

[https://github.com/ppmd/coulomb\\_mc](https://github.com/ppmd/coulomb_mc)

Algorithms L.1, L.2, M.1 and M.2 have been implemented as auto-generated C-code. This allows the pre-computation of constant expressions such as combinatorial factors that arise in the evaluation of the spherical harmonics and unrolling of nested loops such as the ones in lines 5-10 of Algorithm L.2 and lines 4-9 in Algorithm M.2. Finally, the generated code is compiled for a specific chip architecture at runtime to ensure optimal performance. Currently our implementation supports cuboid geometries with free space- and periodic boundary conditions.

##### 4.1. FMM-MC user interface

Recall that the local expansion coefficients  $L_{\ell,\alpha}^{\Delta}$  which are required to compute (changes of) the electrostatic energy of the system of  $N$  particles with charges  $\{q_1, q_2, \dots, q_N\}$  and positions  $\{\mathbf{r}_1, \mathbf{r}_2, \dots, \mathbf{r}_N\}$  are initialised with Algorithm L.3. In the `coulomb_mc` package the charges of all particles are represented as a `PositionDat` instance  $q$ , whereas the positions are stored as a `ParticleDat` object  $r$ . At the beginning of the simulation the user populates  $q$  and  $r$  and uses those to create a `MCFMM_LM` object, which is passed additional information on the domain (such as boundary conditions), the number of levels  $L$  in the grid hierarchy and the number of expansion terms. The constructor of the `MCFMM_LM` class then executes Algorithm L.3 to initialise the values of the coefficients  $L_{\ell,\alpha}^{\Delta}$ . Following this, Algorithm L.1 can be used to compute the change in electrostatic energy which occurs if particle  $j$  transitions from its original position  $\mathbf{r} = \mathbf{r}_j$  to a new position  $\mathbf{r}' = \mathbf{r}'_j$ . In the code this is realised by calling the `propose()` method of the `MCFMM_LM` object and passing it the particle index  $j$  and the new position  $\mathbf{r}' = \mathbf{r}'_j$ . Finally, Algorithm L.2 updates the local expansions  $L_{\ell,\alpha}^{\Delta}$  if the proposed move  $\mathbf{r} \mapsto \mathbf{r}'$  is accepted. Calling the class method `accept()`, which is passed the particle index  $j$  and the new position  $\mathbf{r}'_j$  of particle  $j$  also executes Algorithm L.2 for the transition  $\mathbf{r}_j \mapsto \mathbf{r}'_j$ . It replaces the position of particle  $j$  in the `PositionDat`  $r$  by  $\mathbf{r}'_j$  and updates the expansion coefficients  $L_{\ell,\alpha}^{\Delta}$ .

Listing 1 illustrates the creation of an `MCFMM_LM` object, followed by the computation of the change in energy  $dU = \Delta U$  which would result from moving particle  $j = 7$  to the new position `xj_new = (0.1, 1.0, 5.2)` by calling the `propose()` method. In the last line the move to the new position is accepted by calling the `accept()` method. Note that although the example in Listing 1 assumes that the proposed position is identical to the accepted position, this is not the case in general. Because of this and since the code keeps track of the *total* energy of the system at each step, by default the `accept()` method executes Algorithm L.1 to compute the change in system energy  $\Delta U$ . This can be avoided by passing this change  $\Delta U$  (which - as shown in Listing 1 - might have been computed in a previous call to `propose()` with the new position that is to be accepted) as an additional parameter to the `accept()` method. The corresponding multipole-based Algorithms M.1 - M.3 can be used by creating an `MCFMM_MM` object which keeps track of the multipole expansion coefficients  $M_{\ell,\alpha}^{\Delta}$ . The constructor of this class implements Algorithm M.3. The class methods `propose()` and `accept()` implement Algorithms M.1 and M.2 and can be used in exactly the same way as the corresponding methods of the `MCFMM_MM` class described above.

---

**Listing 1** Illustration of how a FMM\_MC instance is created and then subsequently used to propose and accept moves.

---

```

# Create MCFMM_LM instance with positions in PositionDat r and charges in ParticleDat q.
# Here we use 3 tree levels and 12 expansion terms
MC = MCFMM_LM(r, q, domain, 'pbc', r=3, l=12)
# Perform the initial electrostatic solve
MC.initialise()

# Consider move of particle j=7 to new position rj_new = (0.1, 1.0, 5.2)
j = 7
rj_new = np.array((0.1, 1.0, 5.2))

# Propose a move of charge j to position rj_new
dU = MC.propose((j, rj_new))

# Accept a move of charge j to position rj_new
MC.accept((j, rj_new), dU)

```

---

Note that the aim of `coulomb_mc` is to provide functionality for the calculation of (changes in) the electrostatic energy through the high-level `MCFMM_LM` and `MCFMM_MM` classes, which typically dominates the runtime. It is up to the user to implement the overarching Monte Carlo algorithm which generates proposed new positions  $\mathbf{r}'$  and uses the calculated energy differences to accept or reject particular moves, e.g. in a Metropolis Hastings step.

## 5. Results

In the following we quantify the performance of the algorithms introduced in Sections 3.2 and 3.3 and implemented as described in Section 4. We demonstrate numerically that, as expected, the time spent in each Monte Carlo step increases logarithmically with the number of particles in the system. To assess its overall performance, we also compare the runtime of our code to version 2.06 of the `DL_MONTE` package [14,15] which uses the classical Ewald method to compute electrostatic interactions.

All numerical experiments were carried out on the University of Bath “Balena” HPC cluster. Compute nodes of this machine consist of two Intel E5-2650v2 CPUs, and all timing results are reported for sequential runs on a single core. A snapshot of the source code which can be used to reproduce the results, along with all plotting scripts and raw data is provided at [22]. The code was compiled using version 19.5.281 of the Intel compiler; `DL_MONTE` was compiled with version 17.1.132 of the same compiler.

### 5.1. Configuration and parameter selection

The accuracy of the algorithms described in Sections 3.2 and 3.3 crucially depends on the number of local/multipole expansion terms, which can be quantified by  $p$ , the upper limit in the outer sum in Equations (2) and (3). To provide a fair comparison between the methods introduced in this paper and the Ewald implementation in `DL_MONTE`,  $p$  is adjusted such that acceptance probabilities have errors which are comparable to those in `DL_MONTE`. For a proposed move  $\mathbf{r} \rightarrow \mathbf{r}'$  which results in a change of energy of  $\Delta U$ , the relative error in the acceptance probability  $\delta P$  is defined as

$$\delta P = \frac{|P - P^*|}{|P^*|}. \quad (5)$$

Here  $P = \exp(-\Delta U/(k_B T))$  is the acceptance probability (computed by `DL_MONTE` or an expansion based method) for a given choice of parameters. Assuming that the exact change in energy is  $\Delta U^*$ , the exact acceptance probability is denoted as  $P^* = \exp(-\Delta U^*/(k_B T))$ . For a particular move  $P^*$  is approximated to high accuracy by computing  $\Delta U^*$  with the local expansion based method and 26 expansion terms ( $p = 25$ ).

The configuration for our numerical experiments is based on TEST01 [23] in the `DL_POLY` suite. This setup simulates a simple cubic NaCl crystal of alternating Sodium (Na) and Chloride (Cl) ions with a lattice constant of  $a = 3.3\text{\AA}$ . Fully periodic boundary conditions are used for all numerical experiments, which are performed at a temperature of  $T = 273\text{K}$ .

To estimate the relative errors  $\delta P$  in the acceptance probability, we start with an initial configuration of charges which is constructed by creating a cubic lattice of  $22 \times 22 \times 22 = 10648$  ions as described in TEST01 and perturbing the initial position of each ion by adding a uniform random shift with a maximum size of  $0.01a$  in each spatial direction. Based on this, 1000 moves are proposed (note that no moves are accepted) and for each move the acceptance probabilities  $P$  for `DL_MONTE` and the expansion based approaches are computed along with the “exact” acceptance probability  $P^*$ , which is estimated as described above. This process is repeated for 16 different initial configurations to generate a total of 16000 samples for the quantity  $\delta P$  defined in Equation (5).

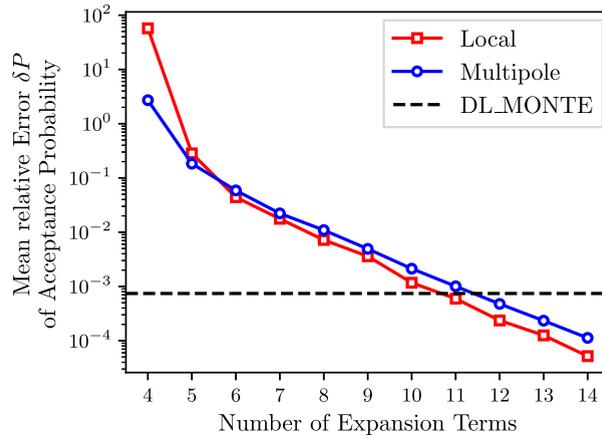


Fig. 3. Relative error  $\delta P$  in acceptance probability for DL\_MONTE and the expansion based methods for a varying number of expansion terms.

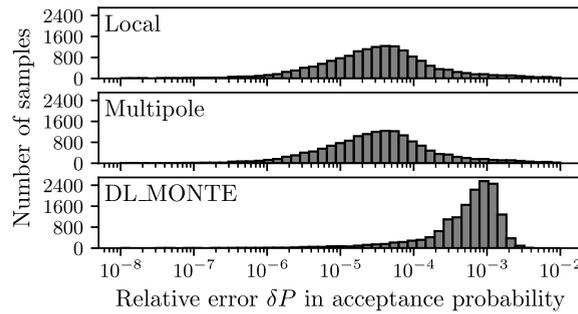


Fig. 4. Histograms of relative error  $\delta P$  in acceptance probability. Results are shown for the local expansion based algorithm (top), the multipole expansion based algorithm (middle) and DL\_MONTE (bottom); 12 expansion terms ( $p = 11$ ) are used for first two methods.

Fig. 3 shows the mean relative error  $\delta P$  (averaged over all 16000 samples) as a function of the number of expansion terms, which varies between 4 and 14 for the expansion based methods. This should be compared to the same quantity computed with DL\_MONTE at a fixed solver tolerance of  $10^{-6}$ , indicated by the horizontal dashed line. As those results show, choosing 12 expansion terms ( $p = 11$ ) results in a comparable mean relative error  $\delta P$  which is smaller than  $10^{-3}$ . Note that for a fixed value of  $p$  the local expansion based method (Algorithm L.1) has a slightly higher accuracy than the method which only uses multipole expansions (Algorithm M.1). To compare the errors of all used methods in more detail, we also inspect the distribution of  $\delta P$  over all 16000 samples. Fig. 4 shows a histogram of the relative error in the acceptance probability, i.e. the number of samples which have a  $\delta P$  that falls into a certain interval  $[\delta P_1, \delta P_2]$ . Results are shown both for DL\_MONTE (again using a solver tolerance of  $10^{-6}$ ) and our expansion based methods with 12 expansion terms. The cumulative density of the probability distribution in Fig. 4 is plotted in Fig. 5. In other words, for a given tolerance  $\epsilon$  on  $\delta P$ , Fig. 5 shows the percentage of samples that have a relative error which does not exceed  $\epsilon$ . As both figures demonstrate, the spread in errors is slightly larger for the expansion based methods: although for those methods a larger fraction of samples have errors well below the tolerance of  $10^{-3}$ , there is a small number of outliers. This, however, is consistent between the two expansion based methods.

Finally, observe that a large relative error  $\delta P$  in the acceptance probability  $P$  will only translate into a large absolute error on  $P$  if  $\Delta U^*$  is also large. It is therefore instructive to also produce a scatter plot of  $\Delta U^*$  against  $\delta P$  for all samples and this is shown in Fig. 6.

### 5.2. Computational complexity

Next, we investigate the growth in computational cost as a function of the number of charges  $N$ . Formally the number of levels  $L$  of the hierarchical tree is  $\mathcal{O}(\log(N))$ . The relative proportion of time spent evaluating the local and multipole expansions in the propose stage (Algorithms L.1 and M.1), update of expansion coefficients (Algorithms L.2 and M.2) and direct, nearest neighbour calculations (Algorithm 4) can be controlled by setting  $L = \lfloor \log_8(\alpha N) \rfloor$  and varying the constant  $\alpha$  (here  $\lfloor \cdot \rfloor$  denotes the floor function defined by  $\lfloor x \rfloor = \max\{n \in \mathbb{N} : n \leq x\}$ ). The optimal value of the parameter  $\alpha$  depends on the computer hardware, the average acceptance rate  $\nu$  and the number of expansion terms. We define the acceptance rate as

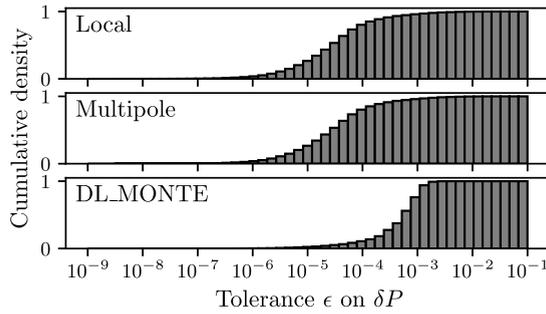


Fig. 5. Cumulative density of the relative error  $\delta P$  in computed probabilities. Results are shown for the local expansion based algorithm (top), the multipole expansion based algorithm (middle) and DL\_MONTE (bottom); 12 expansion terms ( $p = 11$ ) are used for first two methods.

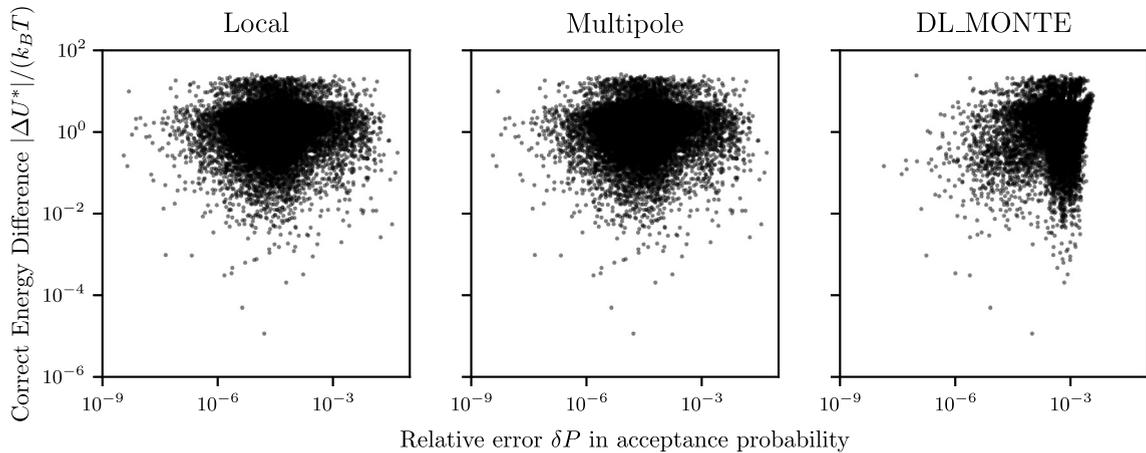


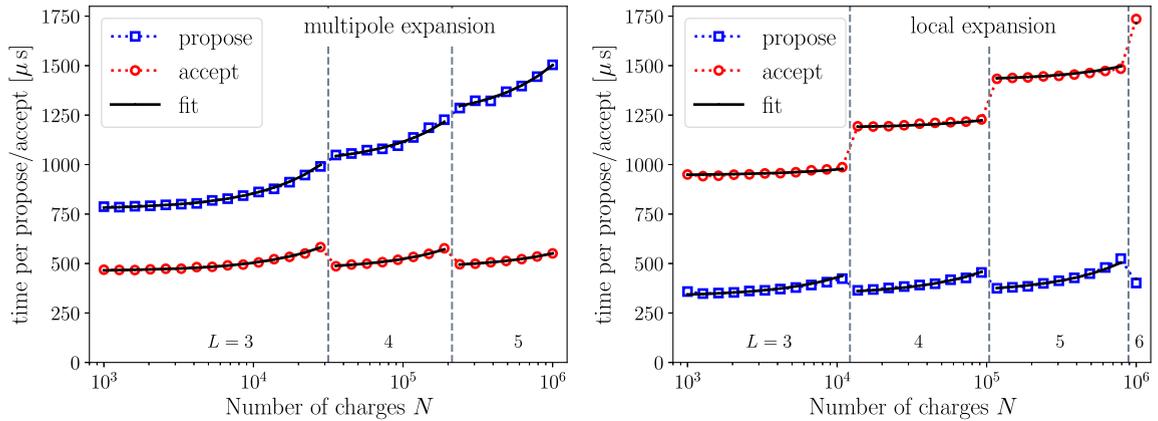
Fig. 6. Scatter plot of  $\Delta U^*/(k_B T)$  against the relative error  $\delta P$  in acceptance probability. Results are shown for the local expansion based algorithm (left), the multipole expansion based algorithm (middle) and DL\_MONTE (right); 12 expansion terms ( $p = 11$ ) are used for first two methods.

$$\nu = \frac{\text{Number of accepted moves}}{\text{Number of proposed moves}} \tag{6}$$

Assuming that on average  $\nu^{-1} = e \approx 2.718$  proposals have to be generated for each accepted move, we find that for our setup and with 12 expansion terms ( $p = 11$ ), the best results are obtained for  $\alpha = 0.327$  when using the local expansion based method (Algorithms L.1, L.2) and  $\alpha = 0.138$  if the multipole method (Algorithms M.1, M.2) is used. This also implies that for a given value of  $N$ , the optimal number of levels for both methods differs by less than 0.5.

To quantify the computational complexity of the propose stage (Algorithms L.1 and M.1) and the accept stage (Algorithms L.2 and M.2) separately, random proposals and accepted moves are generated for problems of increasing size, drawing particle positions and charges from a uniform random distribution with a maximal absolute displacement of 0.25 Å in each spatial direction. We investigated the computational cost of proposing and accepting moves for systems containing between a thousand ( $N = 10^3$ ) and a million ( $N = 10^6$ ) particles. For each  $N$  the initial arrangement of particles is constructed as described in Section 5.1. Fig. 7 shows the average time (measured over 1000 samples) per propose or accept operation as a function of the number  $N$  of charges in the system; the fitted solid lines are discussed in Section 5.4 below. The measured times increase abruptly as the number of levels  $L$  changes as incrementing the number of levels increases the number of expansions that must be evaluated and updated. Although asymptotically we expect all times to grow as  $L \propto \log(N)$ , there are significant differences in the rate of growth and absolute computational cost for the different implementations. While for the multipole based method from [12] proposing a single move is significantly more expensive than accepting it, the opposite is true for our new method based on local expansions. The main reason for this is that the expensive loop over cells in the interaction list has to be executed in the propose stage on the of the multipole based method (Algorithm M.1), whereas the interaction list is traversed in the accept stage of our new method (Algorithm L.2). Overall we therefore expect our new method to be more efficient as the acceptance rate  $\nu$  decreases, and the number of proposals is significantly larger than the number of accepted moves. In Metropolis Hastings simulations this is the case since the acceptance rate is usually significantly less than 1, a typical value is  $1/e \approx 0.3679$ .

Although one would naively expect the runtime of the multipole-based accept (Algorithm M.2) and the local-based propose (Algorithm L.1) to be roughly identical, the measured cost of the latter is slightly smaller. Similarly, Fig. 7 shows that the multipole-based propose (Algorithm M.1) is slightly faster than the local-based accept (Algorithm L.2). This can be explained by details of the implementation. Firstly, the bookkeeping operations for the propose- and accept stages introduce



**Fig. 7.** Time spent in the propose and accept operations as a function of the number of charges  $N$  in the system. Results in the left figure were obtained with the multipole expansion based method in [12] and show the time spent in Algorithm M.1 ( $T_{\text{prop}}^{(\text{mp})}$ , blue squares) and Algorithm M.2 ( $T_{\text{acc}}^{(\text{mp})}$ , red circles). The right figure shows the corresponding numbers for our new method based on local expansion; namely the time spent in Algorithm L.1 ( $T_{\text{prop}}^{(\text{loc})}$ , blue squares) and Algorithm L.2 ( $T_{\text{acc}}^{(\text{loc})}$ , red circles). The step-changes in measured times (marked by dashed vertical lines) correspond to increases in the number of levels  $L$ , which are shown at the bottom of each figure. The fit to the data (solid black lines) is discussed in Section 5.4.

different overheads. For example, when proposing a move the cells  $\alpha_\ell$ ,  $\alpha'_\ell$  have to be computed and when a move is accepted data structures, such as the indirection map that assigns particles to cells, must be updated. Furthermore, in practice reading and writing floating point numbers carries a different cost. For example, in the loop over the interaction list in Algorithm L.2 expansion coefficients are updated, which requires a read and a write operation, however in Algorithm M.1 in the loop over the interaction list expansion coefficients are only read, which avoids any potential write contention. Finally bear in mind that the simulation uses periodic boundary conditions. In our implementation this introduces a small overhead when proposing a move and a slightly larger additional cost when a move is accepted.

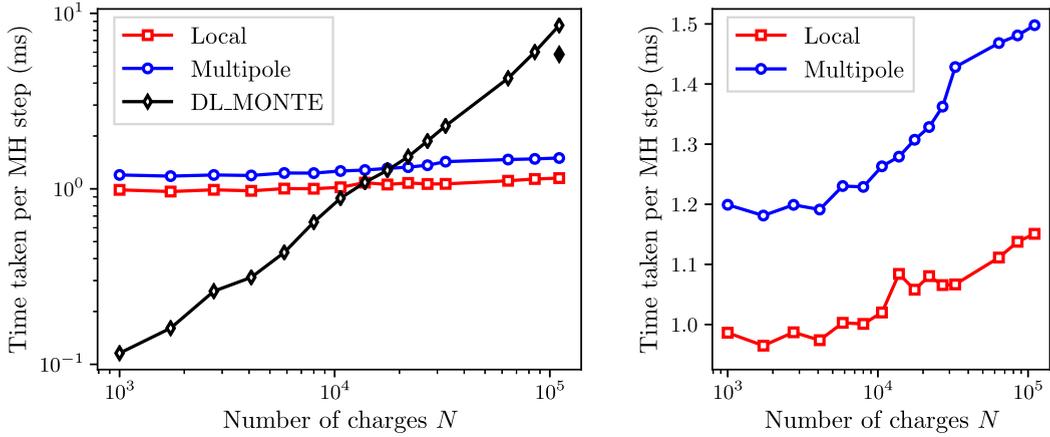
To understand the growth of the runtime for increasing problem size  $N$  in more detail observe the following: when an additional level is added to the octal tree as  $N$  increases, initially there are  $N/N_{\text{cell}} = \mathcal{O}(1)$  charges per cell on the finest level. However, as the total number of charges increases while the number of levels  $L$  is kept fixed, the average number of charges per cell will grow. Hence the cost of Algorithms L.1 and M.1 grows in proportion to  $N/N_{\text{cell}}$  for fixed  $L$  since both require the computation of direct interactions with Algorithm 4. In addition our implementation of Algorithms L.2 and M.2 contains a small number of bookkeeping operations which are formally  $\mathcal{O}(N)$ . As discussed in Section 5.4, the absolute contribution of those operations to the total runtime is very small (and will be amortised by the cost of direct interactions induced by other short-range potentials such a repulsive Lennard-Jones field in real-life simulations). Ignoring this small  $\mathcal{O}(N)$  contribution the asymptotic complexity of Algorithms L.1, L.2, M.1 and M.2 is  $\mathcal{O}(L) = \mathcal{O}(\log(N))$ .

### 5.3. Comparison of full Monte Carlo simulation with DL\_MONTE

Having quantified the time spent in the propose- and accept-stage of our expansion based algorithm separately, we now discuss the growth of the total runtime of an entire Monte Carlo simulation as a function of the problem size. For this we compare the performance of our expansion-based implementations in PPMD against DL\_MONTE. Again we consider an NaCl crystal with the same initial arrangement of charges as described in Section 5.1. To prevent oppositely charged ions from collapsing onto one another over the course of the simulation, a repulsive short-range Lennard-Jones potential with a fixed cutoff of 12 Å is added. For the largest problem sizes (with  $N \approx 10^5$  in a cubic box with a side length of 153.2 Å) this short-range potential adds an additional average cost of 0.18 ms per Monte Carlo step for our expansion based methods. This accounts for approximately 14% of the total average cost per step for the local expansion implementation.

For each proposed move we create a random offset vector  $\delta \mathbf{r} = \mathbf{r}' - \mathbf{r}$ , such that each component  $\delta r_j$  is uniformly distributed in the interval  $[-0.25 \text{ \AA}, +0.25 \text{ \AA}]$ . This resulted in an average acceptance rate of  $\nu \approx 0.438$  for the expansion based methods. Note that the performance of the Ewald implementation in DL\_MONTE is not sensitive to the acceptance rate and that in DL\_MONTE the additional cost of accepting a proposed move is negligible.

Fig. 8 shows the time per MC step for our implementations of the expansion based methods and for DL\_MONTE as a function of the number of charges  $N$ . The size of the system varies between  $N = 10^3$  and  $N = 10^5$  charges and the reported times are averaged over 1000 Metropolis-Hastings steps (i.e. 1000 proposed moves). We do not include the setup times, since those account for an insignificant fraction of the runtime for “real” simulations with a large number of Metropolis-Hastings steps. For  $N < 10^4$  particles DL\_MONTE provides more performance than our implementations, but it is almost an order magnitude slower for the largest considered problem size ( $N = 10^5$ ). Empirically the cost of the expansion based methods grows only relatively slowly with the problem size. For DL\_MONTE, on the other hand, the runtime is approximately proportional to the number of particles. This  $\mathcal{O}(N)$  growth can be explained as follows: for the direct Ewald



**Fig. 8.** Time per Metropolis Hastings step for our implementations of the expansion based methods and DL\_MONTE. In the latter case the cutoff was kept fixed at  $r_c = 12 \text{ \AA}$  for all problem sizes, except for the rightmost data point where results with an optimised cutoff of  $r_c = 28 \text{ \AA}$  are also shown as a solid diamond. The plot on the right shows the same data but with a linear scale on the vertical axis.

summation method with a short-range cutoff  $r_c$ , short-range interactions with  $\mathcal{O}(\rho r_c^3)$  particles have to be computed for each proposed move in a system with density  $\rho$ . To balance errors in the short- and long-range computation, the physical cutoff in Fourier space and  $r_c$  are inversely related. As a consequence, the number of Fourier modes that have to be considered grows like  $\mathcal{O}(V r_c^{-3})$ . Hence, for fixed density where  $V = \rho N$  the combined cost of the short- and long- range contributions is

$$\text{Cost}_{\text{Ewald}}(r_c, N) = \mathcal{O}(r_c^3 + N r_c^{-3}). \quad (7)$$

This cost can be minimised by varying the cutoff such that  $r_c = r_c^{(0)} N^{1/6}$  for some constant  $r_c^{(0)}$ , as discussed for example in [24] (see also [17, Chapter 12]), resulting in a total cost of  $\text{Cost}_{\text{Ewald}}(r_c^{(0)} N^{1/6}, N) = \mathcal{O}(N^{1/2})$  per proposal for the Ewald method. In the current version of DL\_MONTE the short-range cutoff for the Ewald summation has to be identical to the cutoff for Lennard-Jones interactions, which we fixed at  $r_c = 12 \text{ \AA}$  to obtain the majority of the results in Fig. 8. In addition we also quantify how the results would change if the optimal short range cutoff is chosen. At the moment this requires manual fine-tuning and for practical reasons it was not possible to do this for all problem sizes. Equation (7) shows that for a fixed cutoff of  $r_c = 12 \text{ \AA}$  the cost of the Ewald summation is dominated by the long range contribution and grows in proportion to  $N$ . While this is a current limitation of DL\_MONTE and not a fundamental issue, it is worth stressing that even if this problem is overcome, Ewald summation has a computational complexity of  $\mathcal{O}(N^{1/2})$  compared to the  $\mathcal{O}(\log(N))$  complexity of our hierarchical methods. To demonstrate the effect of a more optimal short-range cutoff we varied the cutoff radius  $r_c$  between  $12 \text{ \AA}$  and  $32 \text{ \AA}$  and found that a cutoff of  $r_c = 28 \text{ \AA}$  gives near optimal results. As shown by the rightmost datapoints in Fig. 8, where the results obtained with  $r_c = 28 \text{ \AA}$  are indicated by a solid diamond, this reduces the time per MH step from  $8.0 \text{ ms}$  to  $5.8 \text{ ms}$  for  $N \approx 10^5$  charges. This value might be reduced further in future releases of DL\_MONTE, if the different cutoffs can be varied independently to avoid the evaluation of short range potentials with an unnaturally large cutoff. With an optimally tuned cutoff the Ewald crossover between the Ewald- based method and our hierarchical algorithms would occur for larger problem sizes.

Finally, note that in its current implementation the setup cost of DL\_MONTE grows like  $\mathcal{O}(N^2)$  instead of  $\mathcal{O}(N^{3/2})$  which could be achieved for an optimal  $r_c \propto N^{1/6}$ . Although not considered here, this  $\mathcal{O}(N^2)$  setup time can become computationally significant for systems containing a large number of charges.

#### 5.4. Empirical results for growth of runtime

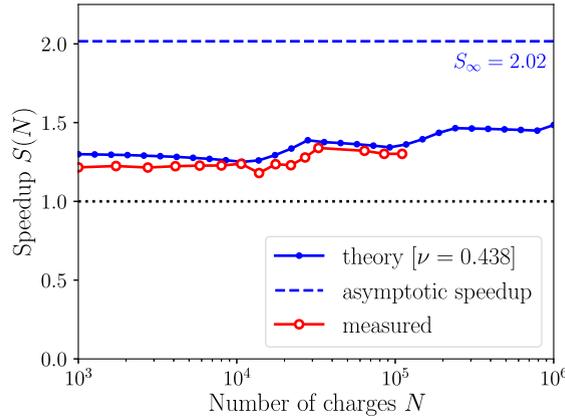
As will be shown below, the average speedup of our method relative to the one in [12] is 1.25 over all considered problem sizes. The smallest measured speedup is 1.18, while the maximal speedup is 1.34, with a median value of 1.23. To model the speedup for large  $N$  and establish an upper limit on the performance gains as  $N \rightarrow \infty$  we fit the measured times in Fig. 7 to

$$T(N) = \tau + \gamma L + \kappa \frac{N}{N_{\text{cell}}} + \sigma N \quad (8)$$

with a least squares approach. In this expression  $L$  is the number of levels of the octree hierarchy for a given value of  $N$  and  $N_{\text{cell}} = 8^{L-1}$  is the number of cells on the finest level for this octree depth. The numerical values of the four fit parameters  $\tau$ ,  $\gamma$ ,  $\kappa$  and  $\sigma$  for the propose- and accept stages of the multipole- and local expansion based methods are given in Table 1.

**Table 1**  
Numerical values of fit parameters for empirical performance model in Equation (8).

| Expansion | Stage   | $\tau$ | $\gamma$ | $\kappa$ | $\sigma$             |
|-----------|---------|--------|----------|----------|----------------------|
| Multipole | propose | 89.15  | 228.40   | 0.50     | $1.49 \cdot 10^{-4}$ |
|           | accept  | 436.71 | 8.12     | 0.27     | $7.11 \cdot 10^{-6}$ |
| Local     | propose | 304.01 | 9.92     | 0.61     | $4.12 \cdot 10^{-5}$ |
|           | accept  | 225.80 | 239.99   | 0.19     | $4.37 \cdot 10^{-5}$ |



**Fig. 9.** Theoretical speedup  $S(N)$  (solid blue line) as defined in Equations (8) and (9) for  $\nu = 0.438$ ; the asymptotic speedup  $S_\infty$  (dash blue line), is given in Equation (10). The measured speedup, in solid red, is computed from the local and multipole timings plotted in Fig. 8.

As the solid curves in Fig. 7 show, Equation (8) models the data extremely well. As expected the coefficient of the  $\mathcal{O}(N)$  term is very small. The empirical model in Eq. (8) allows to predict the expected speedup for a given problem size  $N$  and acceptance rate  $\nu$  as

$$S(N) = \frac{T_{\text{prop}}^{(\text{mp})}(N) + \nu T_{\text{acc}}^{(\text{mp})}(N)}{T_{\text{prop}}^{(\text{loc})}(N) + \nu T_{\text{acc}}^{(\text{loc})}(N)} \tag{9}$$

where superscripts “(loc)” and “(mp)” denote the local and multipole expansion based methods and subscripts “prop” and “acc” label propose and accept operations respectively. Ignoring the  $\mathcal{O}(N)$  term in Equation (8) (this can be justified by the small value of  $\sigma$  in Table 1), the asymptotic speedup is

$$S_\infty = \lim_{N \rightarrow \infty} S(N) = \frac{\gamma_{\text{prop}}^{(\text{mp})} + \nu \gamma_{\text{acc}}^{(\text{mp})}}{\gamma_{\text{prop}}^{(\text{loc})} + \nu \gamma_{\text{acc}}^{(\text{loc})}} \tag{10}$$

Further observe that asymptotically the relative cost of Algorithms L.1 and M.1 is  $\gamma_{\text{prop}}^{(\text{mp})} / \gamma_{\text{prop}}^{(\text{loc})} = 23.03$  while the relative time spent in Algorithms M.2 and L.2 is  $\gamma_{\text{acc}}^{(\text{loc})} / \gamma_{\text{acc}}^{(\text{mp})} = 29.56$ . These ratios deviate from the theoretically expected value of 189 (which, as the reader will recall, is the number of cells in the interaction list) due to details of the implementation. In particular looping over a larger set of cells in the interaction lists allows the compiler to carry out additional optimisation and vectorise the code.

In Fig. 9 we plot the theoretical and measured speedup for an acceptance rate  $\nu = 0.438$ . Here the theoretical speedups (blue) predicted with Equation (8) are computed from the time per propose and accept operation computed in Section 5.2 and fitted with the model in Equation (8). The measured speedups (red) are obtained from the data visualised in Fig. 8. The theoretical upper limit  $S_\infty$  for the speedup defined in Equation (10) is also shown; for the fitted values of  $\gamma$  given in Table 1 and with  $\nu = 0.438$ , we find  $S_\infty = 2.02$ . As Fig. 9 demonstrates, the fit reproduces the measured speedup, which lies in the range [1.18, 1.34], reasonably well. The expected asymptotic speedup for large  $N$  is around 2, demonstrating that our method has the potential to approximately halve the runtime compared to the algorithm in [12]. However, this speedup is only reached for values of  $N$  which are significantly larger than the ones considered in this work; for  $N = 10^6$  we expect a speedup of roughly  $1.5\times$ .

## 6. Conclusion

In this paper we presented a new hierarchical method for accurately including electrostatic interactions in Monte Carlo simulations. Our algorithm has a computational complexity of  $\mathcal{O}(\log(N))$  per Metropolis Hastings step. Compared to the

related technique in [12], our method reduces the average cost of a MH step as the balance of work between the propose and accept operations is more favourable for typical acceptance rates. Numerically we find that runtimes are reduced by around 30% for systems with  $N = 10^5$  charges, with the potential of a speedup of around  $2\times$  for larger values of  $N$ . We demonstrated numerically that our implementation will effectively scale to systems containing  $10^6$  charges whilst maintaining the expected computational complexity of  $\mathcal{O}(\log(N))$  per MH step. As the direct Ewald summation technique has a higher complexity of at least  $\mathcal{O}(\sqrt{N})$ , our implementation becomes more efficient for simulations with more than  $10^4$  particles: for  $N = 10^5$  it is about an order of magnitude faster than the current DL\_MONTE implementation.

There are several avenues for future work. One obvious shortcoming of the present implementation is the lack of parallelisation. While Monte Carlo simulations are “embarrassingly parallel”, and several Markov Chains can be generated in parallel without communications, this increases memory requirements. In our method this issue could be reduced to a certain degree by parallelising over the levels in the grid hierarchy. This is possible since the cost on each level is constant, and computations can be carried out independently, before summing the total energy in the propose stage.

Furthermore, here we have only considered single-particle moves and future work should also investigate other Monte Carlo transitions. For example, in a canonical ensemble the pressure is fixed but the volume of the simulation cell varies. In this case a proposed move could consist of a change of the system volume. In the worst case scenario, the energy change of such a proposed volume move is computed with a full solve of the electrostatic energy of the proposed state. When using an Ewald based approach this system solve incurs an  $\mathcal{O}(N^{\frac{3}{2}})$  cost per MH step, however, with multipole- or local expansion based approaches this can potentially be reduced to an  $\mathcal{O}(N)$  cost.

### CRedit authorship contribution statement

**William Robert Saunders:** Conceptualization, Data curation, Investigation, Methodology, Software, Validation, Visualization, Writing – review & editing. **James Grant:** Conceptualization, Investigation, Methodology, Supervision, Writing – review & editing. **Eike Hermann Müller:** Conceptualization, Methodology, Project administration, Supervision, Validation, Writing – original draft, Writing – review & editing.

### Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

### Acknowledgements

This research made use of the Balena High Performance Computing (HPC) Service at the University of Bath. This project has received funding from the European Union’s Horizon 2020 research and innovation programme under grant agreements No. 646176 and No. 824158.

### References

- [1] G.A. Cisneros, M. Karttunen, P. Ren, C. Sagui, Classical electrostatics for biomolecular simulations, *Chem. Rev.* 114 (1) (2013) 779–814, <https://doi.org/10.1021/cr300461d>.
- [2] M. Jorge, N.A. Seaton, Long-range interactions in Monte Carlo simulation of confined water, *Mol. Phys.* 100 (13) (2002) 2017–2023, <https://doi.org/10.1080/00268970110099585>.
- [3] C.K. Sandalci, C. Koc, S.M. Goodnick, Three-dimensional Monte Carlo device simulation with parallel multigrid solver, *Int. J. High Speed Comput.* 9 (03) (1997) 223–236, <https://doi.org/10.1142/S0129053397000143>.
- [4] P.P. Ewald, Die Berechnung optischer und elektrostatischer Gitterpotentiale, *Ann. Phys.* 369 (3) (1921) 253–287, <https://doi.org/10.1002/andp.19213690304>.
- [5] L. Greengard, V. Rokhlin, A fast algorithm for particle simulations, *J. Comput. Phys.* 73 (2) (1987) 325–348, [https://doi.org/10.1016/0021-9991\(87\)90140-9](https://doi.org/10.1016/0021-9991(87)90140-9).
- [6] L. Greengard, *The Rapid Evaluation of Potential Fields in Particle Systems*, MIT Press, 1988.
- [7] L. Greengard, V. Rokhlin, A new version of the fast multipole method for the Laplace equation in three dimensions, *Acta Numer.* 6 (1997) 229–269, <https://doi.org/10.1017/S0962492900002725>.
- [8] J.C. Crabtree, M. Molinari, S.C. Parker, J.A. Purton, Simulation of the adsorption and transport of CO<sub>2</sub> on faujasite surfaces, *J. Phys. Chem. C* 117 (42) (2013) 21778–21787, <https://doi.org/10.1021/jp4053727>.
- [9] T. Pham, B. Space, Insights into the gas adsorption mechanisms in metal–organic frameworks from classical molecular simulations, *Top. Curr. Chem.* 378 (2020) 14, <https://doi.org/10.1007/s41061-019-0276-x>.
- [10] A.D. Bruce, A.N. Jackson, G.J. Ackland, N.B. Wilding, Lattice-switch Monte Carlo method, *Phys. Rev. E* 61 (2000) 906–919, <https://doi.org/10.1103/PhysRevE.61.906>.
- [11] T.L. Underwood, G.J. Ackland, Lattice-switch Monte Carlo: the fcc–bcc problem, *J. Phys. Conf. Ser.* 640 (2015) 012030, <https://doi.org/10.1088/1742-6596/640/1/012030>.
- [12] T.A. Höft, B.K. Alpert, Fast updating multipole Coulombic potential calculation, *SIAM J. Sci. Comput.* 39 (3) (2017) A1038–A1061, <https://doi.org/10.1137/16M1096189>.
- [13] W.R. Saunders, J. Grant, E.H. Müller, A domain specific language for performance portable molecular dynamics algorithms, *Comput. Phys. Commun.* 224 (2018) 119–135, <https://doi.org/10.1016/j.cpc.2017.11.006>.
- [14] J. Purton, J.C. Crabtree, S. Parker, DL\_MONTE: a general purpose program for parallel Monte Carlo simulation, *Mol. Simul.* 39 (14–15) (2013) 1240–1252, <https://doi.org/10.1080/08927022.2013.839871>.

- [15] A.V. Brukhno, J. Grant, T.L. Underwood, K. Stratford, S.C. Parker, J.A. Purton, N.B. Wilding, DL\_MONTE: a multipurpose code for Monte Carlo simulation, *Mol. Simul.* (2019) 1–21, <https://doi.org/10.1080/08927022.2019.1569760>.
- [16] Z. Gan, Z. Xu, Efficient implementation of the Barnes-Hut octree algorithm for Monte Carlo simulations of charged systems, *Sci. China Math.* 57 (7) (2014) 1331–1340, <https://doi.org/10.1007/s11425-014-4783-5>.
- [17] D. Frenkel, B. Smit, *Understanding Molecular Simulation: From Algorithms to Applications*, vol. 1, Academic Press, San Diego/London, 2001.
- [18] U. Trottenberg, C.W. Oosterlee, A. Schuller, *Multigrid*, Academic Press, London, 2001.
- [19] J. Rottler, A.C. Maggs, A continuum,  $\mathcal{O}(N)$  Monte Carlo algorithm for charged particles, *J. Chem. Phys.* 120 (7) (2004) 3119–3129, <https://doi.org/10.1063/1.1642590>.
- [20] W.R. Saunders, J. Grant, E.H. Müller, I. Thompson, Fast electrostatic solvers for kinetic Monte Carlo simulations, *J. Comput. Phys.* (2020) 109379, <https://doi.org/10.1016/j.jcp.2020.109379>.
- [21] W.R. Saunders, J. Grant, E.H. Müller, Long range forces in a performance portable molecular dynamics framework, in: *Parallel Computing Is Everywhere*, 2018, pp. 37–46.
- [22] W.R. Saunders, J. Grant, E.H. Müller, Code and data release: Monte Carlo simulations with fast and accurate electrostatics, <https://doi.org/10.5281/zenodo.3873307>, Jun. 2020.
- [23] CCP5, DL\_POLY\_4 TEST01, [ftp://ftp.dl.ac.uk/ccp5/DL\\_POLY/DL\\_POLY\\_4.0/DATA/](ftp://ftp.dl.ac.uk/ccp5/DL_POLY/DL_POLY_4.0/DATA/), 2018. (Accessed 1 April 2018).
- [24] J. Kolafa, J.W. Perram, Cutoff errors in the Ewald summation formulae for point charge systems, *Mol. Simul.* 9 (5) (1992) 351–368, <https://doi.org/10.1080/08927029208049126>.