

A random batch Ewald method for particle systems with Coulomb interactions

Shi Jin^{*1,2}, Lei Li^{†1,2}, Zhenli Xu^{‡1,2}, and Yue Zhao^{§1}

¹School of Mathematical Sciences, Shanghai Jiao Tong University, Shanghai, 200240, P. R. China

²Institute of Natural Sciences and MOE-LSC, Shanghai Jiao Tong University, Shanghai, 200240, P. R. China

Abstract

We develop a random batch Ewald (RBE) method for molecular dynamics simulations of particle systems with long-range Coulomb interactions, which achieves an $\mathcal{O}(N)$ complexity in each step of simulating the N -body systems. The RBE method is based on the Ewald splitting for the Coulomb kernel with a random “mini-batch” type technique introduced to speed up the summation of the Fourier series for the long-range part of the splitting. Importance sampling is employed to reduce the induced force variance by taking advantage of the fast decay property of the Fourier coefficients. The stochastic approximation is unbiased with controlled variance. Analysis for bounded force fields gives some theoretic support of the method. Simulations of two typical problems of charged systems are presented to illustrate the accuracy and efficiency of the RBE method in comparison to the results from the Debye-Hückel theory and the classical Ewald summation, demonstrating that the proposed method has the attractiveness of being easy to implement with the linear scaling and is promising for many practical applications.

Key words. Ewald summation, Langevin dynamics, random batch method, stochastic differential equations

AMS subject classifications. 65C35; 82M37; 65T50

1 Introduction

Molecular dynamics simulation is among the most popular numerical methods at the molecular or atomic level to understand dynamical and equilibrium properties of many-body particle systems in many areas such as chemical physics, soft materials and biophysics [9, 18, 17]. However, the long-range interactions such as electrostatic Coulomb interactions pose a major challenge to particle simulations, as one has to take into account all pairs of interactions, leading to $\mathcal{O}(N^2)$ computational cost per iteration for naive discretizations, which is not only computationally expensive but also less accurate considering the presence of boundary conditions in the simulation box. A lot of effort in literature has been devoted to computing the long-range interactions efficiently, and widely studied methods include lattice summation methods such as particle mesh Ewald (PME) [10, 15] and particle-particle particle mesh Ewald (PPPM) [38, 13], and multipole type methods such as treecode [2, 14] and fast multipole methods (FMM) [21, 22, 50]. These methods can reduce the operations per step to $\mathcal{O}(N \log N)$ or $\mathcal{O}(N)$, and have gained big success in practice, but many problems remain

*shijin-m@sjtu.edu.cn

†leili2010@sjtu.edu.cn

‡xuzl@sjtu.edu.cn

§sjtu-15-zy@sjtu.edu.cn

to be solved as the prefactor in the linear scaling can be large, or their implementation is nontrivial, or the scalability for parallel computing is not high.

The mainstream packages [7, 49, 25] for all-atom molecular dynamics simulations mostly use Ewald-type lattice-summation algorithms which are originally proposed by Ewald [16, 11]. This type of methods split the Coulomb kernel into a rapidly decaying function in the real space and a smooth function. The cutoff scheme is introduced for the first part in the real space. The smooth part is approximated by the Fourier series expansion. The classical Ewald achieves an $\mathcal{O}(N^{3/2})$ complexity to sum up all interactions. When the cutoff radius is independent of N and the Fourier series is accelerated by the fast Fourier transform (FFT) with an interpolation to distribute charges on lattices, one obtains the PME method which achieves an $\mathcal{O}(N \log N)$ complexity. The state-of-the-art development of the Ewald-type algorithm includes an optimized choice of volumetric decomposition FFT scheme for large systems on massively parallel supercomputers [31] and efficient methods for Coulomb interactions without full periodicity [20, 42].

In this work, we propose a random batch Ewald (RBE) method of particle systems with Coulomb interactions which enables an $\mathcal{O}(N)$ Ewald method for fast molecular dynamics simulations. The “random mini-batch” idea, namely using the sum over a small random set to approximate some big summation, has its origin in the stochastic gradient descent (SGD) method [44, 6]. This type of ideas have been developed into different methods such as the stochastic gradient Langevin dynamics for Bayesian inference [48], stochastic binary interaction algorithms for the mean field swarming dynamics [1], the random batch method for interacting particle systems [29], and random-batch Monte Carlo simulations [36]. Though the specific implementations are different for different applications, these methods are intrinsically Monte Carlo methods for computing the big summation involved in the dynamics, and the convergence can be obtained due to a time averaging effect [29], obeying the law of large numbers in time.

The RBE method uses the same idea of random minibatch. The new design being different from previous work is that the minibatch is built into the Ewald summation and sampled from the Fourier space. We take a cutoff radius in the real space such that the particles within the radius is of order one, and sample $p = \mathcal{O}(1)$ frequencies in the Fourier expansion of the smooth part of the Ewald splitting. These p frequencies are chosen into the minibatch for the force calculation. The advantages of this approach are threefold. First, the short-range part of the force remains exact and thus the variance of the force can be significantly reduced. Second, the short-range repulsive force due to the van der Waals interaction can be naturally introduced to avoid unphysical configuration. Third and the most important, the importance sampling can be used in the Fourier space in building the minibatch such that the force variance can be further reduced. These strategies combined lead to a simple and efficient RBE method for molecular dynamics, as shown in our numerical examples for calculating typical properties of electrolytes.

The rest of the paper is organized as follows. Section 2 is devoted to an introduction to the setup and the classical Ewald summation, which forms the basis of our method. In Section 3, we introduce the methodology of the RBE and give its detailed implementation. We also provide some theoretic evidence on why the method works and can be efficient. In Section 4, we test the RBE on two typical problems to validate the method. Conclusions are given in Section 5.

2 Overview of the Ewald summation

In this section, we introduce the setup and notations to be used later. Then, we give a brief review of the classical Ewald summation [16, 11].

To approximate electrostatic interactions between charges in an electrolyte of big extent, one often uses a simulation box with periodic boundary conditions (PBCs) [18] to mimic the bulk environment of the electrolyte. Without loss of generality, we consider a cubic box with the edge length L so that the volume of the box is given by $V = L^3$. During the simulation, one calculates interactions of N numerical particles (not necessarily the physical particles)

inside the box with positions \mathbf{r}_i and charge q_i ($1 \leq i \leq N$) satisfying the electroneutrality condition

$$\sum_{i=1}^N q_i = 0. \quad (2.1)$$

Then, one evolves the particles according to Newton's equations

$$\begin{aligned} d\mathbf{r}_i &= \mathbf{v}_i dt, \\ m_i d\mathbf{v}_i &= \mathbf{F}_i(\{\mathbf{r}_i\}) dt + d\boldsymbol{\eta}_i, \end{aligned} \quad (2.2)$$

where $d\boldsymbol{\eta}_i$ represents the coupling with the environment (heat bath) (see Section 3.3).

The forces are computed using $\mathbf{F}_i = -\nabla_{\mathbf{r}_i} U$, where U is the potential energy of the system. Let $\mathbf{r}_{ij} := \mathbf{r}_j - \mathbf{r}_i$ and $r_{ij} = |\mathbf{r}_{ij}|$ be the distance. The potential energy of the system due to Coulomb interactions with PBCs can be written as

$$U = \frac{1}{2} \sum_{\mathbf{n}}' \sum_{i,j=1}^N q_i q_j \frac{1}{|\mathbf{r}_{ij} + \mathbf{n}L|}, \quad (2.3)$$

where $\mathbf{n} \in \mathbb{Z}^3$ ranges over the three-dimensional integer vectors and \sum' is defined such that $\mathbf{n} = 0$ is not included when $i = j$. Due to the long-range nature of the Coulomb potential, this series converges conditionally. Hence, directly computing the interaction energy (2.3) and the corresponding interaction forces using a cutoff approach is less accurate, and one shall introduce more advanced techniques to sum up the infinite series.

The idea of the classical Ewald summation is to separate the series into long-range smooth parts and short-range singular parts. The conditional convergence due to the long-range, but smooth, parts can be dealt with from the Fourier side. To describe the details, one first introduces the error function

$$\text{erf}(x) := \frac{2}{\sqrt{\pi}} \int_0^x \exp(-u^2) du$$

and the error complementary function is $\text{erfc}(x) := 1 - \text{erf}(x)$. Clearly, the Coulomb kernel can be written as

$$\frac{1}{r} = \frac{\text{erf}(\sqrt{\alpha}r)}{r} + \frac{\text{erfc}(\sqrt{\alpha}r)}{r}$$

for any positive constant α , and the potential energy (2.3) can be decomposed as $U := U_1 + U_2$ with

$$U_1 = \frac{1}{2} \sum_{\mathbf{n}}' \sum_{i,j} q_i q_j \frac{\text{erf}(\sqrt{\alpha}|\mathbf{r}_{ij} + \mathbf{n}L|)}{|\mathbf{r}_{ij} + \mathbf{n}L|}, \quad (2.4)$$

$$U_2 = \frac{1}{2} \sum_{\mathbf{n}}' \sum_{i,j} q_i q_j \frac{\text{erfc}(\sqrt{\alpha}|\mathbf{r}_{ij} + \mathbf{n}L|)}{|\mathbf{r}_{ij} + \mathbf{n}L|}. \quad (2.5)$$

The sum in U_2 now converges absolutely and rapidly, and one can truncate it to simplify the computation. The sum in U_1 still converges conditionally in spite of the charge neutrality condition, but since the kernel is smooth, the summation can be treated nicely in the Fourier domain. Define the Fourier transform:

$$\tilde{f}(\mathbf{k}) := \int_{\Omega} f(\mathbf{r}) e^{-i\mathbf{k} \cdot \mathbf{r}} d\mathbf{r},$$

with $\mathbf{k} = 2\pi\mathbf{m}/L$ and $\mathbf{m} \in \mathbb{Z}^3$. The inverse transform gives $f(\mathbf{r}) = (1/V) \sum_{\mathbf{k}} \tilde{f}(\mathbf{k}) e^{i\mathbf{k} \cdot \mathbf{r}}$. Then, U_1 is expressed as (see [18, Chap. 12]),

$$U_1 = \frac{2\pi}{V} \sum_{\mathbf{k} \neq 0} \frac{1}{k^2} |\rho(\mathbf{k})|^2 e^{-k^2/4\alpha} - \sqrt{\frac{\alpha}{\pi}} \sum_{i=1}^N q_i^2, \quad (2.6)$$

where $k = |\mathbf{k}|$ and $\rho(\mathbf{k})$ is given by

$$\rho(\mathbf{k}) := \sum_{i=1}^N q_i e^{i\mathbf{k} \cdot \mathbf{r}_i}, \quad (2.7)$$

which can be viewed as the conjugate of the Fourier transform of the charge density. The divergent $\mathbf{k} = 0$ term is usually neglected in simulations to represent that the periodic system is embedded in a conducting medium which is essential for simulating ionic systems.

By inspection of the expressions above, we may take truncations in both the real and frequency domains. In particular, picking the real space and the reciprocal space cutoffs

$$r_c := s/\sqrt{\alpha}, \quad k_c := 2s\sqrt{\alpha}, \quad (2.8)$$

one then has (see [32])

$$\begin{aligned} U = & \frac{2\pi}{V} \sum_{0 < k \leq k_c} \frac{1}{k^2} |\rho(\mathbf{k})|^2 e^{-k^2/4\alpha} - \sqrt{\frac{\alpha}{\pi}} \sum_{i=1}^N q_i^2 + \\ & \frac{1}{2} \sum_{|\mathbf{r}_{ij} + \mathbf{n}L| \leq r_c} q_i q_j \frac{\text{erfc}(\sqrt{\alpha}|\mathbf{r}_{ij} + \mathbf{n}L|)}{|\mathbf{r}_{ij} + \mathbf{n}L|} + \mathcal{O}\left(\frac{Qe^{-s^2}}{s^2} \left(\frac{s}{\sqrt{\alpha}L^3}\right)^{\frac{1}{2}}\right) =: \widetilde{U}_1 + \widetilde{U}_2 + \mathcal{O}(\cdot), \end{aligned} \quad (2.9)$$

where $Q := \sum_{i=1}^N q_i^2$, \widetilde{U}_1 is defined by the sum of the first two terms, and \widetilde{U}_2 corresponds to the third term. The density of particles in the real space ρ_r and density of frequencies ρ_f in the reciprocal space are given respectively by

$$\rho_r = \frac{N}{L^3}, \quad \text{and} \quad \rho_f = \left(\frac{L}{2\pi}\right)^3. \quad (2.10)$$

The number of interacting particles to be considered for a given particle is thus

$$N_r := \frac{4\pi}{3} r_c^3 \rho_r = \frac{4\pi s^3 N}{3\sqrt{\alpha}^3 L^3},$$

yielding total pairs $N_p = (4\pi/3)s^3 N^2/(\sqrt{\alpha}L)^3$. The number of frequencies to be considered is $N_f = (4\pi/3)k_c^3 \rho_f = (4/3\pi^2)(sL\sqrt{\alpha})^3$. The total work to compute \widetilde{U}_1 is thus $T_f \sim N_f N$ since the computation of $\rho(\mathbf{k})$ needs $\mathcal{O}(N)$ operations. The total work to compute \widetilde{U}_2 is $T_r \sim N_p$. In the usual Ewald summation, one needs to balance these two parts of works, thus $N_p \sim N_f N$. Hence $\sqrt{\alpha} \sim N^{1/6}/L$ is chosen to balance the costs between the real and frequency domains. This then yields the total number of pairs $N_p = \mathcal{O}(N^{3/2})$, and the number of frequencies to be considered is given by $N_f = \mathcal{O}(N^{1/2})$ so that the complexity in the frequency part is $T_f = \mathcal{O}(N^{3/2})$. Hence, the total complexity per iteration is $\mathcal{O}(N^{3/2})$ for the energy computation.

The computation of force can be done directly using

$$\begin{aligned} \mathbf{F}_i = -\nabla_{\mathbf{r}_i} U = & -\sum_{\mathbf{k} \neq 0} \frac{4\pi q_i \mathbf{k}}{V k^2} e^{-k^2/(4\alpha)} \text{Im}(e^{-i\mathbf{k} \cdot \mathbf{r}_i} \rho(\mathbf{k})) \\ & - q_i \sum_{j, \mathbf{n}}' q_j G(|\mathbf{r}_{ij} + \mathbf{n}L|) \frac{\mathbf{r}_{ij} + \mathbf{n}L}{|\mathbf{r}_{ij} + \mathbf{n}L|} =: \mathbf{F}_{i,1} + \mathbf{F}_{i,2}, \end{aligned} \quad (2.11)$$

where we recall $\mathbf{r}_{ij} = \mathbf{r}_j - \mathbf{r}_i$, pointing towards particle j , and

$$G(r) := \frac{\text{erfc}(\sqrt{\alpha}r)}{r^2} + \frac{2\sqrt{\alpha}e^{-\alpha r^2}}{\sqrt{\pi}r}.$$

Note that the force $\mathbf{F}_{i,1}$ is bounded for small \mathbf{k} . In fact, $k \geq 2\pi/L$, so Vk is not small.

Again, we are going to take the truncations as shown in Eq. (2.8). With the choice $\sqrt{\alpha} \sim N^{1/6}/L$, there are $N_f = \mathcal{O}(N^{1/2})$ frequencies to consider. Note the $\rho(\mathbf{k})$ computed can be used for all i , so the complexity for computing the forces $\mathbf{F}_{i,1}$ for all $i = 1, \dots, N$ is $\mathcal{O}(N^{3/2})$. Since there are $\mathcal{O}(N^{1/2})$ particles to consider for each i , the complexity for computing the forces $\mathbf{F}_{i,2}$ for all $i = 1, \dots, N$ is also $\mathcal{O}(N^{3/2})$. The total complexity per iteration is thus $\mathcal{O}(N^{3/2})$.

It is remarked that the PPPM [38, 13] is a fast way to compute the Ewald sum using the FFT. The PPPM chooses parameter α such that $\sqrt{\alpha} \sim N^{1/3}/L$. Using the cutoffs (2.8), the number of frequencies to be considered and the number of particles in real space for a given particle are given respectively by

$$N_f = \mathcal{O}(N), \text{ and } N_r = \mathcal{O}(1). \quad (2.12)$$

Hence, to compute the force, the complexity corresponding to the summation of all frequency components is $\mathcal{O}(N \cdot N_f) = \mathcal{O}(N^2)$ in the direct Ewald summation. To speed up the summation in the Fourier space, one meshes the simulation box, assigns charges on the grid by interpolation and then takes advantage of the FFT to obtain $\rho(\mathbf{k})$ so that the potential can be computed with $\mathcal{O}(N \log N)$ complexity. The potential and forces at the particle locations are then obtained by further interpolation and some numerical difference schemes. Hence, the complexity per iteration is $\mathcal{O}(N \log N)$ for the PPPM.

3 The random batch Ewald

We now aim to develop the stochastic molecular dynamics using the idea of random mini-batch. The implementation of mini-batch (i.e., finding suitable cheap unbiased stochastic approximation) depends on the specific applications. For interacting particle systems in [29], the strategy is the random grouping of particles. By inspection of the Ewald summation ((2.9) and (2.11)), we found that $e^{-k^2/(4\alpha)}$ is summable so that it can be normalized to form a probability distribution. Hence, this allows us to do the importance sampling in the Fourier space. This leads to a random batch strategy for the simulations of molecular dynamics.

3.1 The algorithm

Let us consider the factor $e^{-k^2/(4\alpha)}$ within the first term in Eq. (2.11). Denote the sum of such factors by

$$S := \sum_{\mathbf{k} \neq 0} e^{-k^2/(4\alpha)} = H^3 - 1, \quad (3.1)$$

where

$$H := \sum_{m \in \mathbb{Z}} e^{-\pi^2 m^2 / (\alpha L^2)} = \sqrt{\frac{\alpha L^2}{\pi}} \sum_{m=-\infty}^{\infty} e^{-\alpha m^2 L^2}, \quad (3.2)$$

Here, S is the sum for all three-dimensional vectors \mathbf{k} except 0. The number H is the one for one-dimensional sum. The second equality in Eq. (3.2) is obtained by the Poisson summation formula [3, 12]. Eq. (3.2) can then be simply truncated at $m = \pm 1$ to obtain an approximation,

$$H \approx \sqrt{\frac{\alpha L^2}{\pi}} (1 + 2e^{-\alpha L^2})$$

using the rapid convergence of the series as typical setup in our simulations holds $\alpha L^2 \gg 1$. One can improve the accuracy by using more terms if needed. Then, we have the exact expression for the probability

$$\mathcal{P}_{\mathbf{k}} := S^{-1} e^{-k^2/(4\alpha)}, \quad (3.3)$$

which, with $\mathbf{k} \neq 0$, is a discrete Gaussian distribution and can be sampled efficiently as detailed below.

We apply the Metropolis-Hastings (MH) algorithm (see [24] for details) to sample from the discrete distribution

$$\mathcal{P}(m) \sim H^{-1} e^{-(2\pi m/L)^2/4\alpha}. \quad (3.4)$$

Doing this sampling procedure for three independent experiments will generate the components k_i ($i = 1, 2, 3$) of \mathbf{k} as $\mathbf{k} = 2\pi\mathbf{m}/L$. The samples with $k_1 = k_2 = k_3 = 0$ will be discarded. In the MH procedure, the proposal m^* is generated by first drawing $x^* \sim \mathcal{N}(0, \alpha L^2/(2\pi^2))$, the normal distribution with mean zero and variance $\alpha L^2/(2\pi^2)$, and one then sets $m^* = \text{round}(x^*)$, which is accepted with probability $q(m^*|m)$ in the MH algorithm, and clearly the probability is given by the following explicit expression,

$$\begin{aligned} q(m^*|m) &= \int_{m^*-1/2}^{m^*+1/2} \sqrt{\frac{\pi}{\alpha L^2}} e^{-\pi^2 x^2/\alpha L^2} dx \\ &= \begin{cases} \text{erf}\left(\frac{1/2}{\sqrt{\alpha L^2/\pi^2}}\right) & m^* = 0 \\ \frac{1}{2} \left[\text{erf}\left(\frac{|m^*|+1/2}{\sqrt{\alpha L^2/\pi^2}}\right) - \text{erf}\left(\frac{|m^*|-1/2}{\sqrt{\alpha L^2/\pi^2}}\right) \right] & m^* \neq 0. \end{cases} \end{aligned} \quad (3.5)$$

Since $\mathcal{P}(m^*) \approx q(m^*|m)$, the acceptance rate is very high, which leads to small errors in this sampling procedure. In practical implementation, one can precompute $q(m^*|m) \equiv \bar{q}(m^*)$ for a large enough range of m^* values to speed up the sampling procedure.

We now consider the calculation of the forces in Eq. (2.11) using the random mini-batch strategy. At each step, one picks a batch size p , which is of $\mathcal{O}(1)$, and draws p frequencies \mathbf{k}_ℓ , $1 \leq \ell \leq p$, i.i.d. from the discrete distribution $\mathcal{P}_{\mathbf{k}}$ by the MH sampling method above. The force $\mathbf{F}_{i,1}$ in (2.11) is then approximated by the following random variable:

$$\mathbf{F}_{i,1} \approx \mathbf{F}_{i,1}^* := - \sum_{\ell=1}^p \frac{S}{p} \frac{4\pi \mathbf{k}_\ell q_i}{V k_\ell^2} \text{Im}(e^{-i\mathbf{k}_\ell \cdot \mathbf{r}_i} \rho(\mathbf{k}_\ell)). \quad (3.6)$$

In the molecular dynamics simulations, we use this stochastic force $\mathbf{F}_{i,1}^*$ which is unbiased for the force calculation to replace $\mathbf{F}_{i,1}$. The resulted molecular dynamics is a much cheaper version of the Ewald summation, and we call this stochastic method the Random Batch Ewald (RBE).

Of course, we need to update the p samples after each time iteration. Suppose we have picked a step size Δt and defined the time grid $t_n = n\Delta t$. Then, we renew the batch of frequencies at each time grid point t_n . In real simulations, one will also add the van der Waals potential such as the Lennard-Jones potential so that positive and negative charges will not merge. The force on each particle is then calculated by summing up the contributions of real and Fourier parts, and the Lennard-Jones force (and other forces such as chemical bonds if any). Then, one integrates Newton's equations (2.2) to obtain the position and velocity of the particle in the next time step. Algorithm 1 shows one possible such molecular dynamics method using the RBE with some appropriate thermostat coupled to a heat bath (see Section 3.3 for discussions).

Algorithm 1 (Random-batch Ewald)

- 1: Choose α , r_c and k_c (the cutoffs in real and Fourier spaces respectively), Δt , and batch size p . Initialize the positions and velocities of charges $\mathbf{r}_i^0, \mathbf{v}_i^0$ for $1 \leq i \leq N$.
 - 2: Sample sufficient number of $\mathbf{k} \sim e^{-k^2/(4\alpha)}$, $\mathbf{k} \neq 0$ by the MH procedure to form a set \mathcal{K} .
 - 3: **for** n in $1 : N$ **do**
 - 4: Integrate Newton's equations (2.2) for time Δt with appropriate integration scheme and some appropriate thermostat. The Fourier parts of the Coulomb forces are computed using RBE force (3.6) with the p frequencies chosen from \mathcal{K} in order.
 - 5: **end for**
-

In the case of the leapfrog scheme (equivalent to velocity-Verlet method) and the Andersen thermostat, the loop step in Algorithm 1 is as follows.

- (1) Choose p frequencies from \mathcal{K} without replacement; calculate real and Fourier parts of the electrostatic Coulomb force using RBE (3.6), and other forces such as the Lennard-Jones forces.
- (2) Update the position and velocity of each particle using the following scheme for $n \geq 1$

$$\begin{aligned}\mathbf{v}_i^{n+1/2} &= \mathbf{v}_i^{n-1/2} + \frac{1}{m_i} \mathbf{F}_i^n \times \Delta t, \\ \mathbf{r}_i^{n+1} &= \mathbf{r}_i^n + \mathbf{v}_i^{n+1/2} \times \Delta t.\end{aligned}$$

(Here, $\mathbf{v}_i^{1/2}$ can be obtained via the Euler scheme.)

- (3) Update the velocity $\mathbf{v}_i^{n+1/2}$ of each particle with probability $\nu \Delta t$ by resampling \mathbf{v}_i from the normal distribution $\mathcal{N}(0, I_3 T / m_i)$.

We now analyze the complexity of the RBE method per time step. Similar to the strategy in the PPPM, we may choose α such that the time cost in real space is cheap and the computation in the Fourier space is then accelerated. Compared to the PPPM, the only difference is that the PPPM uses FFT and the RBE uses random mini-batch idea to speed up the computation in the Fourier space. Hence, we make the same choice,

$$\sqrt{\alpha} \sim \frac{N^{1/3}}{L} = \rho_r^{1/3},$$

which is inverse of the average distance between two numerical particles. The complexity for the real space part is $\mathcal{O}(N \cdot N_r) = \mathcal{O}(N)$. Using the random batch approximation (3.6) which is a certain Monte Carlo method for approximating the force, the number of frequencies to be considered is then reduced to

$$N_f = \mathcal{O}(p). \quad (3.7)$$

If we choose *the same batch* of frequencies for all forces (3.6) (i.e., using the same \mathbf{k}_ℓ , $1 \leq \ell \leq p$ for all $\mathbf{F}_{i,1}^*$) in the same time step, since the computed numbers $\rho(\mathbf{k}_\ell)$ can be used for all particles, the complexity per iteration for the frequency part is reduced to $N_f N = \mathcal{O}(pN)$. This implies that the RBE method has linear complexity per time step if one chooses $p = \mathcal{O}(1)$.

3.2 Consistency and stability

In this subsection, we provide some theoretic evidence for the consistency and stability of the RBE algorithm in order to demonstrate its validity.

According to Eq. (3.3), we find that the long wave (low frequency) modes are more likely to be chosen in the random approximation. Since the long wave modes are more important for the periodic effects, this importance sampling strategy could be more effective compared with the uniform sampling across the modes considered. This importance sampling strategy could also possibly reduce the variance so that the random method is more stable. We now provide some theoretic evidence to explain why this method works.

We define the fluctuation in the random batch approximation for the Fourier part of the force on particle i by,

$$\chi_i := \mathbf{F}_{i,1}^* - \mathbf{F}_{i,1}. \quad (3.8)$$

The expectation and variance of the fluctuation can be obtained by direct calculation, which is given by Proposition 1.

Proposition 1. *The fluctuation in force χ_i has zero expectation,*

$$\mathbb{E}\chi_i = 0, \quad (3.9)$$

and that the variance is,

$$\mathbb{E}|\chi_i|^2 = \frac{1}{p} \left(\sum_{\mathbf{k} \neq 0} \frac{(4\pi q_i)^2 S}{V^2 k^2} e^{-k^2/(4\alpha)} |\text{Im}(e^{-i\mathbf{k} \cdot \mathbf{r}_i} \rho(\mathbf{k}))|^2 - |\mathbf{F}_{i,1}|^2 \right). \quad (3.10)$$

The first claim in Proposition 1 implies that the random approximation is consistent or unbiased,

$$\mathbb{E}\mathbf{F}_{i,1}^* = \mathbf{F}_{i,1}, \quad (3.11)$$

where \mathbb{E} means expectation in probability theory (or the ensemble average in the physics community). The second claim says that

$$\mathbb{E}|\chi_i|^2 \lesssim \frac{1}{p} \frac{S}{V} U_1 = \frac{1}{p} \rho_r U_1.$$

If the density $\rho_r = N/V$ is not very big, we expect our stochastic algorithm to work well. Since for $k \gg \sqrt{\alpha}$, the factor $e^{-k^2/(4\alpha)}$ is very small and contributes little to the variance in (3.10). Let us now consider the terms with $k \lesssim \sqrt{\alpha}$. In the dilute solution regime where the Debye–Hückel (DH) theory (see [34], and also Appendix A) is applicable, we expect that $|\text{Im}(e^{-i\mathbf{k} \cdot \mathbf{r}_i} \rho(\mathbf{k}))| \approx 0$. That means the variance is nearly zero. Of course, due to the deviation from the Debye–Hückel theory by thermal fluctuation, this cannot be zero. We expect that $|\text{Im}(e^{-i\mathbf{k} \cdot \mathbf{r}_i} \rho(\mathbf{k}))|$ does not change too much by the thermal fluctuation for $k \ll a^{-1}$ where a is the diameter of the ions (see Appendices A–B). Clearly, if $\sqrt{\alpha} \ll a^{-1}$, the frequencies we consider then satisfy $k \ll a^{-1}$. We then can safely bound

$$|\text{Im}(e^{-i\mathbf{k} \cdot \mathbf{r}_i} \rho(\mathbf{k}))| \leq C.$$

In Appendix B, it is computed under this assumption that

$$\mathbb{E}|\chi_i|^2 \lesssim \frac{1}{p} \rho_r^{4/3}, \quad (3.12)$$

which verifies that the variance of the random force is indeed controlled if the density is not big.

The following result indicates that random mini-batch methods can be valid for capturing the finite time dynamics (we take the Langevin thermostat for illustration and see Section 3.3 for discussions).

Theorem 1. *Let $(\mathbf{r}_i, \mathbf{v}_i)$ be the solutions to*

$$\begin{aligned} d\mathbf{r}_i &= \mathbf{v}_i dt, \\ m_i d\mathbf{v}_i &= [\mathbf{F}_i(\{\mathbf{r}_j\}) - \gamma \mathbf{v}_i] dt + \sqrt{2\gamma/\beta} d\mathbf{W}_i, \end{aligned}$$

where $\{\mathbf{W}_i\}$ are i.i.d. Wiener processes. Let $(\tilde{\mathbf{r}}_i, \tilde{\mathbf{v}}_i)$ be the solutions to

$$\begin{aligned} d\tilde{\mathbf{r}}_i &= \tilde{\mathbf{v}}_i dt, \\ m_i d\tilde{\mathbf{v}}_i &= [\mathbf{F}_i(\{\tilde{\mathbf{r}}_j\}) + \chi_i - \gamma \tilde{\mathbf{v}}_i] dt + \sqrt{2\gamma/\beta} d\mathbf{W}_i, \end{aligned}$$

with the same initial values as $(\mathbf{r}_i, \mathbf{v}_i)$. Suppose that the masses m_i 's are bounded uniformly from above and below. If the forces \mathbf{F}_i are bounded and Lipschitz and $\mathbb{E}\chi_i = 0$, then for any $T > 0$, there exists $C(T) > 0$ such that

$$\mathbb{E} \left[\frac{1}{N} \sum_i (|\mathbf{r}_i - \tilde{\mathbf{r}}_i|^2 + |\mathbf{v}_i - \tilde{\mathbf{v}}_i|^2) \right] \leq C(N, T) \sqrt{\Lambda \Delta t},$$

where Λ is an upper bound for $\max_i (\mathbb{E}|\chi_i|^2)$.

Similar proofs for interacting particle systems can be found in [30, 35, 36], and we omit the proof for the claims here. The constant $C(N, T)$ can be made independent of N in the mean field regime [30]. Clearly, due to the assumption that \mathbf{F}_i is bounded and Lipschitz, the claims above are not helpful for our problem. Anyhow, it can give us some insight how random batch type methods work. Clearly, for a given configuration, a force computed using the RBE is a random approximation to the true force. A single-step evaluation of such random force definitely has no accuracy compared to the true force. The intuition why such methods work is that the effects of random forces accumulate in time. Since the random forces are unbiased, the random errors will roughly cancel out over time. This “law of large numbers” type mechanism in time then makes the random method work. The error bound above is the square root of variance multiplied by Δt , which is the typical error bound given by central limit theorem. Hence, our method is not aiming at computing the forces correctly for a fixed configuration. Instead, we attempt to obtain the evolution of the configurations and the equilibrium distribution with an acceptable error control. We use the RBE method only to speed up MD simulations and obtain configurations, and then use these configurations to compute the true energies, stress tensor (and pressure) using their definitions, without random batch approximation.

A question that may arise is whether one should wait for too many iterations before the “law of large number” mechanism takes effects to capture the long time properties (i.e. whether the random batch type methods will delay the mixing time for the equilibria too much). In [30, 36], it has been shown that when some external confining fields are present, the mixing time for convergence to the global equilibrium with random batch is roughly the same as the one without random batch, as the error controls are uniform in time. When there are no helping external fields such as the cases we are considering here in a periodic box, whether random batch will delay the convergence to the thermal equilibrium is still a theoretically open question. However, when heat bath is present, if the number of particles or modes is statistically large so that a few of them can capture the significant statistical properties, the few chosen representatives may give the correct statistical properties and the random batch methods may capture the correct macroscopic quantities without looping for too many iterations. Hence, we believe the RBE method can capture the long time statistical properties for the many-body systems in contact with heat bath, without increasing the iterations of simulation too much.

As we have seen, the variance of the fluctuation is always multiplied by the step size Δt in the error estimates: $\sqrt{\Lambda \Delta t}$ for the error of trajectories or $\Lambda \Delta t$ for the distributions (see [30] for the weak error estimates regarding first order systems). Hence, the variance somehow measures the stability of the random methods and the boundedness of $\mathbb{E}|\chi_i|^2$ is important for the convergence of the random algorithms. Though the variance is controlled for the RBE, rigorous proof for this method is challenging as the field $\mathbf{F}_{i,2}$ is singular. Building in van der Waals potential into the system can prevent the particles getting too close so the singularity of $\mathbf{F}_{i,2}$ might not be seen, but the rigorous justification of convergence could still be very hard. We will leave the rigorous mathematical analysis for future exploration.

3.3 Discussion on the thermostats

To couple with the heat bath so that the temperature is preserved near the desired value, typical ways include the Andersen thermostat and the Langevin thermostat. Another thermostat used in molecular dynamics in a deterministic approach is the famous Nosé-Hoover thermostat [18, Chap. 6].

In the Andersen thermostat, one assumes the collision frequency between the particle and the heat bath is ν . Then, the time between two collisions for a particular particle satisfies the exponential distribution. Hence, the probability that a particle does not collide with the heat bath during $[t_n, t_n + \Delta t)$ is $1 - \exp(-\nu \Delta t) \approx \nu \Delta t$. Hence, in the Andersen thermostat, at each time grid t_n , for each particle one resets the velocity with probability $\nu \Delta t$. The new velocity is sampled from the Maxwell distribution with temperature T . In Eq. (2.2), $d\eta_i$ corresponds to such resetting noise to model the collision. It is this new velocity that guarantees the correct temperature. Of course, the larger ν is, the better the

temperature can be kept around the desired value. However, too large ν value will bring some unphysical effects [41, 27, 18].

In the Langevin thermostat, the interaction with heat bath is added into the equation directly:

$$m_i d\mathbf{v}_i = (\mathbf{F}_i - \gamma \mathbf{v}_i) dt + \sqrt{2\gamma/\beta} d\mathbf{W}_i.$$

In other words, in (2.2), $d\boldsymbol{\eta}_i = -\gamma \mathbf{v}_i dt + \sqrt{2\gamma/\beta} d\mathbf{W}_i$. The term $-\gamma \mathbf{v}_i dt$ is the friction and $\sqrt{2\gamma/\beta} d\mathbf{W}_i$ is the thermal noise or the fluctuation, both arising from the collision with the heat bath. The fluctuation-dissipation relation requires the strength of the noise to be $\sqrt{2\gamma/\beta}$ so that the system can tend to the correct temperature $T = \beta^{-1}$ (the Boltzmann constant k_B is taken to be 1 for the reduced units). As in the Andersen thermostat, increasing γ can keep the temperature of the system around T better. However, since the temperature enters in through the dynamics, the Andersen thermostat seems to behave better for the temperature control than the Langevin dynamics.

The Nosé-Hoover thermostat uses a Hamiltonian for an extended system of N particles plus an additional coordinate s ([41, 27]):

$$\mathcal{H}_{\text{NH}} = \sum_{i=1}^N \frac{|\mathbf{p}_i|^2}{2m_i s^2} + U(\{\mathbf{r}_i\}) + \frac{p_s^2}{2Q} + L \frac{\ln s}{\beta}.$$

Here, $\mathbf{p}_i = m_i \mathbf{v}_i$ is the momentum of the i th particle. The microcanonical ensemble corresponding to this Hamiltonian reduces to the canonical ensemble for the real variables $\mathbf{p}'_i = \mathbf{p}_i/s$. Hence, one may run the following deterministic ODEs, which are the Hamilton ODEs under \mathcal{H}_{NH} in terms of the so-called real variables,

$$\begin{aligned} \dot{\mathbf{r}}_i &= \frac{\mathbf{p}_i}{m_i}, \\ \dot{\mathbf{p}}_i &= -\nabla_{\mathbf{r}_i} U - \xi \mathbf{p}_i, \\ \dot{\xi} &= \frac{1}{Q} \left(\sum_{i=1}^N \frac{|\mathbf{p}_i|^2}{m_i} - \frac{3N}{\beta} \right). \end{aligned}$$

The time average of the desired quantities such as those in (4.3) will be the correct canonical ensemble average. As one can see, when the temperature of the system, defined by $\sum_{i=1}^N m_i |\mathbf{v}_i|^2 / (3Nk_B)$, is different from T , the extra term $-\xi \mathbf{p}_i$ will drive the system back to temperature T , and thus it may give better behaviors for controlling the temperature.

As we have seen, the random batch approaches will bring in extra variance term. Hence, there is numerical heating effect that increases the temperature by $\Lambda \Delta t$. Due to this reason, the RBE is not suitable for long time simulation under NVE ensemble if without an appropriate conservation scheme for time integration, but it should be good for NVT (as we do in this paper) and other simulations with thermostats. To reduce this artificial temperature, one may on one hand reduce Λ by using bigger batch size or carefully designed importance samplings. The RBE proposed in this work is an importance sampling approach. Also, one may mimic the simulated annealing idea [46, 26, 28] to decrease Δt , which has also been used in the stochastic gradient Langevin dynamics. On the other hand, using suitable thermostat may drive the temperature back to T better. In principle, the Nosé-Hoover is the most effective for preserving temperature. If the frequency ν is chosen suitably in the Andersen thermostat, the temperature can be preserved well too. As we see below in Section 4, if ν and batch size p are slightly bigger and the simulated annealing approach is used, the Andersen thermostat is already enough for the numerical examples we consider. Due to the simplicity, we adopt the Andersen thermostat in this paper to illustrate that the RBE works, while leaving the Nosé-Hoover thermostat for our future development for large systems.

4 Application examples

In this section, we consider two typical application examples to validate the accuracy and efficiency of the proposed method. The first example is the charge distribution in terms

of charge-charge correlation functions in an electrolyte solution with the primitive model where the DH theory can be used to provide a theoretical prediction. The second example is a much harder example with many different species of ions (including a macroion) where charge reversal phenomenon for electric double layer near the surface of a colloidal particle is studied. Both examples indicate that the proposed method is effective and efficient. The calculations are performed in a Linux system with Intel Xeon Scalable Cascade Lake 6248@2.5GHz, 1 CPU core and 4 GB memory.

4.1 Charge correlation functions in electrolyte

In this example, we consider a pure electrolyte monovalent binary ions. The primitive model of the electrolyte is employed, which describes the solvent as mobile ions of uniform sizes embedded in a medium of constant permittivity under a given temperature. The total potential energy of the system is composed of the Coulomb interactions and the short-range van der Waals interaction. The latter is modeled by the shifted Lennard-Jones potential expressed as:

$$V_{\text{LJ}}(r) = \begin{cases} 4\epsilon \left[\left(\frac{\sigma}{r - r_{\text{off}}} \right)^{12} - \left(\frac{\sigma}{r - r_{\text{off}}} \right)^6 \right] + V_{\text{shift}}, & \text{if } r - r_{\text{off}} < R_c \\ 0, & \text{otherwise,} \end{cases} \quad (4.1)$$

where $r_{\text{off}} = (d_i + d_j)/2 - \sigma$ and d_i and d_j are the diameters of two particles respectively, and σ is a positive constant. V_{shift} is taken such that the potential becomes zero when $r - r_{\text{off}} = R_c$.

All the quantities are provided in reduced units (see [18, sec. 3.2]). The diameter of each ion is chosen as $d_i \equiv 0.2$, the reduced temperature is $T = 1.0$ and the reduced dielectric constant is $\epsilon = 1/4\pi$ so that the electric potential of a charge q is given by $\phi(r) = q/r$. For the Lennard-Jones potential in this example, we choose the parameters as $\sigma = 0.2$, $R_c = 4.0$ and $\epsilon = 1$. We fix the particle density to be constant $N/L^3 = 0.3$. Correspondingly, the inverse Debye length in the Debye-Hückel theory (see Appendix A) is $\kappa \approx 1.9416$. We run molecular dynamics simulations to prepare the configuration samples and by taking average of these samples to obtain the charge distribution in terms of charge-charge correlation functions,

$$\rho(r) = g_{++}(r) - g_{+-}(r), \quad (4.2)$$

where g_{++} and g_{+-} are cation-cation and cation-anion pair correlation functions between ions. By the Debye-Hückel theory, the radial distribution of net charge satisfies the following linear relation,

$$\ln(r|\rho(r)|) \approx -1.9416r - 1.1437, \quad r \gg 0.2.$$

Here, due to the setting of the Lennard-Jones potential, we roughly have the parameter a in Appendix A as

$$a = \frac{1}{2}(d_1 + d_2) - \sigma + \sigma = 0.2,$$

and the formula above should be accurate for $r \gg a$.

	α	r_c	n_c	steps	Time (s)
Ewald	0.12	8.0	7	1e6	6067
PPPM	0.55	4.0	15	1e6	3120
RBE	0.55	4.0	$p = 10$	1e6	1267

Table 1: Parameters and computational time for the Ewald, PPPM and RBE results with $N = 300$. The RBE samples from all frequencies and it shows p value in the n_c column.

In the first numerical experiment, we take the length of the periodic box to be $L = 10.0$, and the number of monovalent ions $N = 300$ so that $n_{+1} = n_{-1} = 150$. The Andersen

thermostat is adopted with frequency $\nu = 3$. The parameters are chosen as in Table 1, where the n_c column for RBE lists batch size p as no frequency cutoff is introduced. The parameters are chosen so that the estimated relative force errors for the Ewald method and PPPM are about 10^{-4} by [32] (the parameters are set automatically in LAMMPS software). As discussed already, we choose the same α value for the RBE as that in the PPPM. The batch size p in the RBE is chosen through a convergence test and $p = 10$ gives comparable results already. The results by the RBE, classical Ewald and PPPM methods in comparison with those predicted by the DH theory are shown in Fig. 1. It shows that the error by the RBE is comparable to those by the Ewald and PPPM methods. As Table 1 indicates, the computational time of the RBE is about 1/5 of that for the Ewald method, 2/5 of that for the PPPM method to achieve comparable results in spite that the system size is not very large.

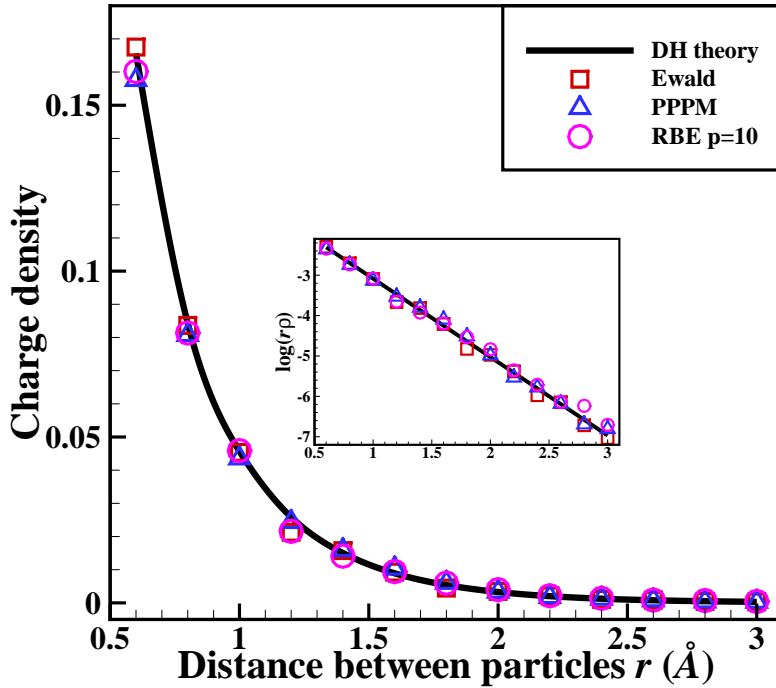


Figure 1: Charge density predicted by simulations using the RBE, classical Ewald and PPPM methods for $L = 10$ and $N = 300$.

We also compute fluctuations of the potential energy (per particle) E_{pot} and pressure P for the three methods to validate the correctness of the configurations. These two quantities are defined by,

$$E_{\text{pot}} = \frac{1}{N}(E_{\text{Coul}} + E_{\text{LJ}}),$$

$$P = \frac{2}{3V} \left(\sum_{i=1}^N \frac{1}{2} m_i v_i^2 - \text{vir} \right), \quad (4.3)$$

where $\text{vir} = (1/2) \sum_{i<j} \mathbf{r}_{ij} \cdot \mathbf{F}_{ij}$ is the virial and E_{Coul} is calculated as (2.9). Pressure is calculated by using Clausius virial theorem with kinetic energy and virial tensor. We recall $\mathbf{r}_{ij} = \mathbf{r}_j - \mathbf{r}_i$, and \mathbf{F}_{ij} is the force of particle j acting on particle i . The potential energy per particle and the average pressure are calculated in LAMMPS using the virial formula [8]. In Fig. 2, the data of every 100 time steps are plotted for the time up to $t = 2000$. We

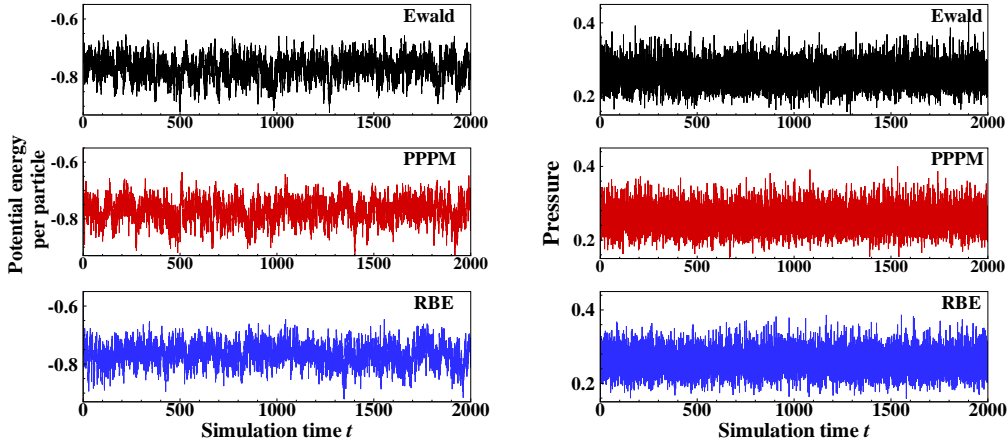


Figure 2: Potential energy per particle and pressure by simulation time using the RBE, classical Ewald and PPPM methods for $L = 10$ and $N = 300$.

calculate the average data of these quantities, $\overline{E}_{\text{pot}}$ and \overline{P} . The relative errors of the RBE compared to the PPPM are both less than 1%.

We increase the size of the system while keeping $\rho_r = N/L^3 = 0.3$ constant to measure the accuracy as well as the computational time. In particular, we choose $N = 300, 600, 1200$ and 2400 , respectively, and the length L is computed correspondingly. In Fig. 3, we show the simulation results for the charge distribution with $p = 10$ for the RBE method. Clearly, the simulation results of the RBE method still agree well with the DH theory for larger r . Particularly, in the embedded subplot we can observe that the linear relation holds up to the error tolerance $e^{-7} \approx 9.1 \times 10^{-4}$, which confirms the accuracy of the RBE method.

	α	r_c	n_c
$N = 100$	0.55	4.0	10
$N = 300$	0.55	4.0	10
$N = 1000$	1.1	3.0	13
$N = 4000$	2.5	2.0	18

Table 2: Parameters for PPPM and RBE, where the RBE method does not have n_c .

Next, we compute the relative accuracy of the potential energy for the RBE method against the PPPM for different densities. In particular, we fix $L = 10$ and consider $N = 100, 300, 1000$ and 4000 , respectively (correspondingly, $\rho_r = 0.1, 0.3, 1.0$ and 4.0). The parameters used in the calculations are shown in Table 2. Note that the RBE method does not have n_c parameter and instead we choose the batch size $p = 10, 20, 50$ and 100 , respectively. The time step is again $\Delta t = 0.002$. The potential energies are computed using 10^4 configurations after equilibrium, sampled every 100 steps. The results are shown in Table 3. Clearly, if we increase the density, we need to use larger batch size to get acceptable accuracy. The RBE with fixed batch size p will have bigger error if the density is increased. Anyhow, even when $\rho_r = 4.0$, using $p = 100$ seems enough to get acceptable results.

ρ_r	$p = 10$	$p = 20$	$p = 50$	$p = 100$
0.1	0.15%	0.13%	0.13%	0.08%
0.3	0.10%	0.08%	0.04%	0.09%
1.0	0.66%	0.18%	0.11%	0.04%
4.0	7.83%	2.38%	0.71%	0.31%

Table 3: Relative error of potential energy for the RBE method against the PPPM method with different densities and batch sizes.




Figure 3: Charge density simulated by the RBE method with increasing system sizes for constant density. The batch size is $p = 10$ for all N .

Lastly, we compare the efficiency for the classical Ewald, PPPM and RBE methods. In Fig. 4 the computational times for the three methods are shown for system size up to $N = 10^6$, where the solid lines present the linear fitting of the data in log-log scale. The results agree with the fact that the complexity per time step for the Ewald summation is of $\mathcal{O}(N^{3/2})$, while the complexity per time step for the RBE is only of $\mathcal{O}(N)$ and the complexity per time step for the PPPM method is a little larger than $\mathcal{O}(N)$. The cost of the RBE is small even when one chooses batch size $p = 100$. The RBE has the best efficiency over a whole range of particle numbers, clearly demonstrating the attractive performance of the algorithm. We remark that a systematic study of the efficiency of the method will be performed in our next work for large-scale all-atom systems, in particular, the comparison with the performance of the PPPM.

4.2 Charge inversion in salty environment

When a highly charged colloid in a solution that contains multivalent counterions, its charge can become overcompensated due to the strong ion correlation between counterions, leading to the charge inversion (or charge reversal) phenomenon. The many-body phenomenon has attracted much attention in the past decades from both experimental [4, 43], theoretical and simulation studies [23, 5, 40, 47, 33, 19], since the charge inversion implies that the effective charge of the colloid-microion complex is abnormally inverted, opposite to the common intuition of understanding from the traditional Poisson-Boltzmann theory.

We follow the setup of Lenz and Holm [33] and consider a highly charged colloid in a solution of asymmetric 3:1 salt with additional 1:1 salt. The colloid has a spherical geometry of diameter $d_0 = 100\text{\AA}$ with a point charge $Q_0 = -300e_0$ at its center. Here $e_0 = 1.6 \times 10^{-19}C$ is the elementary charge. It is placed at the middle of a cubic box with the PBC. The side length of the periodic box is set to $L = 225.8\text{\AA}$ (the volume corresponds to a spherical cell




Figure 4: CPU time per step for the classical Ewald and RBE methods with increasing N .

of radius $R_0 = 140\text{\AA}$). Initially, a total of 200 trivalent counterions, $300 + n_{\text{salt}}$ monovalent coions and n_{salt} monovalent counterions are randomly distributed within the box. These ions have uniform size with a diameter of 4\AA . Clearly, the system satisfies the charge neutrality. The trivalent counterions correspond to the concentration of $c_{(+3)} = 30\text{ mM}$ (i.e. milli mole per liter).

In this example, we implement the methods by the self-written molecular dynamics code with C++. We focus on the accuracy comparison and investigate if or not the RBE can get the correct results with less effort for this relatively complicated many-body phenomenon. Due to the strong charge of the colloid, we take the classical Ewald results as the reference solution. In the simulations, we consider two concentrations for the additional 1:1 salt, i.e., $c_{\text{salt}} = 0\text{ mM}$ and $c_{\text{salt}} = 196\text{ mM}$, where the latter case corresponds to the number of particles $n_{\text{salt}} = 1300$ for each ionic species. The temperature is set to the room temperature $T_* = 298\text{ K}$, and the Bjerrum length ℓ_B is determined by $\ell_B = e_0^2 / (4\pi\epsilon_0\epsilon_r k_B T_*)$, where $\epsilon_r = 78.5$ is the relative dielectric constant of water and ϵ_0 is the vacuum permittivity, resulting in $\ell_B = 7.1\text{\AA}$. The van der Waals interaction is again taken to be part of the Lennard-Jones potential (4.1), where $R_c = 2^{1/6}\text{\AA}$, $\sigma = 1\text{\AA}$ and $\epsilon = 1k_B T_*$. Note that r_{off} is different for different pairs, e.g., $r_{\text{off}} = 52\text{\AA}$ between the colloid and an microion, and $r_{\text{off}} = 4\text{\AA}$ between microions.

To do simulations, we scale all lengths by $L_* = 1\text{\AA}$, temperature by $T_* = 298\text{ K}$, and masses by m_* , the mass of ions which are assumed equal. Then, other quantities can be scaled correspondingly: the energy is scaled by $k_B T_*$, the velocity by $(k_B T_*/m)^{1/2}$, and time by $L_*(m_*/k_B T_*)^{1/2}$, etc.. Consequently, in these reduced units, the room temperature becomes $T = 1$, and the Coulomb interaction between two point charges i and j is given by $U_{ij} = \ell_B q_i q_j / r_{ij}$ where ℓ_B is the scaled Bjerrum length with value 7.1. After we computed the forces using formulas in Section 2 or in Section 3, we should multiply the results with $\ell_B = 7.1$ to get the Coulomb forces for this example.

The molecular dynamics simulations are all performed with the Andersen thermostat,

with $1e5$ steps for the burn-in phase and $6e5$ steps for the sampling to compute ensemble averages. In the burn-in phase, we choose time step

$$\Delta t_n = 0.01 / \ln(1 + n),$$

where n is the number of time steps motivated by the simulated annealing mentioned above. In the sampling phase, we choose $\Delta t = 0.002$. The collision frequency $\nu = 10$ for $c_{\text{salt}} = 0$ mM, and $\nu = 1$ for $c_{\text{salt}} = 196$ mM. The reason to use smaller frequency for $c_{\text{salt}} = 196$ mM is to decrease the artificial diffusion effect introduced by the Andersen thermostat.

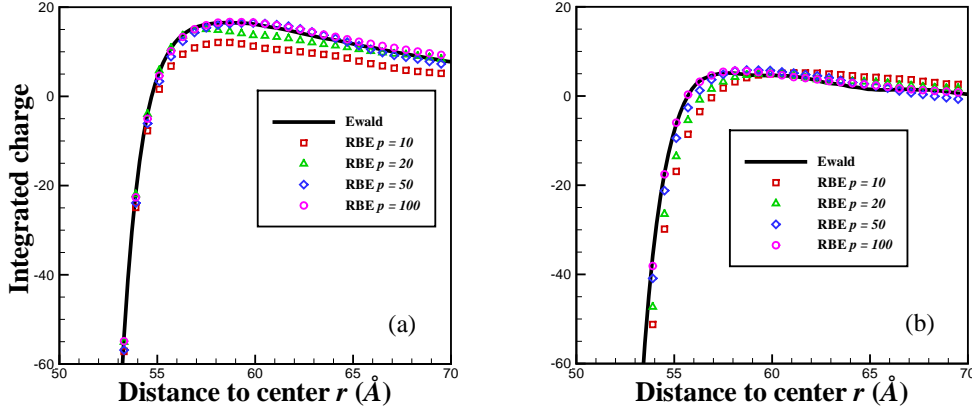


Figure 5: Integrated charge against the distance r from the colloid center when $c_{\text{salt}} = 0$ mM (left) and 196 mM (right): comparison of the Ewald and RBE methods.

	α	r_c	n_c	Time (s)
Ewald ($c = 0$ mM)	0.0014	90.0	8.7	16698
RBE ($c = 0$ mM)	0.0072	40.0	$p = 100$	1167
Ewald ($c = 196$ mM)	0.0014	90.0	8.7	137217
RBE ($c = 196$ mM)	0.0072	40.0	$p = 100$	15258

Table 4: Computational time per $1e5$ simulation steps. The RBE samples from all frequencies and it shows p values in the n_c column.

The settings and running time are shown in Table 4. Clearly, the time consumption of the RBE method is much less (about 1/10 of that for the Ewald method), so the proposed RBE is efficient. The integrated charge distribution, the total charge within the radial direction distance, against the distance r from the colloidal center is plotted in Fig. 5. Regarding the effectiveness, as can be seen from the figure, the RBE can capture the charge reversal phenomenon correctly and obtain acceptable simulation results. As discussed in Proposition 1 and Section 3.3, the force approximation is unbiased, but the randomness results in positive variance leads to numerical heating and systematic error for the equilibrium distribution. As can be seen in Fig. 5, the overcharging effect is weakened for small batches due to this numerical heating. The RBE method converges after $p \gtrsim 100$ and this systematic error is negligible for the system considered here. The inverted charge (maximum of the curve) by the RBE when $p \gtrsim 100$ is in agreement with the Ewald summation and the literature result [33]. This agrees with the discussion above in Sections 3.2–3.3. Since ν is smaller for $c_{\text{salt}} = 196$ mM, the ability of temperature control is reduced and the numerical heating is more obvious for small p values (like $p = 20$). To resolve this, one may consider subtracting the effective temperature due to the random batch from the desired T value, or using better thermostat such as the Nosé-Hoover thermostat. These issues will be explored in our subsequent work.

The charge densities of different kinds of ions are shown in Figs. 6 and 7, for $c = 0$ mM and $c = 196$ mM, respectively. Clearly, the RBE method can compute the densities correctly with acceptable accuracy. Again, larger batch size results in smaller errors. Anyhow, the simulation results seem to be acceptable here for all batch sizes. For these small systems, the Andersen thermostat can already do a satisfactory job, and other temperature preserving techniques can be considered for applications with large systems.




Figure 6: Contribution of different ion types to particle density ρ when $c_{\text{salt}} = 0$ mM.

Overall, according to the two numerical examples above, we find that the RBE method is both accurate and efficient: it can correctly capture the desired physical phenomena while saving the computational time significantly. We also expect the RBE easy to parallelize, and to have good compatibility with a large number of cores. This will be tested in our subsequent work. The extra noise introduced by the random batch can introduce some noticeable numerical heating. One may resolve this by increasing the batch size p or decreasing the time step size to decrease the variance. Some more advanced techniques include subtracting the effective temperature increase or use better thermostats like the Nosé-Hoover. Systematic studies of these will be in our subsequent work, too.

5 Conclusions

In summary, we have developed a novel molecular dynamics method for particle systems with long-range Coulomb interaction using a RBE method which needs only $\mathcal{O}(N)$ operations in each step. The RBE method benefits from a random mini-batch idea for the calculation of the force component in the Fourier series together with an importance sampling for the Fourier modes. We have shown that the algorithm is accurate and efficient by calculating the correlation functions between ions and the charge inversion near the colloidal surface by using two application examples and demonstrated the promising properties for broader applications of the algorithm. Besides the Coulomb systems, the RBE method can be extended to solve other many-body problems such as celestial bodies and complex networks where the long-range interactions also play important roles.

It is remarked that our exploration of the RBE method so far is limited to a few aspects and there are much more work to do in the future. If the system is partially periodic in




Figure 7: Contribution of different ion types to particle density ρ when $c_{\text{salt}} = 196$ mM.

some directions with Directlet or dielectric interface conditions in other directions (e.g., the slab geometries), we believe the extension of our method is straightforward by introducing techniques developed for such problems (see [45, 37, 39] and reference therein). In this work, we have compared the RBE with the classical Ewald summation, and this is actually not fair for the demonstration of the CPU time. More reasonable comparison should be done through the PPPM which is used in many implementations, and this comparison should be performed systematically in addition to the demonstration on the scalability performance in parallel computing. Also, the simulations of this work are based on the primitive model of solvent. This model is simpler by treating water as a continuum medium. This model is very good for the aim of numerical tests of electrostatic algorithms, but the implementation for all-atom simulations shall generate much broader interest for practical uses. All these issues shall be studied in our subsequent works.

Acknowledgement

The work of S. Jin was supported by NSFC grant No. 12031013. The work of L. Li was partially sponsored by NSFC 11901389, Shanghai Sailing Program 19YF1421300 and NSFC 11971314. The work of Z. Xu and Y. Zhao was partially supported by NSFC (grant Nos. 12071288 and 21773165) and the HPC center of Shanghai Jiao Tong University. All the authors are supported by Shanghai Science and Technology Commission (grant No. 20JC1414100).

Appendix

A The Debye–Hückel theory

Consider an electrolyte solution with N ions contained in the cubic box with PBCs, which are idealized as hard spheres of diameter a and carrying charge $\pm q$. The numbers of anions and cations are both $N/2$ to meet electroneutrality condition. Let us fix one ion of charge $+q$ at the origin $\mathbf{r} = 0$ and consider the charge distribution around it.

Inside the region $0 < r < a$ there are no other ions, so the electrostatic potential satisfies the Poisson equation $\varepsilon \nabla^2 \phi = -q\delta(\mathbf{r})$ in this regime, where ε is the permittivity of the solution. Outside this region, the charge of the j th species can be described by the Boltzmann distribution: $\rho_j(\mathbf{r}) = q_j \rho_{\infty,j} e^{-\beta q_j \phi}$ where $j = \pm$ and $q_{\pm} = \pm q$, and $\rho_{\infty,+} = \rho_{\infty,-} = N/(2V)$. Hence, when $r > a$:

$$-\varepsilon \nabla^2 \phi = q\rho_+ e^{-\beta q \phi} - q\rho_- e^{\beta q \phi} \approx \beta q^2 \rho \phi, \quad (\text{A.1})$$

which is the linearized Poisson-Boltzmann equation. By introducing the parameter κ and Debye length λ_D by

$$\kappa \equiv \lambda_D^{-1} = \sqrt{\frac{q^2 \rho}{\varepsilon k_B T}}, \quad (\text{A.2})$$

the solution of the Poisson equation can then be found to be:

$$\phi(r) = \begin{cases} \frac{q}{4\pi\varepsilon r} - \frac{q\kappa}{4\pi\varepsilon(1+\kappa a)}, & r < a, \\ \frac{qe^{\kappa a} e^{-\kappa r}}{4\pi\varepsilon r(1+\kappa a)}, & r > a. \end{cases} \quad (\text{A.3})$$

Hence, the net charge density for $r > a$ is given by

$$\rho(r) = -\varepsilon \nabla^2 \phi(\mathbf{r}) = -\kappa^2 \varepsilon \phi(r). \quad (\text{A.4})$$

Obviously, $\rho(r) < 0$ around the positive charge and

$$\log(r|\rho(r)|) = -\kappa r + \log\left(\frac{\kappa^2 q \exp(\kappa a)}{4\pi(1+\kappa a)}\right)$$

is a linear function of r . The charge density around a negative charge is similarly discussed.

B Variance of the random force under the Debye–Hückel approximation

We consider approximating the charge net density using the Debye–Hückel approximation to estimate,

$$\text{Im}(e^{-i\mathbf{k} \cdot \mathbf{r}_i} \rho(\mathbf{k})) = \text{Im}\left(\sum_{j:j \neq i} q_j \exp(i\mathbf{k} \cdot (\mathbf{r}_j - \mathbf{r}_i))\right). \quad (\text{B.1})$$

We fix the ion q_i at the center. For $r \geq a$, if we use the charge density ρ given by (A.4) and (A.3) to compute this quantity, we get

$$\begin{aligned} \sum_{j:j \neq i} q_j \exp(i\mathbf{k} \cdot (\mathbf{r}_j - \mathbf{r}_i)) &\approx \int_{\mathbb{R}^3 \setminus B(\mathbf{r}_i, a)} \rho(\mathbf{r}) e^{i\mathbf{k} \cdot \mathbf{r}} d\mathbf{r} \\ &= -q_i \frac{\kappa^2 e^{\kappa a}}{2(1+\kappa a)} \int_a^\infty \int_{-1}^1 r^2 \frac{e^{-\kappa r}}{r} \cos(kr z) dz dr \\ &= -q_i \frac{1}{1+\kappa a} \left[\frac{\kappa}{k} \sin(ka) + \cos(ka) \right] / \left(1 + \frac{k^2}{\kappa^2} \right). \end{aligned} \quad (\text{B.2})$$

This term is clearly real and the imaginary part is zero. If one uses the DH theory to compute $\rho(\mathbf{k})$, one may get something bizzard in mathematics. Using the same approximation leads to,

$$e^{-i\mathbf{k} \cdot \mathbf{r}_i} \rho(\mathbf{k}) \approx q_i \left\{ 1 - \frac{1}{1+\kappa a} \left[\frac{\kappa}{k} \sin(ka) + \cos(ka) \right] / \left(1 + \frac{k^2}{\kappa^2} \right) \right\}.$$

This means $\rho(\mathbf{k}) \approx q_i e^{i\mathbf{k} \cdot \mathbf{r}_i} g(k)$ where $g(k)$ is independent of i . The left hand side does not depend on i while the right hand side does. This clearly comes from treating all other

particles except i using the continuum approximation, and i is not special in $\rho(\mathbf{k})$. In spite of the bizzard result for computing $\rho(\mathbf{k})$, we believe that the approximation makes sense when one focuses on computing quantities associated with particle i , and $k \ll a^{-1}$. When $k \ll a^{-1}$, the formula in (B.2) implies that $\text{Im}(e^{-i\mathbf{k}\cdot\mathbf{r}_i}\rho(\mathbf{k})) = 0$. This is understandable: in the equilibrium, provided that all other charges are distributed accurately by the continuum approximation, the net force is zero. In practice, there is thermal fluctuation, and this cannot be zero, but it should be bounded by some number related to the temperature. Moreover, the magnitude of the integral on the right hand side is controlled by a bound uniform in k (recall $|\sin(ax)/x| \leq a$) and we believe this result by continuum approximation can reflect the true magnitude of $\rho(\mathbf{k})$. Hence, when $k \ll a^{-1}$, it is safe to bound $\text{Im}(e^{-i\mathbf{k}\cdot\mathbf{r}_i}\rho(\mathbf{k}))$ by a constant.

Now, if $\sqrt{\alpha} \lesssim a^{-1}$, we can then set $|\text{Im}(e^{-i\mathbf{k}\cdot\mathbf{r}_i}\rho(\mathbf{k}))| \leq C$ in computing (3.10). When $k \gtrsim \sqrt{\alpha}$, we do not assume the bounds on $|\rho(\mathbf{k})|^2$, as such terms will be dominated by $e^{-k^2/(4\alpha)}$. With $\sqrt{\alpha} \sim \rho_r^{1/3}$, one will have $S \approx (\alpha L^2/\pi)^{3/2} \sim N$ by (3.1)–(3.2) so that

$$\begin{aligned} \mathbb{E}|\chi_i|^2 &= \frac{1}{p} \left(\sum_{\mathbf{k} \neq 0} \frac{(4\pi q_i)^2 S}{V^2 k^2} e^{-k^2/(4\alpha)} |\text{Im}(e^{-i\mathbf{k}\cdot\mathbf{r}_i}\rho(\mathbf{k}))|^2 - |\mathbf{F}_{i,1}|^2 \right) \\ &\lesssim \frac{1}{p} \frac{(4\pi q_i)^2 S}{V^2} \int_{2\pi/L}^{\infty} \left(\frac{L}{2\pi} \right)^3 \frac{4\pi k^2}{k^2} e^{-k^2/4\alpha} dk \\ &\approx \frac{8q^2}{p} \frac{S}{V} \sqrt{\alpha\pi} \sim \frac{1}{p} \rho_r^{4/3}. \end{aligned} \tag{B.3}$$

References

- [1] G. ALBI AND L. PARESCHI, *Binary interaction algorithms for the simulation of flocking and swarming dynamics*, Multiscale Modeling & Simulation, 11 (2013), pp. 1–29.
- [2] J. BARNES AND P. HUT, *A hierarchical $O(N \log N)$ force-calculation algorithm*, Nature, 324 (1986), pp. 446–449.
- [3] J. J. BENEDETTO AND G. ZIMMERMANN, *Sampling multipliers and the Poisson summation formula*, Journal of Fourier Analysis and Applications, 3 (1997), pp. 505–523.
- [4] K. BESTEMAN, M. A. G. ZEVENBERGEN, H. A. HEERING, AND S. G. LEMAY, *Direct observation of charge inversion by multivalent ions as a universal electrostatic phenomenon*, Phys. Rev. Lett., 93 (2004), p. 170802.
- [5] H. BOROUDJERDI, Y.-W. KIM, A. NAJI, R. R. NETZ, X. SCHLAGBERGER, AND A. SERR, *Statics and dynamics of strongly charged soft matter*, Phys. Rep., 416 (2005), pp. 129–199.
- [6] L. BOTTOU, *Online learning and stochastic approximations*, On-line Learning in Neural Networks, 17 (1998), p. 142.
- [7] B. R. BROOKS, R. E. BRUCCOLERI, B. D. OLAFSON, D. J. STATES, S. SWANINATHAN, AND M. KARPLUS, *CHARMM: A program for macromolecular energy minimization and dynamics calculations*, J. Comput. Chem., 4 (1983), pp. 187–217.
- [8] D. BROWN AND S. NEYERTZ, *A general pressure tensor calculation for molecular dynamics simulations*, Molecular Physics, (1995).
- [9] G. CICCOTTI, D. FRENKEL, AND I. R. McDONALD, *Simulation of Liquids and Solids: Molecular Dynamics and Monte Carlo Methods in Statistical Mechanics*, North-Holland, Amsterdam, 1987.
- [10] T. DARDEN, D. YORK, AND L. PETERSEN, *Particle mesh Ewald: An $n \log(n)$ method for Ewald sums in large systems.*, J. Chem. Phys., 98 (1993), pp. 10089–10092.

- [11] S. W. DE LEEUW, J. W. PERRAM, AND E. R. SMITH, *Simulation of electrostatic systems in periodic boundary conditions. I. Lattice sums and dielectric constants*, Proceedings of the Royal Society of London. A. Mathematical and Physical Sciences, 373 (1980), pp. 27–56.
- [12] A. DEITMAR AND S. ECHTERHOFF, *Principles of Harmonic Analysis*, Universitext, Springer, Cham, second ed., 2014.
- [13] M. DESERNO AND C. HOLM, *How to mesh up Ewald sums. II. An accurate error estimate for the particle-particle particle-mesh algorithm*, The Journal of Chemical Physics, 109 (1998), pp. 7694–7701.
- [14] Z. H. DUAN AND R. KRASNY, *An Ewald summation based multipole method*, J. Chem. Phys., 113 (2000), pp. 3492–3495.
- [15] U. ESSMANN, L. PERERA, M. L. BERKOWITZ, T. DARDEN, H. LEE, AND L. PEDERSEN, *A smooth particle mesh Ewald method*, J. Chem. Phys., 103 (1995), pp. 8577–8593.
- [16] P. P. EWALD, *Die Berechnung optischer und elektrostatischer Gitterpotentiale*, Ann. Phys., 369 (1921), pp. 253–287.
- [17] R. H. FRENCH, V. A. PARSEGAN, R. PODGORNIK, R. F. RAJTER, A. JAGOTA, J. LUO, D. ASTHAGIRI, M. K. CHAUDHURY, Y.-M. CHIANG, S. GRANICK, S. KALININ, M. KARDAR, R. KJELLANDER, D. C. LANGRETH, J. LEWIS, S. LUSTIG, D. WESOŁOWSKI, J. S. WETTLAUER, W.-Y. CHING, M. FINNIS, F. HOULIHAN, O. A. VON LILIENFELD, C. J. VAN OSS, AND T. ZEMB, *Long range interactions in nanoscale science*, Rev. Mod. Phys., 82 (2010), pp. 1887–1944.
- [18] D. FRENKEL AND B. SMIT, *Understanding molecular simulation: from algorithms to applications*, vol. 1, Elsevier, 2001.
- [19] Z. GAN, X. XING, AND Z. XU, *Effects of image charges, interfacial charge discreteness, and surface roughness on the zeta potential of spherical electric double layers*, J. Chem. Phys., 137 (2012), p. 034708.
- [20] M. GIROTTO, A. P. DOS SANTOS, AND Y. LEVIN, *Simulations of ionic liquids confined by metal electrodes using periodic Green functions*, J. Chem. Phys., 147 (2017), p. 074109.
- [21] L. GREENGARD AND V. ROKHLIN, *A fast algorithm for particle simulations*, J. Comput. Phys., 73 (1987), pp. 325–348.
- [22] ———, *A new version of the Fast Multipole Method for the Laplace equation in three dimensions*, Acta Numerica, 6 (1997), pp. 229–269.
- [23] A. Y. GROSBERG, T. T. NGUYEN, AND B. I. SHKLOVSKII, *The physics of charge inversion in chemical and biological systems*, Rev. Mod. Phys., 74 (2002), pp. 329–345.
- [24] W. K. HASTINGS, *Monte Carlo Sampling Methods Using Markov Chains and Their Applications*, Oxford University Press, 1970.
- [25] 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 (2008), pp. 435–447.
- [26] R. A. HOLLEY, S. KUSUOKA, AND D. W. STROOCK., *Asymptotics of the spectral gap with applications to the theory of simulated annealing.*, Journal of Functional Analysis, 83 (1989), pp. 333–347.
- [27] W. G. HOOVER, *Canonical dynamics: Equilibrium phase-space distributions*, Physical review A, 31 (1985), p. 1695.

- [28] C.-R. HWANG AND S.-J. SHEU, *Large-time behavior of perturbed diffusion markov processes with applications to the second eigenvalue problem for Fokker-Planck operators and simulated annealing*, Acta Applicandae Mathematicae, 19 (1990), pp. 253–295.
- [29] S. JIN, L. LI, AND J.-G. LIU, *Random batch methods (RBM) for interacting particle systems*, J. Comput. Phys., 400 (2020), p. 108877.
- [30] ———, *Convergence of Random Batch Method for interacting particles with disparate species and weights*, SIAM J. Numer. Anal., (to appear).
- [31] J. JUNG, W. NISHIMA, M. DANIELS, G. BASCOM, C. KOBAYASHI, A. ADEDOYIN, M. WALL, A. LAPPALA, D. PHILLIPS, W. FISCHER, C.-S. TUNG, T. SCHLICK, Y. SUGITA, AND K. Y. SANBONMATSU, *Scaling molecular dynamics beyond 100,000 processor cores for large-scale biophysical simulations*, Journal of Computational Chemistry, 40 (2019), pp. 1919–1930.
- [32] J. KOLAFKA AND J. W. PERRAM, *Cutoff errors in the Ewald summation formulae for point charge systems*, Molecular Simulation, 9 (1992), pp. 351–368.
- [33] O. LENZ AND C. HOLM, *Simulation of charge reversal in salty environments: Giant overcharging?*, The European Physical Journal E, 26 (2008), pp. 191–195.
- [34] Y. LEVIN, *Electrostatic correlations: from plasma to biology*, Reports on progress in physics, 65 (2002), p. 1157.
- [35] L. LI, Y. LI, J.-G. LIU, Z. LIU, AND J. LU, *A stochastic version of Stein Variational Gradient Descent for efficient sampling*, Commun. Appl. Math. Comput. Sci., 15 (2020).
- [36] L. LI, Z. XU, AND Y. ZHAO, *A random-batch Monte Carlo method for many-body systems with singular kernels*, SIAM J. Sci. Comput., 42 (2020), pp. A1486–A1509.
- [37] J. LIANG, J. YUAN, E. LUIJTEN, AND Z. XU, *Harmonic surface mapping algorithm for molecular dynamics simulations of particle systems with planar dielectric interfaces*, J. Chem. Phys., 152 (2020), p. 134109.
- [38] B. A. LUTY, M. E. DAVIS, I. G. TIRONI, AND W. F. VAN GUNSTEREN, *A comparison of particle-particle, particle-mesh and Ewald methods for calculating electrostatic interactions in periodic molecular systems*, Mol. Simul., 14 (1994), pp. 11–20.
- [39] O. MAXIAN, R. P. PELAEZ, L. GREENGARD, AND A. DONEV, *A fast spectral method for electrostatics in doubly-periodic slit channels*, arXiv:2101.07088.
- [40] R. MESSINA, C. HOLM, AND K. KREMER, *Charge inversion in colloidal systems*, Comput. Phys. Commun., 147 (2002), pp. 282–285.
- [41] S. NOSÉ, *A molecular dynamics method for simulations in the canonical ensemble*, Molecular physics, 52 (1984), pp. 255–268.
- [42] C. PAN, S. YI, AND Z. HU, *The effect of electrostatic boundaries in molecular simulations: symmetry matters*, Phys. Chem. Chem. Phys., 19 (2017), pp. 4861–4876.
- [43] J. PITTLER, W. BU, D. VAKNIN, A. TRAVESSET, D. J. MCGILLIVRAY, AND M. LÖSCHE, *Charge inversion at minute electrolyte concentrations*, Phys. Rev. Lett., 97 (2006), p. 046102.
- [44] H. ROBBINS AND S. MONRO, *A stochastic approximation method*, The Annals of Mathematical Statistics, (1951), pp. 400–407.
- [45] S. TYAGI, A. ARNOLD, AND C. HOLM, *Electrostatic layer correction with image charges: A linear scaling method to treat slab $2d + h$ systems with dielectric interfaces*, 129 (2008), p. 204102.

- [46] P. J. M. VAN LAARHOVEN AND E. H. L. AARTS, *Simulated Annealing: Theory and Applications.*, D. Reidel Publishing Co., Dordrecht, 1987.
- [47] Z. Y. WANG AND Y. Q. MA, *Insight from Monte Carlo simulations on charge inversion of planar electric double layers in mixtures of asymmetric electrolytes*, J. Chem. Phys., 133 (2010), p. 064704.
- [48] M. WELLING AND Y. W. TEH, *Bayesian learning via stochastic gradient Langevin dynamics*, in Proceedings of the 28th International Conference on Machine Learning (ICML-11), 2011, pp. 681–688.
- [49] L. J. YANG, C. H. TAN, M. J. HSIEH, J. M. WANG, Y. DUAN, P. CIEPLAK, J. CALDWELL, P. A. KOLLMAN, AND R. LUO, *New-generation Amber united-atom force field*, J. Phys. Chem. B, 110 (2006), pp. 13166–13176.
- [50] L. YING, G. BIROS, AND D. ZORIN, *A kernel-independent adaptive fast multipole algorithm in two and three dimensions*, J. Comput. Phys., 196 (2004), pp. 591–626.