Superscalability of the random batch Ewald method

Cite as: J. Chem. Phys. **156**, 014114 (2022); https://doi.org/10.1063/5.0073424 Submitted: 30 September 2021 • Accepted: 19 December 2021 • Accepted Manuscript Online: 19 December 2021 • Published Online: 04 January 2022

🔟 Jiuyang Liang, 🔟 Pan Tan, 🔟 Yue Zhao, et al.

ARTICLES YOU MAY BE INTERESTED IN

Perturbation-adapted perturbation theory The Journal of Chemical Physics **156**, 011101 (2022); https://doi.org/10.1063/5.0079853

Random-batch list algorithm for short-range molecular dynamics simulations The Journal of Chemical Physics **155**, 044108 (2021); https://doi.org/10.1063/5.0056515

Accurate scaling functions of the scaled Schrödinger equation The Journal of Chemical Physics **156**, 014113 (2022); https://doi.org/10.1063/5.0077495





J. Chem. Phys. **156**, 014114 (2022); https://doi.org/10.1063/5.0073424 © 2022 Author(s).

Superscalability of the random batch Ewald method

Cite as: J. Chem. Phys. 156, 014114 (2022); doi: 10.1063/5.0073424 Submitted: 30 September 2021 • Accepted: 19 December 2021 • Published Online: 4 January 2022



Jiuyang Liang,¹ 🔟 Pan Tan,² 🔟 Yue Zhao,¹ 🔟 Lei Li,³ Shi Jin,³ Liang Hong,^{2,a)} and Zhenli Xu^{3,a)} 🔟

AFFILIATIONS

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

²School of Physics and Astronomy and Institute of Natural Sciences, Shanghai Jiao Tong University, Shanghai 200240, People's Republic of China

³School of Mathematical Sciences, Institute of Natural Sciences and MoE-LSC, Shanghai Jiao Tong University, Shanghai 200240, China

^{a)}Authors to whom correspondence should be addressed: hongl3liang@sjtu.edu.cn and xuzl@sjtu.edu.cn

ABSTRACT

Coulomb interaction, following an inverse-square force-law, quantifies the amount of force between two stationary and electrically charged particles. The long-range nature of Coulomb interactions poses a major challenge to molecular dynamics simulations, which are major tools for problems at the nano-/micro-scale. Various algorithms are developed to calculate the pairwise Coulomb interactions to a linear scale, but poor scalability limits the size of simulated systems. Here, we use an efficient molecular dynamics algorithm with the random batch Ewald method on all-atom systems where the complete Fourier components in the Coulomb interaction are replaced by randomly selected minibatches. By simulating the *N*-body systems up to 10^8 particles using 10 000 central processing unit cores, we show that this algorithm furnishes O(N) complexity, almost perfect scalability, and an order of magnitude faster computational speed when compared to the existing state-of-the-art algorithms. Further examinations of our algorithm on distinct systems, including pure water, a micro-phase separated electrolyte, and a protein solution, demonstrate that the spatiotemporal information on all time and length scales investigated and thermodynamic quantities derived from our algorithm are in perfect agreement with those obtained from the existing algorithms. Therefore, our algorithm provides a promising solution on scalability of computing the Coulomb interaction. It is particularly useful and cost-effective to simulate ultra-large systems, which is either impossible or very costly to conduct using existing algorithms, and thus will be beneficial to a broad range of problems at nano-/micro-scales.

Published under an exclusive license by AIP Publishing. https://doi.org/10.1063/5.0073424

I. INTRODUCTION

Molecular dynamics (MD) is one of the most powerful simulation tools in modern science to furnish the atomic-detailed microscopic mechanism underlying experimental findings in a plethora of areas, including physics, chemistry, engineering, biology, and pharmaceutical sciences.^{1–5} Despite the enormous success, the application of MD simulations without specific coarse graining and the enhanced sampling method has been largely limited to moderate size (often below 1×10^6 atoms) and time scale (shorter than 10 ms). These limitations cannot be solved by parallel computing using a large number of computer cores, as the inter-atomic Coulomb interactions^{4,6,7} are long-ranged and require intensive communications between cores, significantly reducing the parallel efficiency, especially when using supercomputers. In the past decades, enormous efforts have been devoted to reduce the computational cost of the Coulomb interaction, and many fast algorithms have been developed, including the lattice summation methods on the basis of fast Fourier transform (FFT),^{8,9} multipole-type methods such as the treecode algorithm¹⁰ and the fast multipole method (FMM),^{11,12} partial differential equation-based multigrid methods,¹³ and Maxwell equation molecular dynamics.¹⁴ These algorithms can reduce the computational complexity to $O(N \log N)$ or even O(N) and have achieved great success in practical applications. However, none of these methods can achieve high scalability for calculating the Coulomb interaction while using a large number of computer cores.¹⁵

The recent random batch Ewald (RBE)¹⁶ method is an alternative O(N) algorithm for electrostatic calculations. It is based on the Ewald splitting, but it avoids the use of the FFT by introducing random mini-batch sampling on the Fourier space to approximate the force contribution from the long-range part. In this work, we develop the RBE for the all-atom molecular dynamics simulation and demonstrate its great computational efficiency and high scalability, especially when applied on large-scale simulations using supercomputers owing to great reduction in global communications. The so-called "random mini-batch" in the RBE originated from the stochastic gradient descent method widely used in machine learning^{17,18} and was first proposed for interacting particle systems with rigorous error estimates.¹⁹ It has succeeded in Monte Carlo simulation on particle systems.²⁰ In the present work, we demonstrate that the RBE-based MD enhances the computational speed by an order of magnitude in comparison to the state-ofthe-art algorithms, including particle-particle particle-mesh (PPPM) and particle-mesh Ewald (PME) methods, and maintains the parallel efficiency of nearly 95% when paralleling up to 10 000 central processing unit (CPU) cores to simulate a large system of 10⁸ atoms. Moreover, a systematic test was conducted on all-atom simulations of three representative systems, namely, bulk water, a micro-phase separated aqueous electrolyte, and a protein solution, revealing that the spatiotemporal information for these systems on a broad range of time and length scales and the thermodynamical quantities are quantitatively reproduced by the RBE-based MD. Thus, as compared to the mainstream algorithms (PPPM or PME^{8,9}), the RBE-based MD furnishes a novel algorithm with substantial improvement in computational efficiency and parallel scalability while maintaining the accuracy of the spatiotemporal information.

This paper is organized as follows: In Sec. II, we introduce the random batch Ewald method and then describe the parallel implementation in details. In Sec. III, we validate the superior CPU performance and superscalability of the RBE method by the simulation on bulk water and perform numerical calculations for three benchmark problems, including bulk water, a micro-phase separated aqueous electrolyte, and a protein solution, to demonstrate the accuracy of the RBE method. Discussions and conclusions are given in Secs. IV and V, respectively.

II. METHODS

A. The random batch Ewald method

In the classical Ewald method,²¹ the Coulomb kernel is split into two components,

$$\frac{1}{r} = \frac{\operatorname{erf}(\sqrt{\alpha}r)}{r} + \frac{\operatorname{erfc}(\sqrt{\alpha}r)}{r}, \qquad (1)$$

where $\operatorname{erf}(\cdot)$ is the error function and $\operatorname{erfc}(\cdot)$ is its complementary function such that the first term is a smooth function and the second one becomes short-ranged. To mimic the bulk environment, a periodic boundary condition is assumed. Without loss of generality, one considers a system of *N* charged particles located at r_j for $j = 1, \ldots, N$ in a cubic box of length *L* and volume $V = L^3$. By applying the Fourier transform on the smooth part of the electrostatic interactions of these charges, one can write the total Coulomb energy as $U = U_1 + U_2$, with

$$U_1 = \frac{2\pi}{V} \sum_{\boldsymbol{k} \neq \boldsymbol{0}} \frac{|\rho(\boldsymbol{k})|^2}{|\boldsymbol{k}|^2} e^{-\frac{|\boldsymbol{k}|^2}{4\alpha}} - \sqrt{\frac{\alpha}{\pi}} \sum_i q_i^2$$
(2)

and

$$U_2 = \frac{1}{2} \sum_{n}' \sum_{ij} q_i q_j \frac{\operatorname{erfc}(\sqrt{\alpha}|\boldsymbol{r}_{ij} + \boldsymbol{n}L|)}{|\boldsymbol{r}_{ij} + \boldsymbol{n}L|},$$
(3)

where $\mathbf{r}_{ij} = \mathbf{r}_j - \mathbf{r}_i$ is the vector starting at particle *i* and pointing toward particle *j*, $\mathbf{k} = 2\pi \mathbf{m}/L$ with $\mathbf{m} \in \mathbb{Z}^3$, and $\rho(\mathbf{k})$ is the structure factor defined by

ρ

$$(\boldsymbol{k}) \coloneqq \sum_{i=1}^{N} q_{i} e^{i\boldsymbol{k}\cdot\boldsymbol{r}_{i}}.$$
(4)

Here, the structure factor is the conjugate of the Fourier transform of the charge density. The Coulomb force acting on the *i*th particle in the form of the Ewald summation is the derivatives of U in two parts,

$$\begin{aligned} \mathbf{F}_{i} &= -\sum_{\mathbf{k}\neq 0} \frac{4\pi q_{i}\mathbf{k}}{V|\mathbf{k}|^{2}} e^{-\frac{|\mathbf{k}|^{2}}{4\alpha}} \operatorname{Im}\left(e^{-i\mathbf{k}\cdot\mathbf{r}_{i}}\rho(\mathbf{k})\right) - q_{i}\sum_{j,n}' q_{j}G(|\mathbf{r}_{ij}+\mathbf{n}L|) \frac{\mathbf{r}_{ij}+\mathbf{n}L}{|\mathbf{r}_{ij}+\mathbf{n}L|} \\ &=: F_{i,1} + F_{i,2}, \end{aligned}$$
(5)

where G(r) is defined by

$$G(r) := \frac{\operatorname{erfc}(\sqrt{\alpha}r)}{r^2} + \frac{2\sqrt{\alpha}e^{-\alpha r^2}}{r^2}$$
(6)

and Im(·) denotes the imaginary part of the independent variable. In Eq. (5), $F_{i,1}$ represents the component in the Fourier space, while $F_{i,2}$ denotes the one in the real space. By proper choice of the parameters (α and the real-space and Fourier-space cutoffs r_c and k_c), the computational complexity is optimized to $O(N^{3/2})$. Moreover, the FFT is often employed to further speed up the evaluation of $F_{i,1}$ such that the real-space cutoff radius can be much smaller than that of the classical Ewald summation, resulting in the core algorithms.^{8,9} A final computational complexity of $O(N \log N)$ can be achieved through these methods for periodic systems.

The RBE avoids the use of the FFT; instead, it employs the random mini-batch strategy to calculate $F_{i,1}$. Here, one picks a small batch of frequencies when evaluating $F_{i,1}$ for a given particle *i*. These frequencies are chosen randomly with an importance-sampling scheme (see below for a brief overview and the work of Jin *et al.*¹⁶ for more details). Let *P* be the batch size and

$$S := \sum_{k \neq 0} \exp(-|\mathbf{k}|^2 / 4\alpha) = H^3 - 1,$$
(7)

with

$$H := \sum_{m \in \mathbb{Z}} e^{-\pi^2 m^2 / (\alpha L^2)} = \sqrt{\frac{\alpha L^2}{\pi}} \sum_{m = -\infty}^{\infty} e^{-\alpha m^2 L^2},$$
 (8)

where the second equality holds due to the Poisson summation formula.²² Then, we have the probability

$$\mathcal{P}_{k} = S^{-1} e^{-\frac{|k|^{2}}{4\alpha}} \tag{9}$$

for $k \neq 0$, which is a discrete Gaussian distribution. Samples obeying this probability can be obtained efficiently by applying the

J. Chem. Phys. **156**, 014114 (2022); doi: 10.1063/5.0073424 Published under an exclusive license by AIP Publishing

Metropolis-Hastings (MH) algorithm²³ to sample from the discrete distribution

$$\mathcal{P}_m := H^{-1} e^{-(2\pi m/L)^2/(4\alpha)},$$
(10)

performing this sampling procedure for three components and discarding the k = 0 sample. At each step of our simulation, one picks a batch size *P*, which is of O(1), and draws *P* frequencies k_{ℓ} , $1 \le \ell \le P$ independent identically distributed (i.i.d) from the discrete distribution \mathcal{P}_k by the above-mentioned MH sampling method. The approximate force then reads

$$\boldsymbol{F}_{i,1}^{*} = -\sum_{\ell=1}^{P} \frac{S}{P} \frac{4\pi \boldsymbol{k}_{\ell} \boldsymbol{q}_{i}}{V |\boldsymbol{k}_{\ell}|^{2}} \operatorname{Im} \left(e^{-i\boldsymbol{k}_{\ell} \cdot \boldsymbol{r}_{i}} \rho(\boldsymbol{k}_{\ell}) \right), \tag{11}$$

which is an unbiased estimator of $F_{i,1}$ in Eq. (5).

Since the number of frequencies *P* used to estimate $\rho(\mathbf{k})$ for each particle is system-size independent, the overall computational cost to calculate the structure factor for the entire simulation system scales as O(N) under a given level of accuracy.¹⁶ The sketch map of the RBE algorithm is present in Fig. 1(a).

In the isothermal-isobaric ensemble simulations (NPT), the virial tensor calculation is needed at each MD step. It can also be done by employing importance sampling in the Fourier space. With the same batch of *P* Fourier modes, the approximate reciprocal space contribution of the virial reads

$$\Xi_{1}^{\beta\gamma} = -\sum_{\ell=1}^{p} \frac{S}{P} \frac{\pi}{V |\boldsymbol{k}_{\ell}|^{2}} |\rho(\boldsymbol{k}_{\ell})|^{2} \cdot \left[\delta_{\beta\gamma} - 2k_{\ell}^{\beta} k_{\ell}^{\gamma} \left(\frac{1}{4\alpha^{2}} + \frac{1}{|\boldsymbol{k}_{\ell}|^{2}} \right) \right], \quad (12)$$

where β and γ are dimensions taken from the three coordinates indicating the corresponding components of the tensor. The details for deriving Eq. (12) are briefly given in Sec. II C.

The steps of the RBE-based MD algorithm are summarized as follows:

- (i) Set parameters α , r_c , and batch size *P*. Load initial positions and strengths of charges.
- (ii) Sample sufficient number of $\mathbf{k} \sim e^{-|\mathbf{k}|^2/4\alpha}$ with $\mathbf{k} \neq \mathbf{0}$, by employing importance sampling to form the total set of frequency samples, \mathcal{K} .
- (iii) Evolve Newton's equations. The real part $F_{i,2}$ of the Coulomb force is directly computed with cutoff r_c , whereas the Fourier part is approximated by $F_{i,1}^*$ using Eq. (11) with the *P* frequencies chosen from \mathcal{K} in order.
- (iv) If the NPT ensemble is employed, compute the real-space virial and the approximated Fourier virial using Eq. (12).

B. Parallel implementation

The RBE is especially suitable for parallelization and vectorization. Here, we present the implementation strategy with hybrid message passing interface (MPI)/OpenMP parallelization for the RBE in both the all-atom NVT and NPT simulations, which supports massively parallel MD simulations of large-scale systems. We use the Intel 512-bit SIMD (AVX-512 architecture) for vectorization implementation, which operates 16 neighbors for singleprecision floating calculation (or eight for double precision) at the same time, and the Intel Parallel Studio for parallelization (including functions of MPI, OpenMP, and AVX-512 instructions). The communication and vectorization procedures are optimized as follows.

Steps (ii) and (iii) of the RBE require a serial importance sampling procedure and a global broadcast operation, whereas their cost is relatively small and can be eliminated by the designed nonjammed communication and computation/communication overlapping. First, for the NVT ensemble, assume that *M* MPI ranks are employed, and *M* independent sampling processes are executed



FIG. 1. Sketch map. (a) The algorithm sketch map of the RBE applied to the electrostatic long-range interaction. (b) Parallel strategy in the Fourier space employing single-instruction multiple-data (SIMD). Only one global operation with O(1) amount of data is required during each step.

in parallel within each rank. Next, the first MPI rank broadcasts the samples to other ranks using the blocking operation. Finally, the equations of motion are calculated [i.e., step (iii)], whereas the samples in other ranks are concurrently broadcasted. This strategy evaluates and updates the samples every M steps, dramatically reducing the global communication cost. Second, for the NPT ensemble, the above strategy no longer works due to the dynamically changed size of the box. We offer an alternative method by the lights of the Multiple-Program Multiple-Data parallelization in GROMACS.²⁴ When M is large, one MPI rank is selected to do only the sampling, which is sampled from the standard normal distribution, and then broadcast the samples and the random variables which are required in the Metropolis step to other ranks. Other ranks receive the samples and multiply by a constant with respect to the instantaneous size of the box on them. An acceptancerejection step is then run at each rank with the same random variables. The communication operation can also overlap with the computation.

Step (iii) also requires the evaluation of the real-part force $F_{i,2}$ of the Coulomb force, and we follow the classical procedure in the MD package of Large-scale Atomic/Molecular Massively Parallel Simulator (LAMMPS).^{25,26} The calculation of the Fourier part $(F_{i,1}^*)$ is rather time-consuming in previous methods, and it is now evaluated using Eq. (11), whose parallel strategy is displayed in Fig. 1(b). The samples and the positions of particles are packaged into 512-bit vectors when the structure factors $\rho(\mathbf{k})$ are evaluated. The OpenMP threads can be selected to involve/uninvolve in loop parallelization when one evaluates $\rho(\mathbf{k})$. Only one global operation, MPI_Allreduce, is required for reducing $\rho(\mathbf{k})$. The approximated force $F_{i,1}^*$ of each particle and the Fourier virial are then obtained from the structure factors.

C. Calculation of the virial

We briefly discuss the calculation of the virial [Eq. (12)]. In a system with periodic boundary conditions, the macroscopic pressure p of a set of N particles contained in a volume V obeys the following relation:

$$p = \frac{2}{3 V} (E_{\rm kin} - \Xi), \tag{13}$$

where $E_{\rm kin}$ is the instantaneous kinetic energy and Ξ is the virial. In the full tensor form, the virial can be written as

$$\Xi = -\frac{1}{2} \sum_{i$$

where r_{ij}^n denotes the distance vector of the nearest image of atom *i* from atom *j* and \otimes denotes the direct product of two vectors. The straightforward implementation of Eq. (14) involves its evaluation in the inner loop of the non-bonded force routine, which results in a significant CPU time consumption. Nevertheless, it is possible to extract the virial calculation from the inner loop. A hybrid method is often employed to evaluate the Ewald-based electrostatic virial such that the real-space virial is treated as the tensor form employing the method in the literature²⁷ and the Fourier virial is computed from the derivative of Lagrangian U_F ,

$$\Xi_{1}^{\beta\gamma} = -\sum_{\eta} \frac{-\partial U_{\rm F}}{\partial h_{\gamma\eta}} h_{\beta\eta}$$
$$= -\sum_{\boldsymbol{k}\neq\boldsymbol{0}} \frac{\pi}{V|\boldsymbol{k}|^{2}} e^{-\frac{|\boldsymbol{k}|^{2}}{4\alpha}} |\rho(\boldsymbol{k})|^{2} \left[\delta_{\beta\gamma} - 2k^{\beta}k^{\gamma} \left(\frac{1}{4\alpha} + \frac{1}{|\boldsymbol{k}|^{2}} \right) \right], \quad (15)$$

where *h* is the tensor indicating the size of the box, β , γ , and η are taken from {*x*, *y*, *z*} indicating the corresponding component. Although the Fourier virial can be cheaply evaluated using Eq. (15), it cannot be directly derived from RBE due to the incomplete data of the structure factors, as only *P* of them is evaluated in Eq. (11). To address this problem, we follow the same idea that employs importance sampling from the Gaussian distribution in the Fourier space. With the same batch of *P* frequencies [see Eq. (11)], the approximate reciprocal space contribution of the virial Ξ_1 reads as Eq. (12).

D. Consistency and stability analysis of employing RBE in the NPT ensemble

We define the fluctuations in the random batch approximation for the Fourier part of the force and the pressure on particle *i* by

$$\boldsymbol{\chi}_i = \boldsymbol{F}_{i,1}(\boldsymbol{r}) - \boldsymbol{F}_{i,1}^*(\boldsymbol{r}) \quad \text{and} \quad \widetilde{\boldsymbol{\chi}}^{\beta\gamma} = p^{\beta\gamma} - \widetilde{p}^{\beta\gamma}, \quad (16)$$

respectively. The expectations and variances of the fluctuation can be obtained by direct calculation, which are given by

$$\mathbb{E}(\boldsymbol{\chi}_{i}) = 0,$$

$$\mathbb{E}|\boldsymbol{\chi}_{i}|^{2} = \frac{1}{P} \left(\sum_{\boldsymbol{k} \neq \boldsymbol{0}} \frac{(4\pi q_{i})^{2} S}{V^{2} |\boldsymbol{k}|^{2}} e^{-\frac{|\boldsymbol{k}|^{2}}{4\alpha}} \left| \operatorname{Im} \left(e^{-i\boldsymbol{k}\cdot\boldsymbol{r}_{i}} \rho(\boldsymbol{k}) \right) \right|^{2} \right) - \frac{1}{P} |\boldsymbol{F}_{i,1}|^{2}, \quad (17)$$

and

$$\mathbb{E}\left(\tilde{\chi}^{\beta\gamma}\right) = 0,$$

$$\mathbb{E}\left(\left|\tilde{\chi}^{\beta\gamma}\right|^{2}\right) = \frac{4S}{9P} \sum_{\boldsymbol{k}\neq\boldsymbol{0}} \frac{\pi^{2} |\rho(\boldsymbol{k})|^{4}}{V^{4} |\boldsymbol{k}|^{4}} e^{-\frac{|\boldsymbol{k}|^{2}}{4\alpha}} \mathcal{Q}_{\beta\gamma}(|\boldsymbol{k}|) - \left|p^{\beta\gamma}\right|^{2},$$
(18)

with

$$\mathcal{Q}_{\beta\gamma}(|\boldsymbol{k}|) = \delta_{\beta\gamma} - 2k^{\beta}k^{\gamma} \left(\frac{1}{4\alpha^{2}} + \frac{1}{|\boldsymbol{k}|^{2}}\right)^{2}.$$
 (19)

Here, \mathbb{E} represents the expectation (or the ensemble average). Equations (17) and (18) imply that the random approximation is consistent or unbiased, namely,

$$\mathbb{E}\boldsymbol{F}_{i,1}^{*} = \boldsymbol{F}_{i,1}, \text{ and } \mathbb{E}\widetilde{\boldsymbol{p}}^{\beta\gamma} = \boldsymbol{p}^{\beta\gamma}.$$
 (20)

Equation (20) illustrates that the averaged effect of the RBE quantities is correct, and the random batch type methods work due to this time averaging effect, which can be regarded as the law of large numbers in time. This is a Monte-Carlo technique, but we have demonstrated that the RBE can correctly reproduce the dynamical properties (see Fig. 3) as the short-range part is calculated exactly. Moreover, the following strong error holds:^{19,28}

$$\mathbb{E}\left[\frac{1}{N}\sum_{i}\left(\left|\boldsymbol{r}_{i}-\tilde{\boldsymbol{r}}_{i}\right|^{2}+\left|\boldsymbol{\upsilon}_{i}-\widetilde{\boldsymbol{\upsilon}}_{i}\right|^{2}\right)\right]^{\frac{1}{2}}\lesssim\sqrt{\Lambda(N)\Delta t},$$
(21)

J. Chem. Phys. **156**, 014114 (2022); doi: 10.1063/5.0073424 Published under an exclusive license by AIP Publishing

where Δt is the time step and

$$\Lambda(N) = \max\left\{\mathbb{E}(|\boldsymbol{\chi}_i|^2), \mathbb{E}(|\boldsymbol{\tilde{\chi}}^{\beta\gamma}|^2)\right\}$$
(22)

is an upper bound of the variance. Note that in the mean field regime,¹⁶ $\Lambda(N)$ is independent of *N*. It was also shown that if $\sqrt{\alpha} \sim \rho_r^{1/3}$ (as it is used in the PPPM²⁹), we have $\mathbb{E} |\chi_i|^2 \leq \rho_r^{4/3}/P$, with $\rho_r = N/V$ being the particle density. Moreover, at the Debye–Hückel limit, the expectation of the squared structure factor is approximated by^{30,31}

$$\langle |\rho(\boldsymbol{k})|^2 \rangle \simeq \frac{k_{\rm B} T V}{\lambda_{\rm D}^2 + 1/|\boldsymbol{k}|^2},$$
 (23)

where $k_{\rm B}T$ is the thermal energy and $\lambda_{\rm D}$ is the Debye length. Assume a symmetric electrolyte with $|q_i| = q$ for all *i* without loss of generality. One has $\lambda_{\rm D} = \sqrt{k_{\rm B}T\varepsilon_0\rho_r/q^2}$, with ε_0 being the vacuum dielectric constant. Since $\mathbf{k} = 2\pi \mathbf{m}/L$ and $\lambda_{\rm D}$ is constant for the fixed particle density $\rho_r = N/V$, one has $\langle |\rho(\mathbf{k})|^2 \rangle \simeq O(V^{1/3}) = O(N^{1/3})$. Following our previous work,¹⁶ one can find that the variances of the approximate force $F_{i,1}^*$ and pressure $\tilde{\rho}^{\beta\gamma}$ decay as $O(P^{-1})$. Note that in the equations of motion of the NPT ensemble,^{32,33} the time derivatives of the momentum and the volume depend linearly on the force and the pressure, respectively. These arguments indicate that the error of the RBE is expected not to grow with the increase of *N* for a fixed density. However, rigorous justification^{16,33} remains an open problem.

III. RESULTS

In this section, we discuss the main results of this paper to demonstrate that the RBE-based MD enhances the computational speed by an order of magnitude in comparison to the state-of-theart algorithms, including PPPM and PME, and maintains promising parallel efficiency. Moreover, a systematic test was conducted on all-atom simulations of three systems, namely, bulk water, a microphase separated aqueous electrolyte, and a protein solution, to demonstrate the accuracy on the spatiotemporal information and the thermodynamical quantities of RBE-based MD. The experimental design for all these simulations and the information of the



Number of CPU Cores C

FIG. 2. Performance comparison. The average CPU time spent per simulation step for four bulk water systems of different sizes of (a) 53 367 atoms, (b) 311 469 atoms, (c) 3 000 000 atoms, and (d) 100 158 744 atoms, as a function of the number of CPU cores, for the RBE and the PPPM. Blue line: RBE with P = 100. Red line: PPPM with $\Delta = 10^{-4}$. The connecting lines between marks are smoothed by employing the B-spline interpolation. The light-colored areas bounded by the dotted lines with appropriate colors show the confidence intervals. The insets in panels show the relative parallel efficiency $\eta(C)$. Note that the insets use log scales in (a) and (b) but linear scales in (c) and (d) for better illustration.

employed hardware and software are given in Appendixes A and B, respectively.

A. CPU performance

The comparison between the RBE and PPPM was made by using the MD engine of LAMMPS on the bulk water assuming the force field of the extended simple point charge (SPC/E) model.³⁴ The parameters of the PPPM are chosen automatically in LAMMPS for a given error level Δ of the relative force.³⁵ The parameter α in the RBE is chosen to be the same as that in the PPPM. The simulations were conducted for two thousand steps to calculate the average CPU time per step, denoted as T(C), where *C* is the number of cores. Four systems of different sizes are used in the test with N = 53 367, 311 469, 3 000 000, and 100 158 744 atoms, and the corresponding results are present in Figs. 2(a)–2(d), respectively. As it can be seen, the computational speed measured by T(C) from the RBE can be an order of magnitude faster than that from the PPPM.

Moreover, the relative parallel efficiency $\eta(C)$ at a given number of cores *C* defined by Eq. (24) is used to characterize the scalability of the algorithm,¹⁵

$$\eta(C) = \frac{C_{\min}}{C} \cdot \frac{T_{\text{best}}}{T(C)},$$
(24)

where $C_{\rm min}$ denotes the minimal nuber of cores used in the calculation and $T_{\rm best}$ is the run time of the fastest method at $C_{\rm min}$. For example, $C_{\rm min} = 1$ for Figs. 2(a) and 2(b) but set to 10 and 120 for Figs. 2(c) and 2(d), respectively, since one processor is too time-consuming and storage-limiting to simulate such large systems as in Figs. 2(c) and 2(d). The relative parallel efficiency illustrates that the RBE remains 95% for up to 10 000 cores when simulating 10⁸ atoms [inset of Fig. 2(d)], significantly outperforming that of the PPPM, which drops to ~ 20% for the same system.

B. Accuracy of the RBE and comparison with the PPPM/PME methods

1. Pure water systems

We calculate four physical quantities on the bulk water simulation including the radial distribution function (RDF), the mean square displacement (MSD), the velocity autocorrelation function (VACF), and the hydrogen bond autocorrelation function (HBACF) to examine the accuracy of the RBE as compared to the PPPM. The



FIG. 3. Comparison of accuracy between the PPPM and RBE on bulk water at the NVT ensemble. Simulation results from the PPPM (red dashed-dotted line) and from the RBE with P = 100 (blue solid line) in bulk water. (a) The radial distribution function of oxygen–oxygen in water molecules. (b) The mean square displacement of the center of mass of water molecules. (c) The velocity autocorrelation function of oxygen in water molecules. (d) The hydrogen bond autocorrelation function.



FIG. 4. The variation of the total potential energy (a) and the temperature variation of bulk water (b) produced by the RBE and PPPM in the NVT ensemble, respectively. As it can be seen, the fluctuations of potential energy and temperature derived from the RBE and PPPM are guantitatively similar.

definition of the HBACF is presented in Appendix C. The RDF of oxygen-oxygen atom pairs furnishes the spatial arrangement of water molecules. The MSD describes the translational motion on the time scale from 1 fs to 1 ns. The VACF and HBACF characterize the short-time vibrational, liberational, and rotational dynamics of water. As it can be seen in Fig. 3, for both spatial arrangement (Å to nm) and dynamical motions of the water molecules (fs to ns), the results derived from the RBE are almost identical to those from the PPPM. The comparisons on the fluctuation of the total potential energy and the temperature of the system over the simulation time are presented in Fig. 4. All these tests are calculated at the NVT ensemble, and we also calculated them at the NPT ensemble and similar results as Figs. 3 and 4 are obtained. In addition, the isochoric and isobaric heat capacities and the relative dielectric constant of bulk water (defined in Appendix D) in the ensembles of NVT and NPT were calculated from the RBE simulations as provided in Fig. 5, showing quantitative agreement with those derived from the PME.

2. LiTFSI ionic liquid

The second benchmark test is an aqueous electrolyte (TFSI) at ultrahigh concentration (5 M/l). We implement the algorithm in the GROMACS package, as the electrolyte requires a special force field which is installed in this MD engine. For comparison, the reference simulations using the PME were also conducted using the GROMACS package. At such high concentration, the electrolyte is microscopically inhomogeneous, separating into two phases (water vs anions) at the length scale 1-2 nm, which are mutually percolated in space.^{36,37} Figure 6(a) illustrates an MD snapshot for the electrolyte, revealing the nano-heterogeneity of the system. The structural information, i.e., the RDFs of the center atom (nitrogen) of the anions, on this concentrated electrolyte derived from both RBE and PME is shown in Fig. 6(b), while the dynamics in the system, including conductivity, viscosity, and diffusion constants, are presented in Figs. 6(c)-6(f). As it can be seen, the spatiotemporal features





J. Chem. Phys. **156**, 014114 (2022); doi: 10.1063/5.0073424 Published under an exclusive license by AIP Publishing



FIG. 6. Comparison of simulation results derived from PME and RBE on concentrated ionic solution. The RBE uses P = 500. (a) Simulation snapshot of the system. The anions are shown in red, and water molecules are represented in blue. The inset shows the structure of the solute. (b) The radial distribution function of nitrogen–nitrogen between anions. (c) The conductivity of the electrolyte. (d) The viscosity of the system. (e) The diffusion constant of Li⁺. (f) The diffusion constant of the center of mass of water molecules.



FIG. 7. Comparison of simulation results derived from PME and RBE on the protein solution. The RBE (blue) uses P = 500. (a) Simulation snapshot of the system. (b) The root mean square deviation (RMSD) of the backbone atoms in the protein. (c) Solvent accessible surface area (SASA). (d) The root mean square fluctuation (RMSF) of residues in the protein. (e) Distribution of the characteristic inter-domain distance between residue C54 and C97.

of the system derived from the two methods are essentially the same.

3. Protein in solution

The third benchmark test is a biological sample, i.e., a protein solution (see more details in the Appendix A). The protein studied is lysozyme [Fig. 7(a)], which is a model system widely used for testing new simulation and experimental protocols for their applications on biological systems. Here, the characteristic structural information of the proteins we are testing is the root mean square deviation (RMSD) of the backbone atoms relative to the initial structure for starting the MD simulations and the surface accessible surface area (SASA) [Figs. 7(b) and 7(c)], while the dynamics of the biomolecule is characterized by the root mean squared fluctuations (RMSF) [Fig. 7(d)]. Definitions of the RMSD, SASA, and RMSF are given in Appendix E. Moreover, we also examine the functional phase space sampled by the two methods. Lysozyme is a two-domain protein, where the two domains conduct a hinge-bending motion in order to facilitate the enzyme to break down the bacterial cell wall.³⁸ The distance between two residues (C54 and C97) is often used to characterize the openingclosing status of the two domains, whose distribution can be used to measure the broadness of the functional phase space and is presented in Fig. 7(e).³⁹ As it can be seen, the structure, dynamics, and functional space of the protein obtained by the PME are accurately reproduced by the RBE. More details on the data production are available at Appendix A.

We remark that there are also many valuable criteria worth an investigation to observe the performance of our algorithm, e.g., the symmetry-preserving mean-field condition presented recently.^{31,40} Moreover, the parameter *P* is determined empirically and is difficult to be given theoretically. This may also be related to these criteria. We will report further demonstrations of these criteria in our subsequent work.

IV. DISCUSSIONS

The RBE method inherits the advantages of Ewald-based methods but employs the random mini-batch idea and importance sampling technique to calculate the Fourier series of the long-range interaction. As a result, it achieves O(N) computational complexity and gains almost linear scalability for parallel computing, outperforming existing electrostatic algorithms. When very few computational cores were employed, the acceleration of the RBE mainly comes from the relatively small batch size of k in calculating $F_{i,1}$. With the increase in the number of CPU processors, the acceleration mainly comes from less communication. This is because six sequential rounds of communication are generally required to perform forward and backward Fourier transforms for the classical FFT, whereas only one global communication is required for the RBE with O(1) data transfer.

When the system is large [e.g., 100×10^6 atoms in Fig. 2(d)], the parallel efficiency of the RBE remains over 95% for 10 000 CPU cores. When the number of atoms assigned to each processor becomes smaller [Figs. 2(a)-2(d)], the parallel efficiency of the RBE goes down but remains much better than that of the PPPM. The reduction in the RBE efficiency in small systems results from the following reasons: First, the most time-consuming part transforms gradually from intra-processor calculation to inter-processor communication for small systems. Second, the frequency of the "loop remainder" in the vectorization increases while the number of particles/batches is not the multiple of the vector width. Different instructions can be used to handle the "loop remainder" in order to prevent out-of-bounds memory access along with other issues; however, additional costs are introduced. Third, the cost for other parts, such as memory load and load balance, will become significant when the one used for arithmetic operation of the Coulomb interactions becomes small.

In addition, the speedup of the RBE is constrained by the cost of the real-space calculation, i.e., the computation of $F_{i,2}$. In many works employing FFT-based methods as their electrostatic solver, the real-space cutoff is balanced such that the costs of real and Fourier spaces are approximately the same. The real-space cutoff for the RBE should be made smaller when one accelerates the calculation in the Fourier space. The calibration of the optimized short-range algorithm for the non-covalent bonds is our ongoing project. Furthermore, if other optimized techniques for real-space cutoffs are combined with the RBE, it is possible to obtain even better acceleration. The research of all-atom simulations coupling the RBE and the recently developed random batch list (RBL) method⁴¹ shall be studied in our subsequent works.

The time comparisons provided in Fig. 2 are only for the calculation of long-range interactions. In practice, the calculation of the Lennard-Jones (LJ) force is often done together with that of the Coulomb force for saving time. Other operations, including thermostat, bond angle, construction of the neighbor list, data statistics, timekeeping, and diagnostic routine, have different requirements in various systems. Generally, costs of these parts are not obvious but will have considerable impact on top of the acceleration of longrange interactions. We look forward to optimize these components in mainstream packages, which may become a bottleneck in the future.

Moreover, the RBE method discussed here is different from any coarse-grained or enhanced sampling methods, e.g., Gaussian accelerated molecular dynamics,^{42,43} replica exchange,⁴⁴ umbrella samplings,^{45,46} and meta dynamics,⁴⁷ as it mainly provides an efficient solution for calculating the Coulomb interactions and accelerates the simulation without losing any dynamical information on all time and length scales. It, thus, can be well combined with any coarse-grained and enhanced sampling methods to further speed up and scale up the simulation systems.

It is remarked that our exploration of the RBE method is limited to full periodic boundary condition. If the system is partially periodic in some directions with Dirichlet 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.^{48–50}

V. CONCLUSIONS

In summary, we have reported an efficient RBE algorithm to evaluate the Coulomb interactions in all-atom molecular dynamics simulations and demonstrated that it can greatly improve the computational efficiency and scalability for large-scale simulations in supercomputers while maintaining the same level of accuracy. These advantages of the RBE algorithm owe to the introduction of the random mini-batch idea, which avoids the use of the FFT and significantly reduces the communication cost and the computational complexity in parallel computation. This novel algorithm will be promising for MD simulations in modern architecture and communication protocols.

ACKNOWLEDGMENTS

The authors acknowledge the Center for High Performance Computing at Shanghai Jiao Tong University for the computing resources. All the authors were supported by the Shanghai Science and Technology Commission (Grant Nos. 20JC1414100 and 21JC1403700). J.L., Y.Z., and Z.X. were supported by the National Natural Science Foundation of China (Grant No. 12071288). P.T. and L.H. were supported by the National Natural Science Foundation of China (Grant Nos. 11974239, 12104295, and 31630002), the Innovation Program of Shanghai Municipal Education Commission, and the Shanghai Jiao Tong University Multidisciplinary Research Fund of Medicine and Engineering (Grant No. YG 2016QN13). S.J. was supported by the National Natural Science Foundation of China (Grant No. 12031013). L.L. was supported by the National Natural Science Foundation of China (Grant Nos. 11901389 and 11971314) and the Shanghai Sailing Program (Grant No. 19YF1421300). The authors thank the anonymous referees for their helpful suggestions, which improved the quality of the manuscript.

AUTHOR DECLARATIONS

Conflict of Interest

The authors have no conflicts to disclose.

Author Contributions

J.L. and P.T. contributed equally to this work.

DATA AVAILABILITY

The data that support the findings of this study are available from the corresponding authors upon reasonable request.

APPENDIX A: EXPERIMENTAL DESIGN

The bulk water system employs the classical SPC/E model in LAMMPS (version 7 August 2019). The SPC/E specifies a threesite rigid water molecule with charges and Lennard-Jones (LJ) parameters assigned to each of the three atoms. Electrostatic interaction is modeled using Coulomb's law, and the dispersion and repulsion forces use the LJ potential. Four cubic simulation boxes of different sizes, 8.16 nm (53 367 atoms), 14.61 nm (311 469 atoms), 31.09 nm (3 000 000 atoms), and 101.01 nm (100 158 744 atoms), respectively, were used and specified with periodic boundary conditions. The equilibration process was carried out for 500 ns in the NPT ensemble at 298 K and 1 bar with the PPPM, followed by 200 ns NVT production MD simulation for data collection with the PPPM and RBE, respectively. The time integration is performed on Nosé-Hoover style non-Hamiltonian equations of motions at a temperature coupling time parameter $T_{damp} = 5$ fs, and the scheme in LAMMPS closely follows the time-reversible measure-preserving

Verlet and rRESPA integrators.⁵¹ The velocity is initially generated according to a Maxwell distribution function at 298 K. All chemical bonds are converted to constraints using the SHAKE algorithm to allow a time step of 1 fs.⁵² During the equilibration process, the short-range part of the Coulomb interaction and the LJ interaction each with a cutoff parameter of 0.9 nm is considered with periodic boundary conditions. The splitting parameter α of the RBE is the same as PPPM's automatic tuning value, and the number of minibatches is set to P = 100. Some important physical properties are investigated to compare the RBE with the PPPM. The RDF, describing how the density of surrounding matter varies as a function of the distance from a point, is a frequently used measurement to analyze the structure of the system.⁵³ The simulation test in the NPT ensemble with RBE was carried out at 298 K and 1 bar, using the Nosé-Hoover thermostat and C-rescale barostat with the coupling time 0.1 ps for temperature coupling and 1 ps for pressure coupling, in GROMACS (version 2021.1). The cutoff radius of the short-range Coulomb interaction and Lennard-Jones interaction is 1.2 nm with the splitting parameter α = 4.2. The number of mini-batches is set to P = 100.

The LiTFSI ionic liquid employs the optimized potentials for liquid simulations all-atom (OPLS-AA)²⁴ force field for Li⁺, the TIP3P model for water molecules,⁵⁴ and the force field⁵⁵ developed for TFSI⁻. The system is equilibrated in the NPT ensemble with the PME at 298 K and 1 bar for 500 ns, followed by 200 ns production MD in the NVT using the Nosé–Hoover thermostat with the PME and RBE, respectively. The system contains 126 424 atoms, including 2560 Li⁺, 2560 TFSI⁻, and 28 488 H₂O. A cubic simulation box of size 11.46 nm was initially used with periodic boundary conditions in GROMACS (version 2020.4). The cutoff radius of the short-range Coulomb interaction and Lennard-Jones interaction is 1.2 nm with the splitting parameter $\alpha = 4.2$. The number of mini-batches is set to P = 500.

The protein solution employs the charmm27²⁴ force field for Lysozyme molecules, and the TIP3P model for water molecules.⁵⁴ The system contains 38 376 atoms, including 12 136 water molecules and the rest being protein molecules. Additional 0.1M NaCl is added into the system to describe the physiological condition. The system is equilibrated in the NPT ensemble with the PME at 298 K and 1 bar for 500 ns, followed by 200 ns production MD in the NVT ensemble using the Nosé–Hoover thermostat with the PME and RBE, respectively. A cubic simulation box of size 7.3 nm was initially used with periodic boundary conditions in GROMACS (version 2020.4). The cutoff of the short-range Coulomb interaction and Lennard-Jones interaction is 1.2 nm with splitting parameter $\alpha = 4.2$. The number of mini-batches is set to P = 500.

APPENDIX B: HARDWARE AND SOFTWARE

The computations in this paper were run on the π 2.0 cluster supported by the Center for High Performance Computing at Shanghai Jiao Tong University. Each CPU node contains two Intel Xeon Scalable Cascade Lake 6248 (2.5 GHz, 20 cores) and 12 × Samsung 16 GB DDR4 ECC REG 2666 memory. The tests using 10 000 CPU cores in this paper employ 250 such nodes. The computing networks are connected using 100 Gbps Intel Omni-Path, which is a high-speed interconnection network technology, and with this network, the communication cost of both the RBE and PPPM are

ARTICLE

significantly reduced. We believe that the hardware employed in such a way accurately reflects the proportion of the cost of near-field and far-field on modern computer cluster architecture. The intelparallel-studio/cluster.2020.1-intel-19.1.1 is used as the compiler and the LAMMPS is compiled using "make intel_cpu_intelmpi." The GROMACS is compiled using the same Intel package.

APPENDIX C: THE DEFINITION OF THE HYDROGEN BOND AUTOCORRELATION FUNCTION (HBACF)

We define the HBACF by the following function:⁵⁶

$$HBACF(t) = \frac{\langle h(0)h(t)\rangle}{\langle h\rangle}dt,$$
 (C1)

where the variable h(t) is one if the specified pair of water molecules is hydrogen bonded at time t, otherwise it is zero. The hydrogen bond is determined by the geometric criterion, i.e., when the distance between the donor hydrogen of one water molecule and the acceptor oxygen of another water molecule is smaller than 0.35 nm, and the angle of hydrogen-donor-acceptor is smaller than 30° . The bracket means averaging over time and all the pairs of water molecules. It is remarked that the HBACF might also depend on the parameters associated with the temperature and pressure control.

APPENDIX D: THE CALCULATION METHOD OF HEAT CAPACITY AND DIELECTRIC CONSTANT

The isochoric heat capacity (C_v) is calculated by the equation

$$C_v = \frac{\sigma_E^2}{2k_{\rm B}T},\tag{D1}$$

where σ_E is the standard deviation of the total energy in the system in the NVT ensemble, k_B is Boltzmann constant, and *T* is the temperature. The isobaric heat capacity (C_p) is calculated by the equation

$$C_p = \frac{\sigma_H^2}{2k_{\rm B}T},\tag{D2}$$

where σ_H is the standard deviation of the enthalpy in the system in the NPT ensemble. The dielectric constant is calculated by the equation

$$\varepsilon_r = 1 + \frac{4\pi}{3k_{\rm B}TV} (\langle M^2 \rangle - \langle M \rangle^2),$$
 (D3)

where V is the mean volume of the system and M is the total dipole moment of the system.

APPENDIX E: THE CALCULATION METHOD OF SASA, RMSD, AND RMSF OF the PROTEIN

The solvent accessible surface area (SASA) of the protein molecule is calculated by rolling a sphere with a radius of the solvent probe of 1.4 Å over the surface of the protein.⁵⁷

The root mean square deviation (RMSD) of a certain structure (t) to a reference structure is calculated by least-square fitting the structure to the reference structure and subsequently calculating the RMSD as

$$RMSD(t) = \left[\frac{1}{M}\sum_{i=1}^{N} m_i |r_i(t) - r_i(0)|^2\right]^{\frac{1}{2}},$$
 (E1)

where $M = \sum_{i=1}^{N} m_i$ and $r_i(t)$ is the position of atom *i* at time *t*.

The root mean square fluctuation (RMSF) is the standard deviation of atomic positions after least-square fitting to a reference structure,

$$RMSF(i) = \sqrt{\frac{1}{\tau} \sum_{t=1}^{\tau} (r_{i,t} - \langle r_i \rangle)^2}$$
(E2)

where $r_{i,t}$ is the position of atom *i* at time *t*, $\langle r_i \rangle$ is the average position of atom *i*, and \mathcal{T} is the total simulation time of the trajectory.

REFERENCES

¹A. T. Brünger, J. Kuriyan, and M. Karplus, Science 235, 458 (1987).

²M. Karplus and G. A. Petsko, Nature 347, 631 (1990).

- ³V. Yamakov, D. Wolf, S. R. Phillpot, A. K. Mukherjee, and H. Gleiter, Nat. Mater. 1, 45 (2002).
- ⁴M. P. Allen and D. J. Tildesley, *Computer Simulation of Liquids* (Oxford University Press, 2017).
- ⁵S. A. Hollingsworth and R. O. Dror, Neuron **99**, 1129 (2018).

⁶R. H. French, V. A. Parsegian, R. Podgornik, R. F. Rajter, A. Jagota, J. Luo, D. Asthagiri, M. K. Chaudhury, Y.-m. Chiang, and S. Granick, Rev. Mod. Phys. 82, 1887 (2010).

⁷D. A. Walker, B. Kowalczyk, M. O. de la Cruz, and B. A. Grzybowski, Nanoscale 3, 1316 (2011).

⁸T. Darden, D. York, and L. Pedersen, J. Chem. Phys. 98, 10089 (1993).

⁹ R. W. Hockney and J. W. Eastwood, *Computer Simulation Using Particles* (CRC Press, 1988).

- ¹⁰J. Barnes and P. Hut, Nature **324**, 446 (1986).
- ¹¹L. Greengard and V. Rokhlin, J. Comput. Phys. 73, 325 (1987).
- ¹²L. Greengard, <u>Science</u> **265**, 909 (1994).
- ¹³U. Trottenberg, C. W. Oosterlee, and A. Schuller, *Multigrid* (Elsevier, 2000).
- ¹⁴A. C. Maggs and V. Rossetto, Phys. Rev. Lett. 88, 196402 (2002).
- ¹⁵A. Arnold, F. Fahrenberger, C. Holm, O. Lenz, M. Bolten, H. Dachsel, R. Halver, I. Kabadshow, F. Gähler, and F. Heber, Phys. Rev. E 88, 063308 (2013).
- ¹⁶S. Jin, L. Li, Z. Xu, and Y. Zhao, SIAM J. Sci. Comput. **43**, B937 (2021).
- ¹⁷H. Robbins and S. Monro, Ann. Math. Stat. **22**, 400 (1951).

¹⁸L. Bottou, in On-Line Learning in Neural Networks (Cambridge University Press, 1999), pp. 9–42.

- ¹⁹S. Jin, L. Li, and J.-G. Liu, J. Comput. Phys. 400, 108877 (2020).
- ²⁰L. Li, Z. Xu, and Y. Zhao, SIAM J. Sci. Comput. **42**, A1486 (2020).
- ²¹ P. P. Ewald, Ann. Phys. **369**, 253 (1921).
- ²² J. J. Benedetto and G. Zimmermann, J. Fourier Anal. Appl. 3, 505 (1997).
- ²³W. K. Hastings, Biometrika 57, 97 (1970).
- ²⁴B. Hess, C. Kutzner, D. van der Spoel, and E. Lindahl, J. Chem. Theory Comput. 4, 435 (2008).
- ²⁵S. Plimpton, J. Comput. Phys. 117, 1 (1995).

²⁶W. Michael Brown, J.-M. Y. Carrillo, N. Gavhane, F. M. Thakkar, and S. J. Plimpton, Comput. Phys. Commun. **195**, 95 (2015).

²⁷A. P. Thompson, S. J. Plimpton, and W. Mattson, J. Chem. Phys. **131**, 154107 (2009).

- ²⁸S. Jin, L. Li, and J.-G. Liu, SIAM J. Numer. Anal. **59**, 746 (2021).
- ²⁹ M. Deserno and C. Holm, J. Chem. Phys. **109**, 7694 (1998).

³⁰J.-P. Hansen and I. R. McDonald, *Theory of Simple Liquids: With Applications to Soft Matter* (Academic Press, 2013).

³¹Z. Hu, arXiv:2110.12795 (2021).

- ³²M. Bernetti and G. Bussi, J. Chem. Phys. 153, 114107 (2020).
- ³³J. Liang, P. Tan, L. Hong, S. Jin, Z. Xu, and L. Li, arXiv:2110.14362 (2021).

³⁴H. J. C. Berendsen, J. R. Grigera, and T. P. Straatsma, J. Phys. Chem. **91**, 6269 (1987).

³⁵J. Kolafa and J. W. Perram, Mol. Simul. 9, 351 (1992).

³⁶O. Borodin, L. Suo, M. Gobet, X. Ren, F. Wang, A. Faraone, J. Peng, M. Olguin, M. Schroeder, and M. S. Ding, ACS Nano 11, 10462 (2017).

³⁷ J. Lim, K. Park, H. Lee, J. Kim, K. Kwak, and M. Cho, J. Am. Chem. Soc. 140, 15661 (2018).

³⁸H. R. Faber and B. W. Matthews, Nature **348**, 263 (1990).

³⁹Y. Wang and H. P. Lu, J. Phys. Chem. B 114, 6669 (2010).

⁴⁰C. Pan, S. Yi, and Z. Hu, Phys. Chem. Chem. Phys. 21, 14858 (2019).

⁴¹ J. Liang, Z. Xu, and Y. Zhao, J. Chem. Phys. 155, 044108 (2021).

⁴²Y. Miao, V. A. Feher, and J. A. McCammon, J. Chem. Theory Comput. **11**, 3584 (2015).

⁴³J. Wang, P. R. Arantes, A. Bhattarai, R. V. Hsu, S. Pawnikar, Y. M. Huang, G. Palermo, and Y. Miao, Wiley Interdiscip. Rev.: Comput. Mol. Sci. 11, e1521 (2021).

⁴⁴Y. Sugita and Y. Okamoto, Chem. Phys. Lett. **314**, 141 (1999).

⁴⁵S. Kumar, J. M. Rosenberg, D. Bouzida, R. H. Swendsen, and P. A. Kollman, J. Comput. Chem. 13, 1011 (1992). ⁴⁶J. Kästner, Wiley Interdiscip. Rev.: Comput. Mol. Sci. 1, 932 (2011).

⁴⁸J. Liang, J. Yuan, E. Luijten, and Z. Xu, J. Chem. Phys. **152**, 134109 (2020).

⁴⁹O. Maxian, R. P. Peláez, L. Greengard, and A. Donev, J. Chem. Phys. **154**, 204107 (2021).

⁵⁰J. Yuan, H. S. Antila, and E. Luijten, J. Chem. Phys. 154, 094115 (2021).

⁵¹M. E. Tuckerman, J. Alejandre, R. López-Rendón, A. L. Jochim, and G. J. Martyna, J. Phys. A: Math. Gen. **39**, 5629 (2006).

⁵²V. Kräutler, W. F. Van Gunsteren, and P. H. Hünenberger, J. Comput. Chem. 22, 501 (2001).

⁵³D. Frenkel and B. Smit, Understanding Molecular Simulation: From Algorithms to Applications (Elsevier, 2001), Vol. 1.

⁵⁴D. J. Price and C. L. Brooks III, J. Chem. Phys. **121**, 10096 (2004).

⁵⁵J. N. Canongia Lopes and A. A. H. Pádua, J. Phys. Chem. B **108**, 16893 (2004).

⁵⁶A. Luzar, J. Chem. Phys. 113, 10663 (2000).

⁵⁷F. Eisenhaber, P. Lijnzaad, P. Argos, C. Sander, and M. Scharf, J. Comput. Chem. 16, 273 (1995).

⁴⁷G. Bussi and A. Laio, Nat. Rev. Phys. 2, 200 (2020).