



HAL
open science

Markov-chain sampling for long-range systems without evaluating the energy

Gabriele Tartero, Werner Krauth

► **To cite this version:**

Gabriele Tartero, Werner Krauth. Markov-chain sampling for long-range systems without evaluating the energy. *The Journal of Chemical Physics*, 2024, Monte Carlo methods, 70 years after Metropolis et al. (1953), 161 (9), pp.094106. 10.1063/5.0225561 . hal-04781028

HAL Id: hal-04781028

<https://hal.science/hal-04781028v1>

Submitted on 13 Nov 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Markov-chain sampling for long-range systems without evaluating the energy

Gabriele Tartero^{1,*} and Werner Krauth^{1,2,3,†}

¹*Laboratoire de Physique de l'École normale supérieure, ENS, Université PSL,
CNRS, Sorbonne Université, Université Paris Cité, Paris, France*

²*Rudolf Peierls Centre for Theoretical Physics, Clarendon Laboratory, Oxford OX1 3PU, UK*

³*Simons Center for Computational Physical Chemistry, New York University, New York (NY), USA*

(Dated: June 4, 2024)

In past decades, enormous effort has been expended to develop algorithms and even to construct special-purpose computers in order to efficiently evaluate total energies and forces for long-range-interacting particle systems, with the particle-mesh Ewald and the fast multipole methods as well as the “Anton” series of supercomputers serving as examples for biomolecular simulations. Cutoffs in the range of the interaction have also been used for large systems. All these methods require extrapolations. Within Markov-chain Monte Carlo, in thermal equilibrium, the Boltzmann distribution can however be sampled natively without evaluating the total interaction potential. Using as an example the Lennard-Jones interaction, we review past attempts in this direction, and then discuss in detail the class of cell-veto algorithms which make possible fast, native sampling of the Boltzmann distribution without any approximation, extrapolation, or cutoff even for the slowly decaying Coulomb interaction. The computing effort per move remains constant with increasing system size, as we show explicitly. We provide worked-out illustrations and pseudocode representations of the discussed algorithms. Python implementations are made available in an associated open-source software repository.

I. INTRODUCTION

One of the aims of computational science is to analyze, and often solve, model systems in fields from subatomic particles to condensed-matter physics and chemistry, to galaxies and to the universe. The power of modern computers appears without limits when compared to the early electronic devices from only two generations ago. Algorithms have also much developed, as has the ease with which computers are interfaced and their output processed. Nevertheless, certain limits persist. The first, broadly speaking, is a limitation in space: One tries to simulate large, possibly infinite systems, but is restricted by the finite extent of computer memories. Second, one may aim for long-time simulations but the necessary evolution equations may not allow one to reach the relevant temporal scales. The third limitation is in the type of couplings between particles or fields. They may be of many-body nature, including quantum and long-range potentials and may require even today require simplifications [1].

In classical condensed-matter particle systems in thermal equilibrium, the limitations on space, time, and what could be called “range” can again be illustrated. In contexts from statistical mechanics to molecular simulation in biochemistry, one may want to analyze large, almost infinite systems, but only in exceptional cases is it possible to simulate them directly [2–6]. Prominent coarse-graining strategies were developed [7] in order to reach larger and larger sizes. On the other hand, it is the change of behavior of finite systems with size, the famous

finite-size scaling, which often provides crucial information on a physical model [8]. One also may want to reach essentially infinite simulation times, as the equilibrium state is approached from an initial configuration only in this limit. Computer computations are by definition of finite duration and, within Markov-chain Monte Carlo, it is only in exceptional cases possible to reach the infinite-time limit directly [9, 10].

The approximation of complex interaction potentials, the above “range” limitation, is what we are concerned with in this paper, again in the framework of equilibrium physics. We consider systems of particles at a given temperature, and governed by the Boltzmann distribution. Such systems can be simulated by molecular dynamics or sampled by the Monte Carlo method. Particle systems with Coulomb or Lennard-Jones interactions will serve as examples. The former decay as $1/r$ with the distance r between charged particles. Its long-range nature derives from the masslessness of the photon. The Lennard-Jones potential describes the interaction between uncharged, non-polar atom. Repulsive at small distances, its attractive $1/r^6$ behavior at large distances describes the London dispersive force [11]. For both the Coulomb and the Lennard-Jones potentials, the long-range nature is thus rooted in profound physical principles, and there is a strong incentive to maintain it in the modeling. However, long-range potentials pose severe problems for computation, which arise, naively, because the total energy of a system of N particles interacting in pairs consists of $N(N-1)/2$ terms. Moving one particle changes $N-1$ terms in the total energy. Likewise, the force on a given particle is composed of $N-1$ terms.

Computational science, in the past decades, has focused on how to evaluate long-ranged potentials or forces. For the case of the Coulomb interaction, the potential or

* gabriele.tartero@phys.ens.fr

† werner.krauth@ens.fr

the forces are computed by interpolating charges to a grid, then solving Poisson’s equation in discretized momentum space using fast Fourier transform. The “Anton” series of supercomputers has been designed in order to optimize the interpolation and force evaluation [12], yielding spectacular speedups [13]. The discretization error of these so-called particle-mesh Ewald methods disappears only for vanishing grid size. The fast multipole method [14] presents another extrapolation of the interaction in terms of moments of the charge distribution. Very often also, the Coulomb interaction is cut off beyond a certain radius, although this profoundly modifies the underlying physics of the models. Discretization and cutoffs both call for extrapolations. For decades, the Lennard-Jones potential was cut off at a finite distance, although artifacts introduced by the cutoff on phase boundaries [15, 16] and on interface effects [17, 18] were pointed out repeatedly. In recent years, this problem has been identified, and the particle-mesh Ewald approach extended to the Lennard-Jones system [19]. However, this creates an avalanche of problems, as the empirical interaction potentials used for the sampling are themselves fitted (originally with a cutoff) and then have to be reparametrized [20].

A Monte Carlo algorithm typically consists in a sequence of simple proposed moves, that are sampled from a certain probability distribution, and that are either accepted or rejected [21]. There is much liberty in which moves to propose, and how to accept or reject them with a “filter”. The only condition to be satisfied is that the Boltzmann distribution be stationary with respect to the move set encoded in a transition matrix. As the proposed move is simple, Markov-chain Monte Carlo is a decision problem in nature.

In this paper, we discuss methods for sampling the Boltzmann distribution $\exp(-\beta U)$ without evaluating U . This approach is conceptually distinct from the previously mentioned computational focus. It leads to the fast *native* sampling (without any cutoffs, discretizations or time-stepping errors) for long-range potentials. In the main part of the paper, we discuss in detail the cell-veto class of algorithms [22], which rely on two key ingredients: First, the aforementioned accept/reject decision is split, using the factorized Metropolis filter [23], into a large number of independent factor decisions, which are coordinated through a consensus principle rather than through the computation of the system energy or its derivatives. Second, the accept/reject decision for a factor is generally undertaken in two steps using two “pebbles” [24] that correspond to the thinning of an inhomogeneous Poisson process [25]: first, in an approximate (conservative) fashion, that provisionally rejects “too many moves”, and, second, in a correction step that accepts some of the provisional rejects such that the overall rejection probability is correct. We provide a step-by-step introduction to this class of algorithms, together with heuristics, pseudocodes and simplified Python implementations. We thus hope to facilitate access to a class of algorithms which have

remained cryptic and have never been presented in the context of decision problems and their possible extensions [26–29].

II. FROM LONG-RANGE POTENTIALS TO THE CELL-VETO ALGORITHM

We concentrate on long-range pair potentials, leaving aside many-body terms that can be treated as well [29, 30]. Several strategies for treating the long-range nature are discussed in Sec. II A, such as solving the accept/reject decision with approximate potentials [26], and as factorizing the total energy into short-range and long-range contributions [27]. The factorized Metropolis filter generalizes this approach (Sec. II B). Together with the concept of bounding potential, it lies at the heart of the native $\mathcal{O}(1)$ cell-veto algorithms, of which we present a naive version in Sec. II C before sharpening it in Sec. III.

A. Potentials, Metropolis algorithm

Specifically, we consider throughout this paper N particles inside a two-dimensional periodic box of size L interacting *via* the Lennard-Jones interaction [31]:

$$U^{\text{LJ}}(r) = 4\epsilon \left[\left(\frac{\sigma}{r}\right)^{12} - \left(\frac{\sigma}{r}\right)^6 \right], \quad (1)$$

At large distances, its attractive tail vanishes as $1/r^6$ (see Fig. 1). Our sample programs are for this two-dimensional system, but they generalize easily to higher dimensions and to arbitrary long-range interactions (see Appendix A). For the most part, a cutoff means that




FIG. 1: Lennard-Jones potential with short-range and long-range parts separated by a cutoff r_c and a shifted cutoff variant which vanishes beyond r_c . For large r , $U^{\text{LJ}}(r) \sim r^{-6}$ is attractive.

one can divide the potential U into a short-range and a long-range contribution:

$$U(r) = \underbrace{U(r)\Theta(r_c - r)}_{U^{\text{short}}} + \underbrace{U(r)\Theta(r - r_c)}_{U^{\text{long}}}, \quad (2)$$

where $\Theta(z)$ is the unit-ramp function, which is zero for negative z and one for positive z . The Boltzmann distribution of the potential U then becomes the product of Boltzmann distributions of the constituents:

$$\pi(\mathbf{X}) = \pi^{\text{short}}(\mathbf{X})\pi^{\text{long}}(\mathbf{X}). \quad (3)$$

Commonly, a cutoff replaces $U(r)$ by its short-range part, and the long-range part is set to zero (see Fig. 1, for the Lennard-Jones potential). In that case, it is customary to shift U^{short} , so that the approximate potential is continuous. More elaborate procedures are common in order to render the approximate potential better behaved around r_c , as often required for molecular dynamics.




FIG. 2: Decisions in MCMC. In the Metropolis filter of Eq. (4), the accept/reject decision for $e^{-\beta U}$ can often be made from an approximation $e^{-\beta U^{\text{short}}}$ with error bounds.




FIG. 3: Multi-step Metropolis algorithm. In Alg. 1, moves are first made with U^{short} only. Initial and final positions are stored. The proposed move $\mathbf{Y}^{\text{old}} \rightarrow \mathbf{Y}^{\text{new}}$ is subject to the Metropolis filter with U^{long} .

In the Metropolis algorithm, a random particle at position $\mathbf{x} \in \mathbf{X}$ is often selected for random displacement to $\mathbf{x}' = \mathbf{x} + \Delta\mathbf{x}$. The move from \mathbf{X} to \mathbf{X}' (in the latter, \mathbf{x} is replaced by \mathbf{x}') is accepted with a probability given by the Metropolis filter

$$p^{\text{Met}}(\mathbf{X} \rightarrow \mathbf{X}') = \min[1, \exp(-\beta\Delta U)]. \quad (4)$$

With a system-size independent cutoff, the number of factors that have to be considered for each move remains constant. These factors can be identified with a standard

cell system and evaluated in $\mathcal{O}(1)$ operations. However, the essential long-range nature of the potential is then sacrificed.

Strategies have been devised to improve the $\mathcal{O}(N)$ scaling of the Metropolis algorithm natively. One natural approach exploits the decision-problem nature of the Metropolis filter. The Bernoulli distribution giving rise to Eq. (4) is sampled with a uniform random number $\text{ran}(0,1) = \Upsilon$, and the move is accepted if $\Upsilon < \exp(-\beta U)$ and rejected otherwise. For a given dynamical cutoff r_c , the accept/reject decision may be taken on the basis of $U^{\text{short}}(r_c)$ if the neglected long-range contribution $U^{\text{long}}(r_c)$ can be proven not to change the decision (see Fig. 2). The cutoff r_c is then increased if required, adding particles in concentric layers around \mathbf{x} or \mathbf{x}' . Under mild conditions on the maximum local density, this approach can be made rigorous [26]. However, the sheer number of evaluations (which can easily reach 10^{10} or 10^{12}) renders it difficult to implement. Also, dynamical cutoffs do not allow for long-range potentials which, in d -dimensional space, decay as r^{-d} or slower, as the contribution of U^{long} then remains unbounded for large r_c [22, 32].

In the multi-time-step Metropolis algorithm [27], the Lennard-Jones system again splits the potential U into long-range and short-range contributions, performing a series of n_s provisional moves on the basis of U^{short} only. This amounts to proposing a move of $k < N$ particles (with coordinates \mathbf{Y} , see Fig. 3) from \mathbf{Y}^{old} to \mathbf{Y}^{new} . These “short-range” moves are then accepted/rejected with the Metropolis filter for U^{long} only. Calling for simplicity the original configuration \mathbf{X} (with the k particles at \mathbf{Y}^{old}) and the target configuration \mathbf{X}' (with the k particles at \mathbf{Y}^{new}), the composite probabilities are

$$p(\mathbf{X} \rightarrow \mathbf{X}') = p^{\text{short}}(\mathbf{X} \rightarrow \mathbf{X}') \min \left[1, \frac{\pi^{\text{long}}(\mathbf{X}')}{\pi^{\text{long}}(\mathbf{X})} \right] \quad (5)$$

$$p(\mathbf{X}' \rightarrow \mathbf{X}) = p^{\text{short}}(\mathbf{X}' \rightarrow \mathbf{X}) \min \left[1, \frac{\pi^{\text{long}}(\mathbf{X})}{\pi^{\text{long}}(\mathbf{X}')} \right].$$

Taking the ratio in Eq. (5) and noting that p^{short} , as a sequence of n_s reversible moves, satisfies the detailed-balance condition with respect to π^{short} :

$$\pi^{\text{short}}(\mathbf{X})p^{\text{short}}(\mathbf{X} \rightarrow \mathbf{X}') = \pi^{\text{short}}(\mathbf{X}')p^{\text{short}}(\mathbf{X}' \rightarrow \mathbf{X}),$$

we see that the full transition matrix satisfies detailed balance with respect to the full Boltzmann distribution $\pi = \pi^{\text{short}}\pi^{\text{long}}$

$$\pi(\mathbf{X})p(\mathbf{X} \rightarrow \mathbf{X}') = \pi(\mathbf{X}')p(\mathbf{X}' \rightarrow \mathbf{X}), \quad (6)$$

so that the algorithm is correct (see Alg. 1 (**multi-step-metropolis**) for an implementation).

In Alg. 1 (**multi-step-metropolis**), short-range moves can be implemented in $\mathcal{O}(1)$ with an appropriate cell system [21, Sect. 2.4.1]). The construction of the sets \mathbf{Y}^{old} and \mathbf{Y}^{new} is also $\mathcal{O}(1)$ per element, so that the

complexity of the entire short-range loop is $\mathcal{O}(n_s)$. However, the complexity of the long-range decision in Alg. 1 is $\mathcal{O}(n_s \times N)$, with an acceptance probability that plummets with increasing n_s , a parameter that should therefore be chosen to be as small as possible. This leads to $n_s = 1$, and this defeats the initial intention.

```

procedure multi-step-metropolis
input  $\mathbf{X}$  (configuration at time  $t$ )
 $\mathbf{Y}^{\text{old}} \leftarrow \emptyset$ ;  $\mathbf{Y}^{\text{new}} \leftarrow \emptyset$ 
for  $i = 1, \dots, n_s$ : (short-range steps)
   $\mathbf{x} \leftarrow \text{choice}(\mathbf{X})$  (random particle)
   $U_{\mathbf{x}}^{\text{short}} \leftarrow \sum_{\mathbf{x}'' \in \mathbf{X} \setminus \{\mathbf{x}\}: |\mathbf{x}'' - \mathbf{x}| < r_c} U(|\mathbf{x}'' - \mathbf{x}|)$ 
   $\mathbf{x}' \leftarrow \mathbf{x} + \Delta \mathbf{x}$  (with  $|\Delta \mathbf{x}| < \delta$ )
   $U_{\mathbf{x}'}^{\text{short}} \leftarrow \sum_{\mathbf{x}'' \in \mathbf{X} \setminus \{\mathbf{x}\}: |\mathbf{x}'' - \mathbf{x}'| < r_c} U(|\mathbf{x}'' - \mathbf{x}'|)$ 
   $\Upsilon \leftarrow \text{ran}(0, 1)$ 
  if  $\Upsilon < \exp[-\beta(U_{\mathbf{x}'}^{\text{short}} - U_{\mathbf{x}}^{\text{short}})]$ :
     $\mathbf{X} \leftarrow \{\mathbf{x}'\} \cup \mathbf{X} \setminus \{\mathbf{x}\}$ ;  $\mathbf{Y}^{\text{new}} \leftarrow \mathbf{Y}^{\text{new}} \cup \{\mathbf{x}'\}$ 
    if  $\mathbf{x} \notin \mathbf{Y}^{\text{new}}$ :  $\mathbf{Y}^{\text{old}} \leftarrow \mathbf{Y}^{\text{old}} \cup \{\mathbf{x}\}$ 
    else:  $\mathbf{Y}^{\text{new}} \leftarrow \mathbf{Y}^{\text{new}} \setminus \{\mathbf{x}\}$ 
   $\Delta U^{\text{long}} \leftarrow \sum'_{\substack{\mathbf{x} \in \mathbf{Y}^{\text{new}}, \mathbf{x}' \in \mathbf{X} \\ |\mathbf{x} - \mathbf{x}'| > r_c}} U(|\mathbf{x} - \mathbf{x}'|) - \sum'_{\substack{\mathbf{x} \in \mathbf{Y}^{\text{old}}, \mathbf{x}' \in \mathbf{X} \setminus \mathbf{Y}^{\text{new}} \\ |\mathbf{x} - \mathbf{x}'| > r_c}} U(|\mathbf{x} - \mathbf{x}'|)$ 
   $\Upsilon \leftarrow \text{ran}(0, 1)$  (long-range decision)
  if  $\Upsilon > \exp(-\beta \Delta U^{\text{long}})$ :
     $\mathbf{X} \leftarrow \mathbf{Y}^{\text{old}} \cup \mathbf{X} \setminus \mathbf{Y}^{\text{new}}$ 
output  $\mathbf{X}$  (configuration at time  $t + 1$ )

```

Algorithm 1: multi-step-metropolis. Composite iteration of the algorithm of Ref. [27]. The calculation of $U_{\mathbf{x}}^{\text{short}}$ is $\mathcal{O}(1)$ with the use of a grid. The prime in \sum' eliminates double counting of pairs $(\mathbf{x}, \mathbf{x}')$ and $(\mathbf{x}', \mathbf{x})$. For large n_s , this algorithm has many long-range rejections.

Algorithm 1 (multi-step-metropolis), in spite of its serious limitations, illustrates the split of U into $U^{\text{short}}(r)$ and $U^{\text{long}}(r)$ that effectively leads to a factorization of the interaction. The concept will be extended in the following sections, where it leads to native $\mathcal{O}(1)$ long-range algorithms. The split of the interactions also echoes Hamiltonian Monte Carlo [33–35], where provisional moves $\mathbf{X} \rightarrow \mathbf{X}'$ are proposed from a sequence of molecular-dynamics iterations rather than from a number of Metropolis steps. In Hamiltonian Monte Carlo, accumulated finite-time-step errors are also eliminated by a Metropolis filter, and the Boltzmann distribution is again sampled without approximations. This works because the errors are small, but it fails for large long-range systems because the errors are extensive in N even for $n_s = 1$.

B. Factorized Metropolis filter

The multi-time-step Metropolis approach of Alg. 1 separates the interactions into two factors, one short-range and one long-range. More generally, potentials U can

often be written as

$$U(\mathbf{X}) = \sum_{M \in \mathcal{M}} U_M, \quad (7)$$

with factors M in a set \mathcal{M} , so that the Boltzmann distribution is

$$\pi(\mathbf{X}) = \exp[-\beta U(\mathbf{X})] = \prod_{M \in \mathcal{M}} \exp(-\beta U_M). \quad (8)$$

We can then replace the Metropolis filter of Eq. (4), written as

$$\mathcal{P}^{\text{Met}}(\mathbf{x} \rightarrow \mathbf{x}') = \min \left[1, \prod_{M \in \mathcal{M}} \exp(-\beta \Delta U_M) \right], \quad (9)$$

with the factorized Metropolis filter [23]

$$\mathcal{P}^{\text{fact}}(\mathbf{x} \rightarrow \mathbf{x}') = \prod_{M \in \mathcal{M}} \underbrace{\min[1, \exp(-\beta \Delta U_M)]}_{\mathcal{P}_M^{\text{fact}}(\mathbf{x} \rightarrow \mathbf{x}')}. \quad (10)$$

That the factorized filter also satisfies the detailed-balance condition can be shown just as in Eqs (5) and (6).

In our context, the use of pair factors $M = \{i, j\}$ with $U_M = U(|\mathbf{x}_i - \mathbf{x}_j|)$ is natural (even though the concept is easily generalizable [23, 24, 32]). A proposed move of a single particle then involves $N - 1$ factors (that is, $N - 1$ pairs). It is accepted by consensus, that is, if all factors accept it (see Alg. 2 (factorized-metropolis) and Fig. 4; see Ref. [24] for a general discussion of the consensus principle). We will show below how to implement this by evaluating $\mathcal{O}(1)$ pairs, thus with only $\mathcal{O}(1)$ operations.

```

procedure factorized-metropolis
input  $\mathbf{X}$  (configuration at time  $t$ )
 $\mathbf{x} \leftarrow \text{choice}(\mathbf{X})$  (random particle)
 $\mathbf{x}' \leftarrow \mathbf{x} + \Delta \mathbf{x}$  (with  $|\Delta \mathbf{x}| < \delta$ )
for  $\mathbf{x}'' \in \mathbf{X} \setminus \{\mathbf{x}\}$ :
   $\Upsilon \leftarrow \text{ran}(0, 1)$ 
  if  $\Upsilon > \exp[-\beta(U_{\mathbf{x}''\mathbf{x}'} - U_{\mathbf{x}''\mathbf{x}})]$ : goto 1
 $\mathbf{X} \leftarrow \{\mathbf{x}'\} \cup \mathbf{X} \setminus \{\mathbf{x}\}$ 
1 output  $\mathbf{X}$  (configuration at time  $t + 1$ )

```

Algorithm 2: factorized-metropolis. Proposing a move, and accepting it in case of consensus, also samples the Boltzmann distribution of Eq. (8).

C. Bounding the potential, cell-veto algorithm

The factorized Metropolis filter replaces the accept/reject decision based on the total energy $U(x)$ with independent decisions for all factor potentials U_M with $M \in \mathcal{M}$. This replacement carries enormous potential for speedup as, for a given distant pair of particles \mathbf{x} and \mathbf{x}'' , we need not systematically evaluate




FIG. 4: Factorized Metropolis filter using the consensus principle (see Alg. 2). (a): The active particle \mathbf{x} may interact with all other particles. (b,c,d, ...): Every factor of the factorized Metropolis filter in Eq. (10) independently decides whether to veto a proposed move. The conjunctions \wedge denote the consensus: the move $\mathbf{x} \rightarrow \mathbf{x}'$ is accepted if no factor vetoes it.

all the rejection probabilities [22], applying a principle called “thinning” [25]. Heuristically, let us suppose that, at a distance between \mathbf{x} and \mathbf{x}'' , the veto probability $q = 1 - \exp(-\beta\Delta U_M)$ can be bounded by $\epsilon \ll 1$. Rather than check the factor at each proposed move, we check it with probability ϵ (on average every $1/\epsilon$ times). Only when we check it, we evaluate q , and veto with probability q/ϵ . Because of $q = \epsilon \times (q/\epsilon)$, the overall veto probability is correct, but the bound has allowed us to drastically reduce the number of evaluations.




FIG. 5: For a pair of far-away particles \mathbf{x} and \mathbf{x}'' , the factor rejection probability for a move $\mathbf{x} \rightarrow \mathbf{x}'$ can be bounded by a constant q_Z depending on the target cell Z relative to the active cell 0.

The cell-veto algorithm [22] puts the above heuristic into practice by dividing the periodic simulation box into cells that usually contain at most one particle (see Fig. 5). Rather than the particle \mathbf{x}'' , one addresses a target cell Z relative to the active cell 0 containing the active particle \mathbf{x} . The upper bound q_Z for the rejection plays the role of ϵ in the heuristic. It satisfies:

$$q_Z \geq \max_{\substack{\mathbf{x} \in 0, \mathbf{x}'' \in Z \\ \mathbf{x}': |\mathbf{x}' - \mathbf{x}| < \delta}} \{1 - \exp[-\beta(U_{\mathbf{x}''\mathbf{x}'} - U_{\mathbf{x}''\mathbf{x}})]\}. \quad (11)$$

‘To yield the correct rejection probability $[1 - \exp[-\beta(U_{\mathbf{x}''\mathbf{x}'} - U_{\mathbf{x}''\mathbf{x}})]] q_Z$, a veto called by cell Z , with probability q_Z , must be confirmed with probability

$$\frac{[1 - \exp[-\beta(U_{\mathbf{x}''\mathbf{x}'} - U_{\mathbf{x}''\mathbf{x}})]]}{q_Z}, \quad (12)$$

Two complications must be solved. First, a cell Z neighboring the active cell 0 does not allow for a finite upper bound q_Z . Second, the cell Z may contain more than one particle. Both cases can be easily treated with the second requiring a set of “surplus” particles. The “two-pebble” decision [24], first for the cell Z , then for the particle $\mathbf{x}'' \in Z$, is illustrated in Alg. 3 (cell-veto), which also implements a first loop over neighbor and surplus particles, as in Alg. 2 (factorized-metropolis), and a cell-veto loop over far-away cells that form a set \mathcal{F} (see Fig. 6). This loop over the elements of \mathcal{F} is naive.

procedure cell-veto

```

input  $\mathbf{X}$  (configuration of  $N$  particles)
 $\mathbf{x} \leftarrow \text{choice}(\mathbf{X})$ 
 $\mathbf{x}' \leftarrow \mathbf{x} + \Delta\mathbf{x}$  (with  $|\Delta\mathbf{x}| < \delta$ )
for  $\mathbf{x}'' \in \{\text{neighbor}\} \cup \{\text{surplus}\}$ : (see Fig. 6b)
   $\Upsilon \leftarrow \text{ran}(0, 1)$ 
  if  $\Upsilon > \exp[-\beta(U_{\mathbf{x}''\mathbf{x}'} - U_{\mathbf{x}''\mathbf{x}})]$ : goto 1
for  $Z \in \mathcal{F}$ , non-empty: (loop over far-away cells  $\mathcal{F}$ )
   $\Upsilon_1 \leftarrow \text{ran}(0, 1)$ 
  if  $\Upsilon_1 > q_Z$ :
     $\Upsilon_2 \leftarrow \text{ran}(0, 1)$ 
     $\mathbf{x}'' \leftarrow \text{particle in cell } Z$ 
    if  $\Upsilon_2 < \frac{1 - \exp[-\beta(U_{\mathbf{x}''\mathbf{x}'} - U_{\mathbf{x}''\mathbf{x}})]}{q_Z}$ : goto 1
 $\mathbf{X} \leftarrow \{\mathbf{x}'\} \cup \mathbf{X} \setminus \{\mathbf{x}\}$ 
1 output  $\mathbf{X}$ 

```

Algorithm 3: cell-veto. Naive cell-veto algorithm. The far-away cells Z use a two-pebble veto, one for the cell, one for the particle (see patch in Alg. 4).

III. CELL-VETO LOOP

The loop over the far-away cells \mathcal{F} in the naive Alg. 3 (cell-veto) scales linearly with N . It can be avoided by sampling the subset $\mathcal{S}_{\text{veto}}$ of cells vetoing the proposed move (Sec. III A). Two different scenarios are possible. For infinitesimal moves, when the cell-veto algorithm is implemented for a continuous-time Markov process, as




FIG. 6: Naive cell-veto algorithm for a move $\mathbf{x} \rightarrow \mathbf{x}'$, as implemented in Alg. 3 (**cell-veto**). (a): The box is divided into cells which rarely contain more than one particle. (b): Neighbor and surplus particles are handled directly, as in Alg. 2. (c,d, ...): Cell-veto loop. The far-away cells \mathcal{F} are iterated over. Each cell Z , then the particle \mathbf{x}'' inside it, may veto the move.

in the event-chain Monte Carlo algorithm [36], the veto probabilities are themselves infinitesimal and, most of the time, the set $\mathcal{S}_{\text{veto}}$ is empty. Otherwise, if $\mathcal{S}_{\text{veto}} \neq \emptyset$, this set contains a single cell, which can be sampled in $\mathcal{O}(1)$ using Walker’s algorithm [37] (Sec. III B). In contrast, for finite moves, the set $\mathcal{S}_{\text{veto}}$ may contain more than one element, in which case it can be sampled without iterating over all cells [28]. We present an algorithm which reduces the sampling problem to the infinitesimal-move case (Sec. III C).

A. The set $\mathcal{S}_{\text{veto}}$ of veto cells

In Alg. 3 (**cell-veto**), all cells are looped over, but the veto must be confirmed in a second step only for the subset $\mathcal{S}_{\text{veto}} \subset \mathcal{F}$ of vetoing cells. The “cell-accept” decisions are confirmed automatically [24]. Clearly, the probability of any such subset is

$$\pi(\mathcal{S}_{\text{veto}}) = \prod_{\substack{Z \in \mathcal{S}_{\text{veto}} \\ \text{veto}}} q_Z \prod_{\substack{Z \notin \mathcal{S}_{\text{veto}} \\ \text{no veto}}} (1 - q_Z). \quad (13)$$

Given $\mathcal{S}_{\text{veto}}$, the loop over far-away cells becomes superfluous, and the cell-veto algorithm can be patched as in Alg. 4.

B. Infinitesimal moves, event-chain Monte Carlo

In the non-reversible event-chain Monte Carlo algorithm [23, 36, 38], the move

$$\mathbf{x} \rightarrow \mathbf{x}' = \mathbf{x} + \mathbf{v}dt \quad (14)$$

is infinitesimal and it is repeated until vetoed by a factor M (in our case a pair). The particle originally at \mathbf{x} thus moves in continuous time along a ray indicated by the velocity \mathbf{v} in Eq. (14). This velocity can often be taken

procedure **cell-veto(patch)**

```

...
(treating neighbor and surplus particles as in Alg. 3)
...
 $\mathcal{S}_{\text{veto}} \leftarrow$  sampled from  $\pi(\mathcal{S}_{\text{veto}})$  (see Eq. (13))
for  $Z \in \mathcal{S}_{\text{veto}}$ , non-empty: (patch of cell-veto loop)
   $\Upsilon \leftarrow \text{ran}(0, 1)$ 
   $\mathbf{x}'' \leftarrow$  (unique) particle  $\in Z$ 
  if  $\Upsilon < \frac{1 - \exp[-\beta(U_{\mathbf{x}''\mathbf{x}'} - U_{\mathbf{x}''\mathbf{x}})]}{q_Z}$ : goto 1
 $\mathbf{X} \leftarrow \{\mathbf{x}'\} \cup \mathbf{X} \setminus \{\mathbf{x}\}$ 
1 output  $\mathbf{X}$ 

```

Algorithm 4: **cell-veto(patch)**. Patch of Alg. 3. The iteration over the the far-away cells \mathcal{F} is replaced by the iteration over the vetoing cells $\mathcal{S}_{\text{veto}} \subset \mathcal{F}$.

of unit norm $|\mathbf{v}| = 1$ and the inverse move need not be proposed. After the veto, the other element in the pair factor becomes the active particle, and this situation that can be easily generalized to larger factors describing for example many-body interactions or groups of atoms inside a pair of molecules [29, 30, 39, 40].

Within the active cell “0”, a large number of subsequent infinitesimal moves along the ray are not vetoed by any cell, with the probability

$$1 - \underbrace{\sum_{Z \in \mathcal{F}} q_Z}_{q_{\mathcal{F}}}, \quad (15)$$

giving rise to a waiting-time distribution

$$\mathbb{P}(t)dt = q_{\mathcal{F}} \exp(-q_{\mathcal{F}}t) dt. \quad (16)$$

This waiting time for the next cell event can be sampled as $t = \log[\text{ran}(0, 1)]/q_{\mathcal{F}}$. At the event time, the cell responsible for the veto can be sampled in $\mathcal{O}(1)$ using Walker’s algorithm, as is illustrated in Figs 7 and 8 for far-away cells $\mathcal{F} = \{A, \dots, P\}$. In a periodic simulation box, the Walker table concerns relative cell locations, and

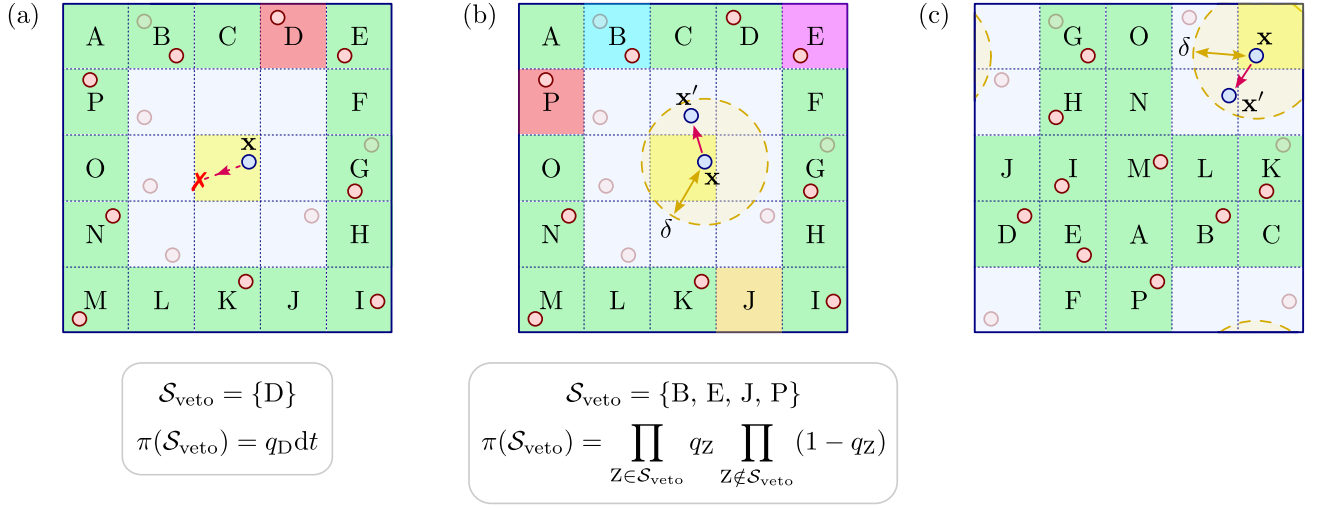


FIG. 7: Subsets of cells. (a): An infinitesimal move, vetoed by at most a single cell. (b): A finite move $\mathbf{x} \rightarrow \mathbf{x}'$, vetoed by a non-trivial subset of cells. (c): With periodic boundary conditions, the set \mathcal{F} is defined relative to the active cell.




FIG. 8: Walker's algorithm for sampling the subset $\mathcal{S}_{\text{veto}} \subset \mathcal{F} = \{A, B, \dots, P\}$ of vetoing cells (see Fig. 7). The probabilities shown correspond to $\{q_A, \dots, q_P\}$ for infinitesimal moves or $\{\lambda_A, \dots, \lambda_P\}$, for finite moves. (a): The probabilities reflect the spatial symmetry of the cells. (b): The probabilities are rearranged into a rectangle, with at most two elements stacked on top of each other. After a $\mathcal{O}(|\mathcal{F}|)$ preparation (from (a) to (b)), each sample takes two random numbers, so is $\mathcal{O}(1)$ for $|\mathcal{F}| \rightarrow \infty$.

it is prepared at the beginning of the simulation and then translates together with the active particle (see Fig. 7c). The method outlined above can be integrated into Alg. 4 for a native constant-time event-chain Monte Carlo algorithm for the Lennard-Jones model and many other models (see Fig. 10 for run-time data, and Appendix A for a Python implementation).

C. Finite displacements

For a finite proposed move $\mathbf{x} \rightarrow \mathbf{x}'$, in discrete time and with finite cell-veto probabilities q_Z for $Z \in \mathcal{F}$, more than one cell can veto, leading to a non-trivial subset $\mathcal{S}_{\text{veto}}$. While the cell-accepts are definite, the vetoes are not. They must be confirmed by a second ‘‘pebble’’. The set $\mathcal{S}_{\text{veto}}$, with the probability distribution given in Eq. (13), can be sampled naively by iterating over the elements of \mathcal{F} .

To sample the set $\mathcal{S}_{\text{veto}}$ more efficiently, we consider a

procedure poisson-veto

input $\{q_A, \dots, q_P\}$ (q_Z : veto probability of cell Z)
 $\mathcal{S}_{\text{veto}} \leftarrow \emptyset$
for $Z \in \{A, \dots, P\}$:
 $\left\{ \begin{array}{l} \lambda_Z \leftarrow -\log(1 - q_Z) \\ t \leftarrow -\log[\text{ran}(0, 1)] / \lambda_Z \\ \text{if } t < 1: \mathcal{S}_{\text{veto}} \leftarrow \mathcal{S}_{\text{veto}} \cup \{Z\} \end{array} \right.$
output $\mathcal{S}_{\text{veto}}$

Algorithm 5: poisson-veto. Sampling the set $\mathcal{S}_{\text{veto}}$ through $|\mathcal{F}|$ Poisson processes. See patch in Alg. 6 for an equivalent, faster, version with a single Poisson process.

Poisson process of intensity λ_Z in the time interval $[0, 1]$, with a distribution of the number of events as

$$\mathbb{P}(n \text{ events}) = \frac{\lambda_Z^n}{n!} e^{-\lambda_Z}. \quad (17)$$

If we identify q_Z with the probability that one or more

events take place in the time interval $[0, 1]$, we have:

$$q_Z = \mathbb{P}(\geq 1 \text{ events}), \quad (18)$$

$$= 1 - \mathbb{P}(0 \text{ events}) = 1 - e^{-\lambda_Z}, \quad (19)$$

which gives $\lambda_Z = -\log(1 - q_Z)$. The waiting-time distribution of this Poisson process is

$$\mathbb{P}(t)dt = \lambda_Z \exp(-\lambda_Z t) dt, \quad (20)$$

and events from this process can be sampled as $t = \log[\text{ran}(0, 1)]/\lambda_Z$. Naively, we might loop over the cells in \mathcal{F} , translate the cell-veto probabilities q_Z into the intensities λ_Z of associated Poisson processes and check whether the waiting times are smaller than 1 (see Alg. 5 (`poisson-veto`)). Integrated into Alg. 4 (`cell-veto(patch)`), this exactly reproduces the factorized Metropolis algorithm.




FIG. 9: Sampling the veto cells $\mathcal{S}_{\text{veto}}$ (see Fig. 7b). (a): Using subsequent Poisson processes for $\{A, \dots, P\}$ (see Alg. 5 (`poisson-veto`)). (b): Using a single Poisson process, together with Walker’s algorithm. Each cell may veto repeatedly (see Alg. 6 (`poisson-veto(patch)`)).

As the Poisson processes in Alg. 5 (`poisson-veto`) are independent, one may run them in parallel (see Fig. 9). A single Poisson process with intensity $\lambda = \sum_{Z \in \mathcal{F}} \lambda_Z$ then controls the time at which one of the $|\mathcal{F}|$ Poisson processes, that is, one of the cells, holds a veto. Walker’s algorithm can again be used to identify this cell (see Alg. 6 (`poisson-veto(patch)`)). This algorithm avoids the iteration over the set \mathcal{F} of far-away cells, and as Walker’s algorithm is of complexity $\mathcal{O}(1)$, the complexity of our algorithm is essentially $\mathcal{O}(|\mathcal{S}_{\text{veto}}|)$. The time interval of the Poisson process may be adjusted so that $\mathcal{S}_{\text{veto}}$ contains only few elements. Algorithm 6 (`poisson-veto(patch)`) can be integrated into Alg. 4 for a native constant-time Monte Carlo algorithm for the Lennard-Jones model (see Fig. 10 for run-time data, and Appendix A for a Python implementation).

procedure `poisson-veto(patch)`

input $\{q_A, \dots, q_P\}$ (q_Z : veto probability of cell Z)
 $\{\lambda_A, \dots, \lambda_P\} \leftarrow \{-\log(1 - q_A), \dots, -\log(1 - q_P)\}$
 $\lambda \leftarrow \sum_{Z \in \mathcal{F}} \lambda_Z$ (total intensity of all Poisson processes)
 $\mathcal{S}_{\text{veto}} \leftarrow \emptyset$
 $t \leftarrow 0$

while True:

$t \leftarrow t - \log[\text{ran}(0, 1)]/\lambda$ (waiting time to next event)
if $t > 1$: **break**
 $Z \leftarrow \text{walker}(\{\lambda_A, \dots, \lambda_P\})$
 $\mathcal{S}_{\text{veto}} \leftarrow \mathcal{S}_{\text{veto}} \cup \{Z\}$

output $\mathcal{S}_{\text{veto}}$

Algorithm 6: `poisson-veto(patch)`. Efficient sampling of the set $\mathcal{S}_{\text{veto}}$ of veto cells through a single Poisson processes.

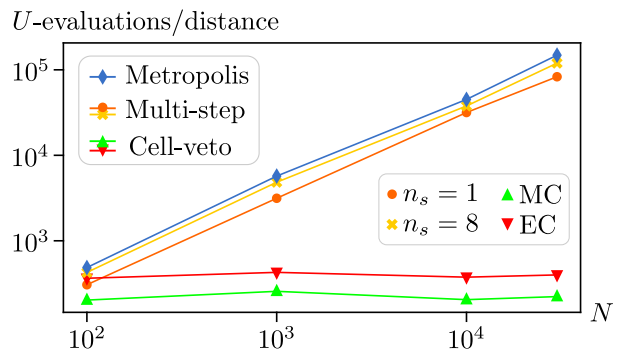


FIG. 10: Number of pair-potential evaluations per unit-distance traveled by a single particle. The results were obtained using the Python programs of Appendix A. The lines are guides to the eye.

IV. CONCLUSIONS

In this paper, we have discussed the role of long-range potentials within molecular simulation and, especially, within the Monte Carlo approach to sampling. In a situation where enormous effort has been dedicated to the efficient evaluation of the total energy and its gradient, we point out that Monte Carlo sampling of the Boltzmann distribution $\exp(-\beta U)$ does not generally require knowledge of the total energy U . The cell-veto algorithms that we presented here use a “two-pebble” decision that is akin to the thinning of Poisson processes, and they can be extended in many directions. In related research [29, 32, 41], we have shown that the SPC/Fw water model frequently used in biomolecular simulations can be sampled exactly in the canonical ensemble, without relying on thermostats, discretization, cutoffs, and grids. It remains to be seen whether our native methods can be applied more generally to molecular applications and whether it can solve the problems posed by truncating, smoothing out, or discretizing interactions in long-time simulations. In our examples, the long-range nature of the potentials relate to the physics of the photon (rather than the meson) for the Coulomb interaction, and the London dispersion force for the Lennard-Jones potential

of great importance for the physics of interfaces and surfaces in soft condensed matter.

ACKNOWLEDGMENTS

This research was supported by a grant from the Simons Foundation (Grant 839534, MET). We thank N. Bou-Rabee, T. Schlick, Y. Sugita, and S. Vionnet for helpful discussions.

Appendix A: Computer programs

The present work is accompanied by the `MCLongRange` software package, which is published as an open-source project under the GNU GPLv3 license. `MCLongRange` is available on GitHub as a part of the `JeLLyFysh` organization. The package contains Python implementations of many algorithms discussed in this paper and has been used to produce the numerical results of Fig. 10. The url of the repository is <https://github.com/jellyfysh/MCLongRange.git>.

-
- [1] T. Schlick, *Molecular Modeling and Simulation: An Interdisciplinary Guide*. Springer, 2002.
- [2] H. Eissfeller and M. Opper, “New method for studying the dynamics of disordered spin systems without finite-size effects,” *Phys. Rev. Lett.*, vol. 68, no. 13, p. 2094–2097, 1992.
- [3] J. Jordan, R. Orús, G. Vidal, F. Verstraete, and J. I. Cirac, “Classical Simulation of Infinite-Size Quantum Lattice Systems in Two Spatial Dimensions,” *Phys. Rev. Lett.*, vol. 101, no. 25, 2008.
- [4] K. Van Houcke, E. Kozik, N. Prokof’ev, and B. Svistunov *Phys. Procedia*, vol. 6, p. 95, 2010.
- [5] K. Van Houcke, F. Werner, E. Kozik, N. Prokof’ev, B. Svistunov, M. J. Ku, A. S. and L.W. Cheuk, A. Schirrotzek, and M. W. Zwierlein, “Feynman diagrams versus Fermi-gas Feynman emulator,” *Nature Phys.*, vol. 8, p. 366, 2012.
- [6] R. Rossi, “Determinant Diagrammatic Monte Carlo in the Thermodynamic Limit,” *Phys. Rev. Lett.*, vol. 119, p. 045701, 2017.
- [7] M. Levitt, “Birth and Future of Multiscale Modeling for Macromolecular Systems (Nobel Lecture),” *Angew. Chem. Int. Ed.*, vol. 53, p. 10006–10018, Aug. 2014.
- [8] J. Cardy, *Scaling and Renormalization in Statistical Physics*. Cambridge University Press, 1996.
- [9] D. Aldous and P. Diaconis, “Shuffling Cards and Stopping Times,” *Am. Math. Mon.*, vol. 93, no. 5, pp. 333–348, 1986.
- [10] J. G. Propp and D. B. Wilson, “Exact sampling with coupled Markov chains and applications to statistical mechanics,” *Random Structures Algorithms*, vol. 9, no. 1-2, pp. 223–252, 1996.
- [11] A. J. Stone, *The Theory of Intermolecular Forces, 2nd edition*. Oxford University Press, 2013.
- [12] D. E. Shaw, M. M. Deneroff, R. O. Dror, J. S. Kuskin, R. H. Larson, J. K. Salmon, C. Young, B. Batson, K. J. Bowers, J. C. Chao, M. P. Eastwood, J. Gagliardo, J. P. Grossman, C. R. Ho, D. J. Ierardi, I. Kolossváry, J. L. Klepeis, T. Layman, C. McLeavey, M. A. Moraes, R. Mueller, E. C. Priest, Y. Shan, J. Spengler, M. Theobald, B. Towles, and S. C. Wang, “Anton, a Special-purpose Machine for Molecular Dynamics Simulation,” in *Proceedings of the 34th Annual International Symposium on Computer Architecture*, ISCA ’07, (New York, NY, USA), pp. 1–12, ACM, 2007.
- [13] D. E. Shaw, P. Maragakis, K. Lindorff-Larsen, S. Piana, R. O. Dror, M. P. Eastwood, J. A. Bank, J. M. Jumper, J. K. Salmon, Y. Shan, and W. Wriggers, “Atomic-Level Characterization of the Structural Dynamics of Proteins,” *Science*, vol. 330, no. 6002, pp. 341–346, 2010.
- [14] L. Greengard and V. Rokhlin, “A fast algorithm for particle simulations,” *J. Comput. Phys.*, vol. 73, no. 2, pp. 325–348, 1987.
- [15] B. Smit and D. Frenkel, “Vapor–liquid equilibria of the two-dimensional Lennard-Jones fluid(s),” *J. Chem. Phys.*, vol. 94, no. 8, pp. 5663–5668, 1991.
- [16] B. Smit, “Phase diagrams of Lennard-Jones fluids,” *J. Chem. Phys.*, vol. 96, no. 11, pp. 8639–8640, 1992.
- [17] Y. A. Lei, T. Bykov, S. Yoo, and X. C. Zeng, “The Tolman Length: Is It Positive or Negative?,” *J. Am. Chem. Soc.*, vol. 127, no. 44, p. 15346–15347, 2005.
- [18] A. Tröster, M. Oettel, B. Block, P. Virnau, and K. Binder, “Numerical approaches to determine the interface tension of curved interfaces from free energy calculations,” *J. Chem. Phys.*, vol. 136, no. 6, 2012.
- [19] C. Tempra, O. H. S. Ollila, and M. Javanainen, “Accurate Simulations of Lipid Monolayers Require a Water Model with Correct Surface Tension,” *J. Chem. Theory Comput.*, vol. 18, no. 3, p. 1862–1869, 2022.
- [20] Y. Yu, A. Krämer, R. M. Venable, B. R. Brooks, J. B. Klauda, and R. W. Pastor, “CHARMM36 Lipid Force Field with Explicit Treatment of Long-Range Dispersion: Parametrization and Validation for Phosphatidylethanolamine, Phosphatidylglycerol, and Ether Lipids,” *J. Chem. Theory Comput.*, vol. 17, no. 3, p. 1581–1595, 2021.
- [21] W. Krauth, *Statistical Mechanics: Algorithms and Computations*. Oxford University Press, 2006.
- [22] S. C. Kapfer and W. Krauth, “Cell-veto Monte Carlo algorithm for long-range systems,” *Phys. Rev. E*, vol. 94, p. 031302, 2016.
- [23] M. Michel, S. C. Kapfer, and W. Krauth, “Generalized event-chain Monte Carlo: Constructing rejection-free global-balance algorithms from infinitesimal steps,” *J. Chem. Phys.*, vol. 140, no. 5, p. 054116, 2014.
- [24] G. Tartero and W. Krauth, “Concepts in Monte Carlo sampling,” *Am. J. Phys.*, vol. 92, no. 1, pp. 65–77, 2024.
- [25] P. A. W. Lewis and G. S. Shedler, “Simulation of non-homogeneous Poisson processes by thinning,” *Nav. Res. Logist. Q.*, vol. 26, no. 3, pp. 403–413, 1979.
- [26] F. Müller, H. Christiansen, S. Schnabel, and W. Janke, “Fast, Hierarchical, and Adaptive Algorithm for Metropolis Monte Carlo Simulations of Long-Range Interacting Systems,” *Phys. Rev. X*, vol. 13, no. 3, 2023.

- [27] B. Hetényi, K. Bernacki, and B. J. Berne, “Multiple “time step” Monte Carlo,” *J. Chem. Phys.*, vol. 117, no. 18, pp. 8203–8207, 2002.
- [28] M. Michel, X. Tan, and Y. Deng, “Clock Monte Carlo methods,” *Phys. Rev. E*, vol. 99, p. 010105, 2019.
- [29] P. Höllmer, A. C. Maggs, and W. Krauth, “Molecular simulation from modern statistics: Continuous-time, continuous-space, exact.” <https://arxiv.org/abs/2305.02979>, 2023.
- [30] J. Harland, M. Michel, T. A. Kampmann, and J. Kierfeld, “Event-chain Monte Carlo algorithms for three- and many-particle interactions,” *EPL*, vol. 117, no. 3, p. 30001, 2017.
- [31] J. E. Lennard-Jones, “Cohesion,” *Proc. Phys. Soc.*, vol. 43, no. 5, p. 461, 1931.
- [32] M. F. Faulkner, L. Qin, A. C. Maggs, and W. Krauth, “All-atom computations with irreversible Markov chains,” *J. Chem. Phys.*, vol. 149, no. 6, p. 064113, 2018.
- [33] S. Duane, A. D. Kennedy, B. J. Pendleton, and D. Roweth, “Hybrid Monte-Carlo,” *Phys Lett B*, vol. 195, pp. 216–222, 1987.
- [34] R. M. Neal, “MCMC using Hamiltonian dynamics,” in *Handbook of Markov Chain Monte Carlo* (S. Brooks, A. Gelman, G. Jones, and X.-L. Meng, eds.), pp. 113–162, Chapman and Hall/CRC, 2011.
- [35] N. Bou-Rabee and J. M. Sanz-Serna, “Geometric integrators and the Hamiltonian Monte Carlo Method,” *Acta Numer.*, vol. 27, pp. 113–206, 2018.
- [36] W. Krauth, “Event-Chain Monte Carlo: Foundations, Applications, and Prospects,” *Front. Phys.*, vol. 9, p. 229, 2021.
- [37] A. J. Walker, “An Efficient Method for Generating Discrete Random Variables with General Distributions,” *ACM Trans. Math. Softw.*, vol. 3, no. 3, pp. 253–256, 1977.
- [38] E. P. Bernard, W. Krauth, and D. B. Wilson, “Event-chain Monte Carlo algorithms for hard-sphere systems,” *Phys. Rev. E*, vol. 80, p. 056704, 2009.
- [39] M. Klement and M. Engel, “Efficient equilibration of hard spheres with Newtonian event chains,” *J. Chem. Phys.*, vol. 150, no. 17, p. 174108, 2019.
- [40] M. Michel, A. Durmus, and S. Sénécal, “Forward Event-Chain Monte Carlo: Fast Sampling by Randomness Control in Irreversible Markov Chains,” *J. Comput. Graph. Stat.*, vol. 29, no. 4, pp. 689–702, 2020.
- [41] P. Höllmer, L. Qin, M. F. Faulkner, A. C. Maggs, and W. Krauth, “JeLLyFysh-Version1.0 — a Python application for all-atom event-chain Monte Carlo,” *Comput. Phys. Commun.*, vol. 253, p. 107168, 2020.