

Velocity Jumps for Molecular Dynamics

Nicolaï Gouraud, Louis Lagardère, Olivier Adjoua, Thomas Plé, Pierre Monmarché,*
and Jean-Philip Piquemal*



Cite This: *J. Chem. Theory Comput.* 2025, 21, 2854–2866



Read Online

ACCESS |



Metrics & More

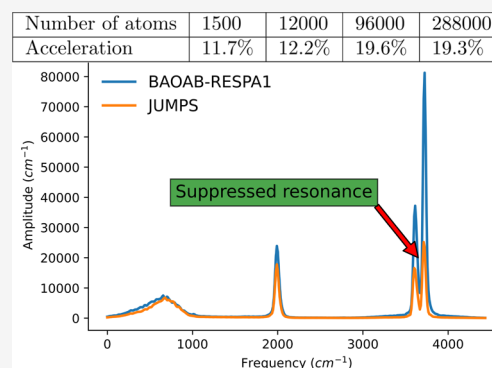


Article Recommendations



Supporting Information

ABSTRACT: We introduce the Velocity Jumps approach, denoted as JUMP, a new class of Molecular dynamics integrators, replacing the Langevin dynamics by a hybrid model combining a classical Langevin diffusion and a piecewise deterministic Markov process, where the expensive computation of long-range pairwise interactions is replaced by a resampling of the velocities at random times. This framework allows for an acceleration in the simulation speed while preserving sampling and dynamical properties such as the diffusion constant. It can also be integrated in classical multi-time-step methods, pushing further the computational speedup, while avoiding some of the resonance issues of the latter thanks to the random nature of jumps. The JUMP, JUMP-RESPA and JUMP-RESPA1 integrators have been implemented in the GPU-accelerated version of the Tinker-HP package and are shown to provide significantly enhanced performances compared to their BAOAB, BAOAB-RESPA and BAOAB-RESPA1 counterparts, respectively.



1. INTRODUCTION

Molecular dynamics (MD)^{1,2} is a popular numerical tool to infer macroscopic properties of matter from simulations at a microscopic level, with applications ranging from material sciences to biology.^{3–5} Because of the discrepancy between the microscopic and macroscopic time and space scales, in order to obtain relevant results, one needs to simulate complex molecular systems with a high precision and for a very long time.^{6–11} In recent years, this has been allowed by the development of empirical force fields (such as CHARMM,¹² AMBER,¹³ OPLS,¹⁴ GROMOS,¹⁵ AMOEBA,¹⁶ SIBFA¹⁷ and others^{18–20}), as well as the rise of high-performance computing with GPUs,^{21–24} allowing to run massively parallel codes. Nowadays, many molecular simulation softwares are available (LAMMPS,²⁵ NAMD,²⁶ GROMACS,²⁷ AMBER,²⁸ Tinker-HP²⁹ and others^{9,30–32}) and a lot of attention is drawn to improving the efficiency of the simulation algorithms.

In the framework of Langevin dynamics, multi-time-step methods^{33–36} are now a standard way to make faster simulations than classical integrators such as BAOAB,^{37–40} but they are limited by resonance effects that bound the maximum usable time step at a given accuracy.^{41,42} Various alternative strategies have been proposed such as the Generalized Langevin Equation (GLE)⁴³ or the stochastic isokinetic extended phase-space algorithm.^{44–46} They yield significant acceleration but rely either on some empirical fitting⁴³ or have an important negative impact on the dynamical properties such as the diffusion coefficient,⁴⁴ which is an indication of a limitation in the effective sampling rate. In the last case, the computational gain is reduced by the

fact that a longer trajectory is necessary to keep the same amount of sampling. To speed up molecular dynamics without resorting to some fitting and while not affecting too much the dynamics,^{43,47} one possibility is to combine well-chosen integrators within a multisplit approach like for BAOAB-RESPA1 to push forward the stability limit.³⁴

Here we will focus on an alternative approach that involves replacing the Langevin dynamics by a hybrid model combining a classical Langevin diffusion and a piecewise deterministic Markov process, in which some long-range, computationally expensive force are treated in an adaptive way at random times. The initial idea of this strategy was first introduced in⁴⁸ in a particular case, and further studied theoretically in.⁴⁹ This hybrid model still samples exactly from μ but can be simulated using a numerical splitting scheme that requires fewer gradient computations per time step, with a precision of the same order (in the time step) as the classical splitting schemes of the Langevin diffusion such as BAOAB. Moreover, it can be tuned to be arbitrarily close to the Langevin dynamics in terms of stochastic trajectories (see [50, Theorem 3.6]), which makes it suitable to estimate the dynamical properties of the process (with, of course, a trade-off between the accuracy of these dynamical properties and the numerical cost of the simulation,

Received: January 7, 2025

Revised: February 26, 2025

Accepted: February 28, 2025

Published: March 6, 2025



as would be with any numerical approximation of the Langevin equation). Finally, this versatile framework is parallelizable (allowing GPU implementations) and can be combined with the multi-time-step methods, pushing further the computational speedup while avoiding some of the resonance issues of the latest.^{41,42} This new integrator has been implemented in the Tinker-HP software,²⁹ both on CPU and GPU versions,²¹ the latter allowing simulations on much larger systems.

The paper is organized as follows. Section 2 is devoted to the description of the numerical method: we define in Sections 2.1 and 2.2 the continuous-time hybrid process, Sections 2.3 and 2.4 describe the simulation techniques and how they yield a computational advantage. Then, in Section 3, we see how our approach can be applied to Coulomb and van der Waals interactions in a classical force field. Finally, Section 4 is devoted to numerical experiments, both on CPU and GPU versions of the Tinker-HP software.²⁹ The Supporting Information gives the proof of some of the results and other technical details.

1.1. Notations. In all the following, $\|\cdot\|$ denotes the Euclidean norm in \mathbb{R}^d (or the Frobenius norm when it is applied to matrices), and \cdot the standard dot product. For all $x \in \mathbb{R}$, $(x)_+ = \max(0, x)$ denotes the positive part of x .

2. NUMERICAL METHOD

2.1. Velocity Jump Langevin Process. Let us consider a system of N interacting atoms. We denote $(x, v) \in \mathbb{R}^{6N}$ their positions and velocities, $M = \text{diag}(m_1 I_3, \dots, m_N I_3)$ their mass matrix, $\beta = 1/(k_B T)$ the inverse temperature of the system (where k_B is the Boltzmann constant), $U: \mathbb{R}^{3N} \rightarrow \mathbb{R}$ the potential energy function encoding the interactions between the particles and finally $H(x, v) = U(x) + \frac{1}{2} v^T M v$ the Hamiltonian of the system, corresponding to its total energy.

In the canonical ensemble, the system is described by a probability measure that gives, for any given state $(x, v) \in \mathbb{R}^{6N}$, the probability that the system is in the configuration (x, v) . This probability measure is called the Boltzmann-Gibbs measure, defined by

$$d\mu(x, v) = \frac{1}{\mathcal{Z}_\mu} \exp(-\beta H(x, v)) dx dv \quad (1)$$

where $\mathcal{Z}_\mu = \int_{\mathbb{R}^{6N}} \exp(-\beta H(x, v)) dx dv$. Macroscopic quantities are then described as expectations of an observable with respect to this Gibbs measure. In many cases, computing them analytically is impossible, since they are high-dimensional integrals involving an unknown constant \mathcal{Z}_μ . However, it is possible to approximate $\mathbb{E}_\mu[\varphi]$ by simulating a long trajectory of a Markov process $(X_s, V_s)_{s \geq 0}$ that is ergodic with respect to μ , which means that for any φ (in some suitable class of functions), almost surely

$$\frac{1}{t} \int_0^t \varphi(X_s, V_s) ds \xrightarrow[t \rightarrow \infty]{} \int_{\mathbb{R}^{6N}} \varphi(x, v) d\mu(x, v) = \mathbb{E}_\mu[\varphi]$$

A popular process that has this property (under mild conditions on the potential U) is the Langevin dynamics, defined as the solution of the following SDE

$$\begin{cases} dX_t = V_t dt \\ dV_t = -M^{-1} \nabla U(X_t) dt - \gamma V_t dt - M^{-1/2} \sqrt{2\gamma\beta^{-1}} dW_t \end{cases} \quad (2)$$

where $(W_t)_{t \geq 0}$ is a standard Brownian motion in \mathbb{R}^{3N} and $\gamma > 0$ is a friction parameter.

In practice, in many molecular simulations, periodic boundary conditions are enforced (typically to ensure bounded solutions in systems that do not have a confining potential), which means that the position lies in the periodic flat torus $\mathbb{T}^{3N} = \mathbb{R}^{3N}/\mathbb{Z}^{3N}$. The measure (1) and the SDE (2) are then understood in $\mathbb{T}^{3N} \times \mathbb{R}^{3N}$.

The following general procedure was first described in.⁴⁸ The first step is to decompose the forces as a sum of $K \geq 1$ vector fields F_i

$$\nabla U(x) = \sum_{i=0}^K F_i(x)$$

such that, typically, F_0 gathers the computationally inexpensive components (e.g., short-range interactions exhibiting fast variation), while the F_i terms, for $i \geq 1$, represent longer-range forces, which are more numerically intensive than F_0 (as each atom interacts with all others through these forces, unlike the short-range ones).

We now introduce the velocity jump Langevin process $(X_t, V_t)_{t \geq 0}$ as the following. The dynamics follows the Langevin diffusion associated with the force F_0

$$\begin{cases} dX_t = V_t dt \\ dV_t = -M^{-1} F_0(X_t) dt - \gamma V_t dt - M^{-1/2} \sqrt{2\gamma\beta^{-1}} dW_t \end{cases} \quad (3)$$

with the velocity V_t undergoing additional random jumps at rate $\lambda_i(x, v)$ following a jump kernel $q_i(x, v, dv')$ (that describes the probability distribution of the velocity after a jump), both defined below in Section 2.2, as also explained in⁴⁸ in a particular case. The jump mechanism (λ_i, q_i) depends on the force F_i in a way that ensures that the equilibrium measure of the process is indeed the canonical measure (1). Informally, the process follows (3) but, between times t and $t + \delta$ for a small δ , with probability $\lambda_i(X_t, V_t)\delta + o(\delta)$, its velocity v is resampled according to the kernel $q_i(X_t, V_t, dv')$. It can be seen as a hybrid between a diffusion and a piecewise deterministic Markov process (PDMP) such as considered in,^{51–53} which have recently drawn much attention in numerical probabilities and Bayesian statistics, see.^{52,54–58} As we will explain, simulating this process in a certain way yields a computational advantage over the simulation of the Langevin diffusion, because the computation of the force F_i will only be required when a jump of type i is proposed, which will not happen at each time step for every i . We refer to the Supporting Information Section S1, for a more rigorous definition.

2.2. Jump Mechanism. The velocity jump Langevin process defined in Section 2.1 was described in⁴⁸ in the particular case where the jump mechanism corresponds to the Bouncy Particle Sampler (BPS),⁵⁹ i.e., with

$$\lambda_i(x, v) = \beta(v \cdot F_i(x))_+ \quad (4)$$

and deterministic jumps $q_i(x, v, dv') = \delta_{R_i(x,v)}(dv')$, where

$$R_i(x, v) = v - 2 \frac{v \cdot F_i(x)}{F_i(x) \cdot M^{-1} F_i(x)} M^{-1} F_i(x) \quad (5)$$

Although this jump mechanism gives good results in a certain framework,⁴⁸ it also suffers some limitations, that we will explain below, in Section 3. We now define a more general jump mechanism based on,⁵⁰ that can be seen as an interpolation between the BPS and Hamiltonian dynamics. Let $\varepsilon(x) > 0$ be a positive function of the position. For each $1 \leq i \leq K$, We define the following jump rate

$$\lambda_i(x, v) = \frac{\sqrt{\beta}}{\varepsilon(x)} |M^{-1/2} F_i(x)| \Theta(\eta) \quad (6)$$

where

$$\eta = \varepsilon(x) \sqrt{\beta} \frac{v \cdot F_i(x)}{|M^{-1/2} F_i(x)|}$$

and $\Theta(\eta) = \mathbb{E}[(\eta + G)_+]$ with $G \sim \mathcal{N}(0, 1)$ a standard Gaussian variable. The transition kernel $q_i(x, v, dv')$ is such that at the moment of a jump, the velocity is updated as

$$v \leftarrow v - \frac{2\varepsilon(x)}{\sqrt{\beta}(1 + \varepsilon^2(x))} (\eta + \tilde{G}) \frac{M^{-1} F_i(x)}{|M^{-1/2} F_i(x)|} \quad (7)$$

where \tilde{G} is a one-dimensional random variable with density

$$f_\eta(y) = \frac{1}{\Theta(\eta) \sqrt{2\pi}} (\eta + y)_+ e^{-y^2/2}$$

In practice, depending on the context, the parameter ε will be chosen to be either a positive constant or proportional to the norm of the force: $\varepsilon(x) = \varepsilon_0 |F_i(x)|$ with ε_0 a positive constant. It is shown in⁵⁰ that in particular, $\varepsilon = \infty$ corresponds to the BPS (i.e., the jump process described in eqs (4) and (5)), and that when $\varepsilon \rightarrow 0$, the corresponding velocity jump process (with a free transport on positions) converges to the Hamiltonian dynamics associated with the force F_v , i.e., the process solving

$$\begin{cases} dX_t = V_t dt \\ dV_t = -M^{-1} F_i(X_t) dt \end{cases}$$

Moreover, as explained in Supporting Information Section S4, the definitions (6) and (7) ensure that the Gibbs measure μ is left invariant by the velocity jump Langevin process defined in Section 2.1, and the ergodicity of the process in a theoretical setting is proven in.⁴⁹ In other words, when ε is small enough, the velocity jump Langevin process is close, in terms of stochastic trajectories, to the Langevin dynamics, which yields (as we will see in Section 4), a good preservation of dynamical properties of the system on top of a computational speedup without loss of accuracy in sampling.

2.3. Splitting Schemes and Integration to Multi-Time-Step Methods. **2.3.1. JUMP Integrator.** In practice, since the solution of the SDE (3) cannot be simulated exactly, time discretization methods need to be used to approximate the process. Similarly to the BAOAB^{38–40} scheme, the continuous-time dynamics of the velocity jump Langevin process can be approximated following a Trotter/Strang splitting scheme. We will describe three of them, and refer to the Section S2, for more technical details.

The first splitting consists in decomposing the dynamics into a free transport (A), an acceleration due to F_0 (B), a jump (J)

and a friction-dissipation part (O). For a given time-step δ , one step of the numerical scheme is given by the succession of steps BJAOAJB, where

- $v \leftarrow v - \frac{\delta}{2} M^{-1} F_0(x)$ (force F_0)
- $v \leftarrow W_{\delta/2}$ (jumps)
- $x \leftarrow x + \frac{\delta}{2} v$ (free transport)
- $v \leftarrow e^{-\gamma\delta} v + \sqrt{\beta^{-1}(1 - e^{-2\gamma\delta})} M^{-1/2} G$ with $G \sim \mathcal{N}(0, I_d)$ (friction/dissipation),

and where $(W_s)_{s \geq 0}$ is the piecewise-constant process initialized at $W_0 = v$, that jumps at rate λ_i according to q_i for each $i \in [[1, K]]$. Such a scheme is similar to BAOAB, but where two half-time jump steps are added between the forces and the transport part. As explained and proved in,⁴⁹ this palindromic form gives rise to a second-order scheme in the time-step, that is, the discretization bias on the invariant measure is of order δ^2 . This was the procedure used in⁴⁸ for the particular jump process defined in eqs 4 and (5). For numerical stability reasons that we will detail in Section 3, it will be sometimes useful to merge the (B) and (J) parts into one (C) step, and run a splitting scheme of the form CAOAC, where the (C) corresponds to

$$(C) v \leftarrow \tilde{W}_{\delta/2}$$

with $(\tilde{W}_s)_{s \geq 0}$ is defined as $(W_s)_{s \geq 0}$, but such as between jumps, it undergoes the constant acceleration $-M^{-1} F_0(x)$ instead of being constant.

In fact, the BJAOAJB and CAOAC schemes can be combined: we can for instance treat some of the forces F_i in (C), and the others separately in (J), that is, doing a scheme of the form C'J'AOAJ'C', where (J') treats, through a piecewise-constant velocity jump process, the forces F_i for $i \in [[1, K']]$ (with $K' < K$) and (C') treats the F_i for $i \in [[K' + 1, K]]$ and F_0 in the same way as the step (C) described above. This situation will also be encountered in the following applications, see Section 3. We will now refer to those splitting schemes (BJAOAJB, CAOAC or C'J'AOAJ'C') as the JUMP integrator.

2.3.2. JUMP-RESPA and JUMP-RESPA1 Integrators. All these steps can be integrated in a classical multi-time-step framework. Let us say that F_0 is itself a sum of two distinct terms $F_0 = F_{0,0} + F_{0,1}$, where, for instance, $F_{0,1}$ corresponds to a slow-varying many-body force, that cannot be treated efficiently with velocity jumps, but that does not need to be computed as often as the fast-varying forces gathered in $F_{0,0}$. In that case, let Q be the transition kernel corresponding to a splitting scheme (like BJAOAJB, CAOAC or C'J'AOAJ'C') associated with $\nabla U - F_{0,1}$ with a certain time-step δ . Let $\Delta = n\delta$ (with $n \in 2\mathbb{N}$) be a larger time step. The principle of the JUMP-RESPA scheme is to use the following splitting

1. $v \leftarrow v - \frac{\Delta}{2} M^{-1} F_{0,1}(x)$.
2. Perform n times the step Q with time step δ .
3. $v \leftarrow v - \frac{\Delta}{2} M^{-1} F_{0,1}(x)$.

This procedure can itself be iterated to include more different time steps, such as BAOAB-RESPA1³⁴ that has three layers of time steps. By analogy, We will refer to this situation as the JUMP-RESPA1 integrator. More details are given in Section S2.

Let us now see how to simulate efficiently the processes $(W_s)_{s \geq 0}$ and $(\tilde{W}_s)_{s \geq 0}$, and how this gives rise to a computational advantage.

2.4. Efficient Simulation of Jumps. Let us start with the process $(\tilde{W}_i)_{i \geq 0}$ (corresponding to the (C) step in the CAOAC scheme). We suppose that each jump rate λ_i is bounded by a constant $\bar{\lambda}_i$, and denote $\bar{\lambda} = \sum_i \bar{\lambda}_i$. Then an easy computation on the generator (see Section S3) shows that the jumps can be simulated exactly this way: starting from (x, v) ,

1. Draw \mathcal{E} a standard exponential random variable, and let $T = \mathcal{E}/\bar{\lambda}$ be the next jump time proposal.
2. Draw I such that $\mathbb{P}(I = i) = \bar{\lambda}_i/\bar{\lambda}$. Propose a jump of type I at time T .
3. The jump is accepted with probability $\lambda_I(x, v - M^{-1}F_0(x)T)/\bar{\lambda}_I$, in which case the velocity is resampled at time T according to $q_I(x, v - M^{-1}F_0(x)T, dv')$, otherwise the velocity at time T is simply $v - M^{-1}F_0(x)T$ (there is no jump).

The construction is then repeated by induction over the jump time proposals. The case of $(W_s)_{s \geq 0}$ is very similar: the only difference is that since the process is constant between jumps, in the step 3, a jump of type I is accepted with probability $\lambda_I(x, v)/\bar{\lambda}_I$, in which case the velocity is resampled according to $q_I(x, v, dv')$. This method, called Poisson thinning^{60,61} (which is exact, in the sense that it does not induce any further time discretization error) is the key to the computational advantage of the algorithm: contrary to a classical numerical scheme of the Langevin diffusion (such as BAOAB) where the gradient ∇U is computed at each time step, here we only need to evaluate F_i when a jump of type i is proposed (at the step 3 above), which does not happen at each time step for each i . However, notice that similarly to BAOAB and contrary to multi-time-step splitting methods, there is still a unique time step in the discretization and no additional parameters to tune: the frequency at which F_i is evaluated is random and is adapted to each force.

In many cases, we will use the same bound for all the forces F_i , so the $\bar{\lambda}_i$ terms will be uniform in i : $\bar{\lambda}_i = \lambda^*$ (so that $\bar{\lambda} = K\lambda^*$). In those cases, the step 2 above is simply a uniform draw among all the i . The Algorithm 1 describes the (C) step of the CAOAC scheme, and the Algorithm 2 describes the (J) step of the BJAOAB scheme.

In fact, as shown in Section S6, in the general case, the bounds $\bar{\lambda}_i$ are local (they depend the velocity v) and need to be re-evaluated after each jump. This is not a problem, since the only requirement in order to have a computational advantage is that those bounds do not depend on F_i . However, in the case of the Bouncy Particle Sampler (i.e., the jump mechanism defined in (4) and (5)), the norm of the velocity does not change after a jump, so $\bar{\lambda}$ stays constant throughout the time step. In those cases, when we simulate the process up to a time δ , we do not need to know exactly the jump times, but only the number of jumps and the order in which they occur. Classical properties of the exponential distribution ensure that during the time interval $[0, \delta]$, the number of jump proposals (step 1 above) follows a Poisson distribution of parameter $\bar{\lambda}\delta$. Then the simulation procedure for the (J) part of BJAOAB is the following: starting from (x, v) ,

1. Draw $M \sim \text{Poisson}(\bar{\lambda}\delta)$ (the total number of jump proposals).
2. For each $j \in \llbracket 1, M \rrbracket$

- (a) Draw I such that $\mathbb{P}(I = i) = \bar{\lambda}_i/\bar{\lambda}$. The j -th proposal is a jump of type I .
- (b) Accept the proposal (i.e., resample the velocity according to $q_I(x, v, dv')$) with probability $\lambda_I(x, v)/\bar{\lambda}_I$.

This situation is described in Algorithm 3.

Algorithm 1 Simulation of (C) part.

Simulates the jump process $(\tilde{W}_s)_{s \geq 0}$ for a time $\delta/2$, starting from (x, v) .

```

1: procedure VELOCITY-UPDATE( $\delta, x, v$ )
2:    $T \leftarrow 0, t \leftarrow 0$                                 ▷  $t$  is the current time,  $T$  is the time of jump proposals
3:    $a \leftarrow -M^{-1}F_0(x)$                                 ▷  $a$  is the acceleration due to the short range forces
4:   while  $t \leq \delta/2$  do
5:      $T \leftarrow \mathcal{E}(K\lambda^*)$                                 ▷ Draw the next jump proposal time: an exponential law of parameter  $\bar{\lambda}$ 
6:     if  $T + t > \delta/2$  then
7:        $v \leftarrow v + (\delta/2 - t)a$ 
8:     else
9:        $t \leftarrow t + T$                                 ▷ Update the current time
10:       $v \leftarrow v + Ta$                                 ▷ Update the velocity until the jump proposal
11:       $I \leftarrow \text{RANDOM}(\llbracket 1, K \rrbracket)$                     ▷ Choose a jump type
12:       $U \leftarrow \text{RANDOM}([0, 1])$ 
13:       $\lambda \leftarrow \lambda_I(x, v)$                                 ▷ compute the  $I$ -th jump rate at the current state
14:      if  $U \leq \lambda/\lambda^*$  then
15:         $v \leftarrow \text{SAMPLING}(q_I(x, v, dv'))$             ▷ the velocity is resampled
16:      end if
17:    end if
18:  end while
19: end procedure

```

Algorithm 2 Simulation of (J) part (general case).

Simulates the piecewise-constant jump process $(W_s)_{s \geq 0}$ for a time $\delta/2$, starting from (x, v) .

```

1: procedure JUMP( $\delta, x, v$ )
2:    $T \leftarrow 0, t \leftarrow 0$                                 ▷  $t$  is the current time,  $T$  is the time of jump proposals
3:   while  $t \leq \delta/2$  do
4:      $T \leftarrow \mathcal{E}(K\lambda^*)$                                 ▷ Draw the next jump proposal time: an exponential law of parameter  $\bar{\lambda}$ 
5:     if  $T + t > \delta/2$  then
6:        $t \leftarrow \delta/2$ 
7:     else
8:        $t \leftarrow t + T$                                 ▷ Update the current time
9:        $I \leftarrow \text{RANDOM}(\llbracket 1, K \rrbracket)$                     ▷ Choose a jump type
10:       $U \leftarrow \text{RANDOM}([0, 1])$ 
11:       $\lambda \leftarrow \lambda_I(x, v)$                                 ▷ compute the  $I$ -th jump rate at the current state
12:      if  $U \leq \lambda/\lambda^*$  then
13:         $v \leftarrow \text{SAMPLING}(q_I(x, v, dv'))$             ▷ the velocity is resampled
14:      end if
15:    end if
16:  end while
17: end procedure

```

Algorithm 3 Simulation of (J) part (constant bound case)

Special case: a jump does not affect the bound on the jump rate.

Simulates the piecewise-constant jump process $(W_s)_{s \geq 0}$ for a time $\delta/2$, starting from (x, v) .

```

1: procedure JUMP( $\delta, x, v$ )
2:    $M \leftarrow \text{POISSON}(K\lambda^*\delta/2)$ 
3:   for  $j \in \llbracket 1, M \rrbracket$  do
4:      $I \leftarrow \text{RANDOM}(\llbracket 1, K \rrbracket)$                                 ▷ For each jump proposal  $j$ , choose randomly a jump type  $I$ 
5:      $U \leftarrow \text{RANDOM}([0, 1])$ 
6:      $\lambda \leftarrow \lambda_I(x, v)$                                 ▷ compute the  $I$ -th jump rate at state  $(x, v)$ 
7:     if  $U \leq \lambda/\lambda^*$  then
8:        $v \leftarrow \text{SAMPLING}(q_I(x, v, dv'))$             ▷ the velocity is resampled
9:     end if
10:   end for
11:   return
12: end procedure

```

Regarding the simulation of the jump itself, it is shown in⁵⁰ that the random variable \tilde{G} that appears in (7) can be easily simulated with rejection sampling. Assuming that we can sample a distribution with density g that satisfies $f(x) \leq Cg(x)$ and such that the ratio $\frac{f(x)}{Cg(x)}$ is easy to compute for any x , the procedure is the following

1. Draw a proposal Y with density g .
2. Draw U uniform in $[0, 1]$. if $U \leq \frac{f(Y)}{Cg(Y)}$ then the proposal is accepted.
3. Otherwise, repeat from step 1.

The mean number of proposals before acceptance is C , so the smaller the constant the faster the procedure is. Hence, following⁵⁰ (to which we refer for details), in our case, g is a:

Gamma law	for $\eta < -2.5$
exponential law	for $-2.5 \leq \eta < -1$
Rayleigh law	for $-1 \leq \eta \leq 0$
mixed Rayleigh – Gaussian law	for $\eta > 0$

3. APPLICATION TO ELECTROSTATIC AND VAN DER WAALS INTERACTIONS

In a classical (nonpolarizable) force field such as OPLS,¹⁴ AMBER⁶² or CHARMM,¹² the nonbonded parts of the interaction potential are the Lennard–Jones and electrostatic contributions. Since the Lennard–Jones term decreases very fast, in periodic boundary conditions one can simply compute the forces in a radius up to a certain cutoff. The electrostatic part, decreasing at a much slower rate, is more problematic. A classical way of treating the long-range interactions is by using Ewald summation techniques, for instance through the SPME method⁶³ (although other approaches exist, see for instance⁶⁴). In this method, the Coulomb potential is decomposed in a direct and a reciprocal part, the latter being computed in the Fourier space with a fast Fourier transform. In practice, the algorithm BOUNCE described in,⁴⁸ i.e., a BJAQJB splitting on the process with the BPS jump mechanism, has been implemented in Tinker-HP²⁹ and tested in periodic boundary conditions with jumps on the long-range parts of the van der Waals interactions and the direct part of the electrostatic interactions. The reciprocal part was treated with a classical multi-time-step method, the reference system propagator algorithm (RESPA).³⁵ Since electrostatics are much stronger than van der Waals at a long-range, the BPS electrostatic jumps can only be done on a very small range. Jumping on a larger range yields too many velocity bounces, in the sense that the numerical scheme with the classical time step of 1 fs is very unstable. One solution would be to decrease the time step, which is of course counterproductive, since the goal is to reduce the final cost of the calculations.

On the contrary, the jump mechanism defined in (6) and (7) yields milder jumps than the orthogonal reflections of the BPS. Moreover, as discussed earlier, the process can be seen as an interpolation between bounces and diffusions: in the situations where BOUNCE is numerically unstable, choosing a small parameter ε makes the process closer to classical Hamiltonian dynamics. In addition to that, combining the (B) and (J) parts (i.e., doing the CAOAC scheme instead of BJAQJB) stabilizes further the discretized process.

These two elements combined, namely a more general jump process and the fusion of the “jump” and “force” steps, allow to replace a larger range of the potential by jumps than in the BOUNCE algorithm, yielding more computational gain, while preserving dynamical properties of the system such as the diffusion coefficients. Finally, the BPS jumps of⁴⁸ on the van der Waals forces can also be integrated in the general procedure, and two multi-time-step versions of the algorithm, namely JUMP-RESPA and JUMP-RESPA1 combine all type of jumps and yields a computational gain over the BAOAB-RESPA and BAOAB-RESPA1³⁴ algorithm.

3.1. Direct Electrostatic Jumps. As reminded in the Section S5, the Ewald summation splits the electrostatic interactions into a direct and a reciprocal part. In Periodic Boundary Conditions (PBC), this decomposition can be chosen in a way that enforces the “minimum image

convention”: in the direct space, each atom interacts only with the closest image of an other atom. We assume in the following that this condition is satisfied. By denoting $r_{ij} = \|x_i - x_j\|$ the distance between the atoms i and j , the direct electrostatic energy associated with each pair (i, j) is

$$U_{ij}(x) = \frac{q_i q_j}{r_{ij}} \operatorname{erfc}(\alpha r_{ij})$$

In order to decompose this term into a short and a long-range contribution, let r_c be a cutoff (smaller than the “Ewald cutoff”) that separates the direct and reciprocal part, whose value is enforced by the choice of the parameter α , $\chi(r)$ a switching function that goes smoothly from 1 to 0 around the cutoff (its precise formula is given in Section S6) and let

$$U_{ij} = \chi(r_{ij})U_{ij} + (1 - \chi(r_{ij}))U_{ij}$$

We now denote $F_{ij} = \nabla_{x_i}(1 - \chi(r_{ij}))U_{ij} \in \mathbb{R}^3$, the component of the long-range force of the atom j acting on the atom i . Therefore, the direct part of electrostatic interactions is decomposed into

$$\nabla U_{\text{direct}} = F_0 + \sum_{i \neq j} A_i F_{ij}$$

where $A_i \in \mathcal{M}_{3N \times 3}$ is the matrix with all coefficients are equal to zero except $A(3i - 2, 1) = A(3i - 1, 2) = A(3i, 3) = 1$. Therefore, a jump mechanism is associated with each pair interaction $A_i F_{ij}$. In this context, (6) and (7) have simpler expressions than in the general case described in Section 2.

Here, $\left| M^{-1/2} F_{ij}(x) \right| = \frac{|F_{ij}(x)|}{\sqrt{m_i}}$ (with m_i the mass of the atom i), which yields the following jump rate:

$$\lambda_{ij}(x, v) = \frac{\sqrt{\beta}}{\varepsilon(x) \sqrt{m_i}} |F_{ij}(x)| \Theta \left(\varepsilon(x) \sqrt{\beta m_i} \frac{v_i \cdot F_{ij}(x)}{|F_{ij}(x)|} \right)$$

and a jump only modifies the atom i

$$v_i \leftarrow v_i - \frac{2\varepsilon(x)}{1 + \varepsilon^2(x)} \left(\varepsilon(x) \frac{v_i \cdot F_{ij}(x)}{|F_{ij}(x)|} + \frac{\tilde{G}}{\sqrt{\beta m_i}} \right) \frac{F_{ij}(x)}{|F_{ij}(x)|}$$

where \tilde{G} has been generated from the distribution f_η with $\eta = \varepsilon(x) \sqrt{\beta m_i} \frac{v_i \cdot F_{ij}(x)}{|F_{ij}(x)|}$. As detailed in the Section S6, in this case it is easy to obtain explicit analytic bounds on $\|F_{ij}\|$ that are required in the thinning method. Moreover, for each i , those bounds are uniform in j which, as explained in Section 2.4, simplifies the thinning procedure.

Since the jump mechanisms associated with F_{ij} only involve the i -th atom, the jump processes associated with the different atoms can be seen as independent Markov chains that can be simulated in parallel, which in particular allows massively parallel implementations on GPU, see Section 4.

Finally, as mentioned in the introduction of this section, in this case it is more convenient to perform a CAOAC scheme instead of BJAQJB. If we treat a large portion of the interactions with jumps, the rate λ_{ij} will be higher, yielding many velocity jumps, so merging the (J) and (B) steps of the splitting scheme has the effect of stabilizing the dynamics. The Algorithm 4 below describes precisely the procedure when ε is constant. The case of an adaptive $\varepsilon(x) = \varepsilon_0 |F_i(x)|$ is exactly the

same but with a slightly different expression for the jump rate, the bound and the kernel.

Algorithm 4 Step (C): a general velocity update (with constant ε)

Simulates during a time $\delta/2$ the PDMP treating the acceleration due to the short range forces and the jumps due to the long range forces. The position x is fixed.

Input: N the number of atoms, (x, v) their positions and velocities, δ the time step

```

1: procedure VELOCITYUPDATE( $N, \delta, x, v$ )
2:   for  $i \in [1, N]$  do
3:      $T \leftarrow 0, t \leftarrow 0$  ▷  $t$  is the current time,  $T$  is the time of jump proposals
4:      $v_i \leftarrow v_i^0$  ▷  $v_i^0$  is the initial velocity of atom  $i$ 
5:      $a_i \leftarrow -\nabla U_{\text{short}}(x)/m_i$  ▷  $a$  is the acceleration due to the short range forces
6:      $L_i \leftarrow L_i$  ▷ The uniform bound on the  $|F_{ij}|, 1 \leq j \leq N$ 
7:     while  $t \leq \delta/2$  do
8:        $B_i \leftarrow (|v_i|^2 + (\delta/2 - t)^2 |a_i|^2 + 2(\delta/2 - t)(v_i \cdot a_i)_+)^{1/2}$  ▷ bound on  $|v_i|$  during the time  $[t, \delta/2]$ 
9:        $\bar{\lambda} \leftarrow L_i \times N \left( \beta B_i + \frac{1}{\sqrt{2\pi m_i}} \right)$  ▷ The bound on the jump rate
10:       $T \leftarrow \mathcal{E}(\bar{\lambda})$  ▷ Draw an exponential law of parameter  $\bar{\lambda}$ 
11:      if  $T + t > \delta/2$  then
12:         $v_i \leftarrow v_i + (\delta/2 - t)a_i$ 
13:      else
14:         $t \leftarrow t + T$ 
15:         $v_i \leftarrow v_i + T a_i$ 
16:         $\eta \leftarrow \varepsilon \sqrt{\beta m_i} \frac{F_{ij}(x)}{|F_{ij}(x)|}$ 
17:         $J \leftarrow \text{RANDOM}([1, N])$ 
18:         $U \leftarrow \text{RANDOM}([0, 1])$ 
19:         $p \leftarrow \lambda_{ij}(x, v) / \bar{\lambda} = \frac{|F_{ij}(x)| \Theta(\eta)}{L_i (B_i \varepsilon \sqrt{m_i} + 1 / \sqrt{2\pi})}$  ▷ the probability of acceptance
20:        if  $U \leq p$  then
21:           $\tilde{G} \leftarrow \text{REJECTION SAMPLING}(\eta)$  ▷ Generate  $\tilde{G}$ 
22:           $v_i \leftarrow v_i - \frac{2\varepsilon}{\sqrt{\beta m_i}(1 + \varepsilon^2)} \left( \eta + \tilde{G} \right) \frac{F_{ij}(x)}{|F_{ij}(x)|}$  ▷ Performs a jump
23:        end if
24:      end if
25:    end while
26:  end for
27:  return
28: end procedure

```

3.1.1. Optimization of the Bound: The Ring Technique.

Let r_E be the Ewald cutoff, i.e., the maximum radius that is taken into account in the direct part and r_c the cutoff between the short-range interactions treated in the acceleration part and the long-range interactions treated with the jumps. In order to improve the bound on the jump rate (and therefore reduce the number of jump proposals), we can introduce an intermediate cutoff r_M ($r_c < r_M < r_E$) and compute two different bounds: a bound on the $\|F_{ij}\|$ such that $r_c - h \leq r_{ij} \leq r_M$ (h is the switching parameter) denoted L_M and a bound on the $\|F_{ij}\|$ such that $r_M \leq r_{ij} \leq r_E$ denoted L_L . For each i , we denote $N_M(i)$ and $N_L(i)$ the number of atoms j in these two cases, respectively. Now, the term $N \times L_i$ in the bound on the jump rate $\bar{\lambda}_i$ can be replaced by $N_M(i) L_M + N_L(i) L_L$. The algorithm (again in the case where ε is constant) can then be modified the following way:

1. Draw a proposition at rate $\bar{\lambda}_i$.
2. Draw $Y \in \{M, L\}$ such that $\mathbb{P}(Y = M) = \frac{L_M N_M(i)}{N_M(i) L_M + N_L(i) L_L}$ and $\mathbb{P}(Y = L) = \frac{L_L N_L(i)}{N_M(i) L_M + N_L(i) L_L}$.
3. Draw uniformly j in the right neighbor-lists (namely the “middle-range” or the “long-range” list).
4. Accept a jump of type (i, j) with probability $\frac{|F_{ij}(x)| \Theta(\eta)}{L_Y (B_i \varepsilon \sqrt{m_i} + 1 / \sqrt{2\pi})}$.

Remark. This modification requires to build new “middle” and “long” neighbor-lists, in order to know $N_L(i)$ and $N_M(i)$ for each atom and to draw uniformly among them.

3.2. Lennard–Jones Jumps. The same procedure can be used for van der Waals forces. In this case, the potential is given by

$$U_{\text{vdW}}(x_1, \dots, x_N) = \sum_{i \neq j} \varepsilon_{ij} \left[\left(\frac{\sigma_{ij}}{r_{ij}} \right)^{12} - \left(\frac{\sigma_{ij}}{r_{ij}} \right)^6 \right]$$

for some parameters $\sigma_{ij}, \varepsilon_{ij} > 0$. Since each term of this sum decreases very fast in r_{ij} , there is no splitting in direct and

reciprocal parts and all the terms beyond a certain cutoff r_{vdW} (i.e., for $r_{ij} > r_{\text{vdW}}$) are neglected. Again, we can define the pairwise term

$$U_{ij}^{\text{vdW}}(x) = \varepsilon_{ij} \left[\left(\frac{\sigma_{ij}}{r_{ij}} \right)^{12} - \left(\frac{\sigma_{ij}}{r_{ij}} \right)^6 \right]$$

choose a cutoff $r_c < r_{\text{vdW}}$ and use the switching function χ to separate each term into a short-range and a long-range part:

$$U_{ij}^{\text{vdW}} = \chi(r_{ij}) U_{ij}^{\text{vdW}} + (1 - \chi(r_{ij})) U_{ij}^{\text{vdW}}$$

From now, we define $F_{ij} = \nabla_{x_i} (1 - \chi(r_{ij})) U_{ij}$ and build a jump mechanism associated with each couple (i, j) . Here, since the van der Waals forces are weaker than the Coulomb ones at long-range, there are less stability issues and we can choose the parameter giving the least jumps: $\varepsilon = \infty$, which is the case that corresponds to the BPS. The jump rate associated with the force F_{ij} is

$$\lambda_{ij}(x, v) = \beta(v \cdot F_{ij}(x))_+$$

and the jumps are deterministic: $q_{ij}(x, v, dv') = \delta_{R_{ij}(x, v)}(dv')$ where

$$R_{ij}(x, v) = v - 2 \frac{v \cdot F_{ij}(x)}{|F_{ij}(x)|^2} F_{ij}(x)$$

If we denote L_i a bound on all the $\|F_{ij}\|$ for $1 \leq j \leq N$, the jump rate can be bounded by

$$\lambda_{ij}(x, v) \leq \beta |v| L_i := \lambda_{ij}^*$$

A jump being an orthogonal reflection of v against $F_{ij}(x)^\perp$, it does not modify $\|v\|$ and the bound λ_{ij}^* can stay the same during all the time step. It is therefore advantageous to use a BJAOAJB splitting with the Algorithm 3 described in 2.4. Moreover, as in the case of jumps on the direct electrostatic part, λ_{ij}^* is the same for all j , we can define $\bar{\lambda}_i = N \lambda_{ij}^*$ and the process associated with the velocity of each atom can be simulated independently. This yields the Algorithm 5, that was also described in ref 48.⁴⁸

Algorithm 5 Step (J): BPS for van de Waals interactions

Simulates the piecewise-constant jump process for a time $\delta/2$, starting from (x, v) .

Input : N the number of atoms, (x, v) their positions and velocities, δ the time step

```

1: procedure JUMP( $N, \delta, x, v$ )
2:   for  $i \in [1, N]$  do
3:      $M_i \leftarrow \text{POISSON}(N \beta |v_i| L_i \delta/2)$  ▷ Draw the number of jump proposals
4:     for  $k \in [1, M_i]$  do
5:        $J \leftarrow \text{RANDOM}([1, N])$  ▷ For each jump proposal  $k$ , choose randomly a jump of type  $J$ 
6:        $U \leftarrow \text{RANDOM}([0, 1])$ 
7:        $p \leftarrow \frac{|v_i \cdot F_{ij}(x)|}{|v_i| L_i}$  ▷ compute the probability of accepting the jump
8:       if  $U \leq p$  then
9:          $v_i \leftarrow R_{ij}(x, v)$  ▷ the velocity bounces
10:      end if
11:    end for
12:  end for
13:  return
14: end procedure

```

3.3. Conclusions: Three New Integrators. To conclude, using the jump mechanisms described in this section gives rise to three new integrators. First, the JUMP integrator consists in running either a BJAOAJB scheme, where the (J) part treats the Lennard-Jones jumps, a CAOAC scheme, where the (C) part treats the direct electrostatic jumps (along with the acceleration due to the rest of the forces) or, combining both, a C'J'AOAJC' scheme, where (J') treats the Lennard-Jones jumps, and (C') the electrostatic direct jumps. Regarding the

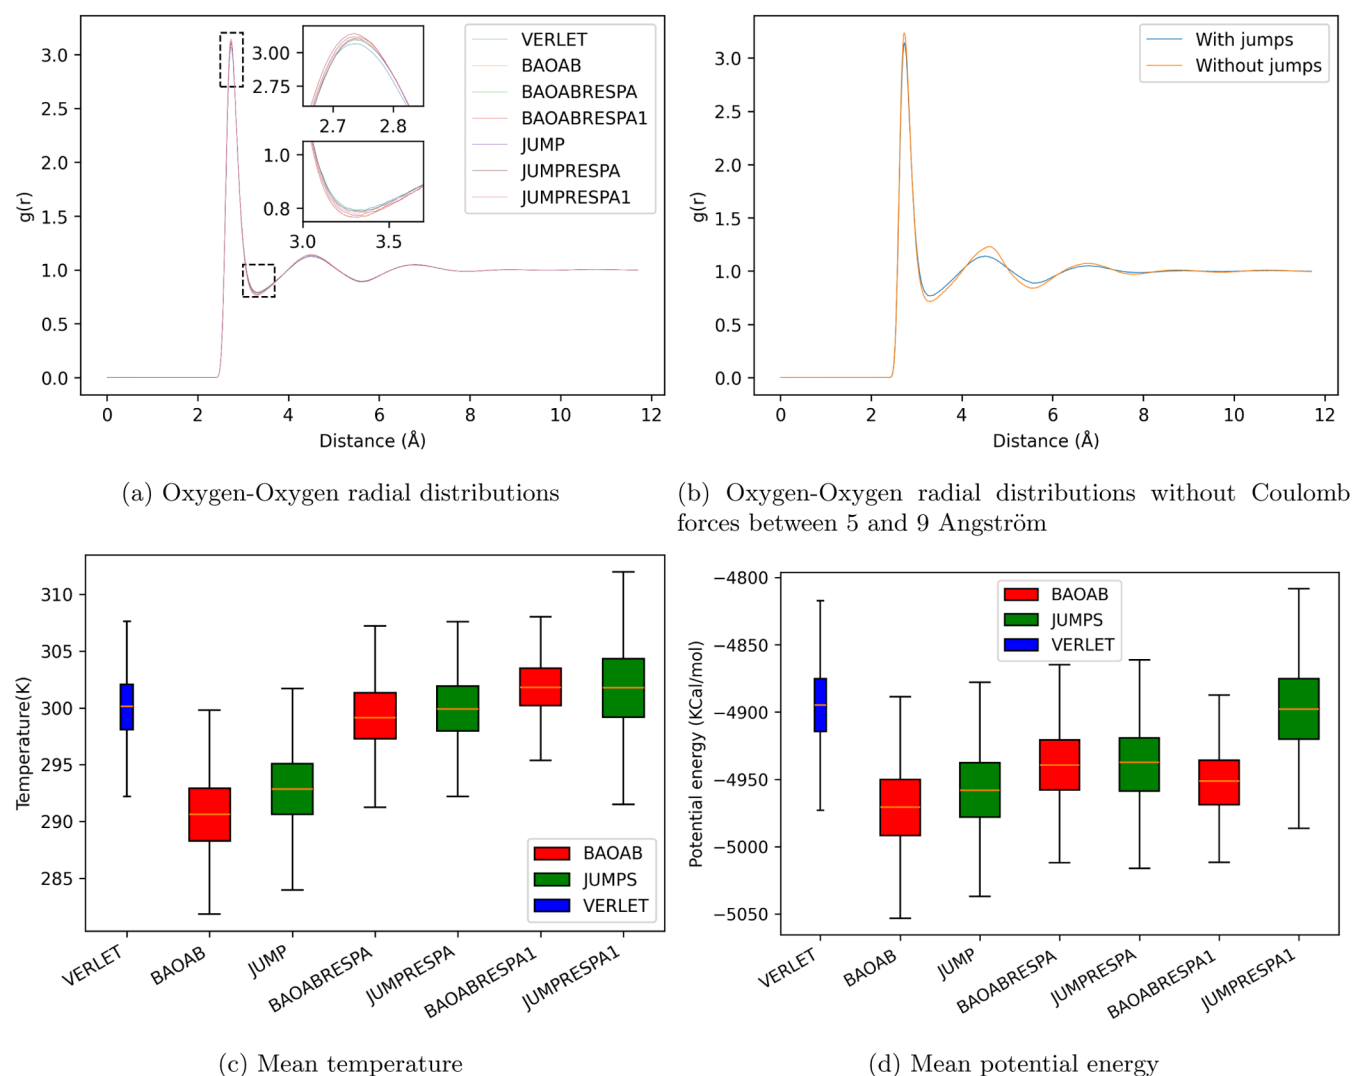


Figure 1. Sampling properties of the JUMP algorithms for a system of 500 water molecules.

multi-time-steps versions, JUMP-RESPA simply consists in treating in a small time-step δ all the intramolecular forces of the potential (bond stretching, angle bending, bond-angle cross terms, out-of-plane bending and torsional rotations), and leaving all the nonbonded forces (namely the van der Waals and Coulomb interactions) in a larger time step $\Delta = n\delta$ (with $n \in 2\mathbb{N}$), including the jumps. However, when using jumps, this is not very natural, as it increases the issues related to time discretization (unstability or numerical bias) while having no effect on the complexity of the jump parts, since with the thinning method, a step-size $n\delta$ requires on average n times more jumps (hence computations) than a step-size δ .

This remark leads to the following scheme: in JUMP-RESPA1, a small time step δ treats all the intramolecular forces, an intermediate time step $\kappa = m\delta$ treats the short-range van der Waals, the short-range direct electrostatics and the jumps (that take into account the long-range van der Waals forces and the long-range direct electrostatics), and a large time step $\Delta = n\kappa$ treats the reciprocal electrostatics, the many-body force that is, in practice (especially in large systems) the most numerically expensive force to compute. If we applied the algorithm to a polarizable force field such as AMOEBA,¹⁶ the expensive polarization part would also be treated in this largest time step.

In comparison, in BAOAB-RESPA1, the long-range van der Waals and long-range direct space electrostatics are treated in the larger time step. Here, in JUMP-RESPA1, those are treated by the jumps, in the intermediate time step. As discussed above, indeed, the jump process has the property of not requiring more computations when put in a smaller time step: the jump rate stays exactly the same.

4. NUMERICAL EXPERIMENTS

All variations of the algorithms described in Section 3 have been implemented in Tinker-HP software for molecular dynamics, both on the CPU and the parallel GPU versions. The software will be made freely available to academic users within the next release of the Tinker-HP code.^{21,29,65,66} Computations were done on the computing cluster of the

Table 1. Hydration Free Energies of Water, Sodium, and Potassium in 216 Water Molecules

integrator (kcal/Mol)	BAOAB	JUMPS
$\Delta G_{\text{hydrat}} \text{H}_2\text{O}$	6.73 ± 0.09	6.77 ± 0.08
$\Delta G_{\text{hydrat}} \text{Na}^+$	75.48 ± 0.09	75.61 ± 0.03
$\Delta G_{\text{hydrat}} \text{K}^+$	58.81 ± 0.04	58.82 ± 0.02




Figure 2. Diffusion coefficients for a system of 500 water molecules.

Theoretical Chemistry Laboratory (LCT) of Sorbonne University, Paris, on a 22 cores Intel(R) Xeon(R) Silver 4116 CPU @ 2.10 GHz for the CPU version, and a Nvidia Quadro GV100 for the GPU version. All numerical tests were done on water systems in the NVT ensemble with temperature $T = 300\text{K}$, of sizes ranging from 216 to 96000 molecules, with SPC/Fw model in the OPLS-AA force-field, in periodic boundary conditions. Simulation parameters are given at the end of this section, and Section S7, gives more details on the implementation.

Tuning of Parameters. As explained in Section 2.4, treating certain forces with jumps allows to reduce the total number of computations, and therefore accelerates the simulation. However, replacing a too large range of interactions by jumps yields either numerical instabilities (as it is the case for the BOUNCE algorithm described in⁴⁸) or a significant loss in the autodiffusion constant, which indicates a reduction of the sampling rate. As mentioned in³⁴ the ideal situation is when the acceleration rate is higher than the loss rate in the diffusion constant.

In order to do that, the choice of an adaptive ϵ helps the conservation of the diffusion, as the jump process will yield

milder jumps in the shorter-range regions (where the forces are stronger). Moreover, the section of pairwise interaction that are replaced by jumps should not be too large, or in too short-range regions. In other words, there should not be too many jumps, either because of too many jump processes running in parallel (if too many interactions are treated by jumps) or because the jump rates are too high (if those interactions are too strong). Those settings, although they can push further the computational speedup, always yields a loss in autodiffusion.

In other words, the best choice is to always choose an adaptive ϵ parameter, with a small enough ϵ_0 (so that the trajectories stay close to the Langevin process), but not too small so there is still a computational advantage, and to replace either a small portion of relatively short-range electrostatic interactions (for example between 5 and 7 Å), or a larger portion of longer-range interactions (for example between 7 and 11 Å).

The first setting seems to be the best for small systems, and the second one is more appropriate for larger systems, especially on the GPU version. Indeed, treating very long-range interactions with jumps allows to reduce the reciprocal Coulomb force, which is very expensive in large systems (it can take up to 50% of the total computation time). Moreover, since the reciprocal Coulomb force is partly responsible for resonance issues in multi-time-step integrators, reducing this part allows to extend the stability limit of the external time step for the JUMP-RESPA1 algorithm, which allows for a further increase in simulation speed.

Precision in Sampling and Dynamical Properties.

Figure 1 shows the oxygen–oxygen radial distributions, the mean temperature and mean potential energy of 500 water molecules, comparing the new JUMP, JUMP-RESPA and JUMP-RESPA1 integrators to its classical BAOAB, BAOAB-RESPA and BAOAB-RESPA1 counterparts, as well as the VERLET integrator for reference, with 2ns of simulation. To illustrate the non-negligible influence of the forces that are replaced by jumps, this figure also contains a radial distribution where long-range direct electrostatic interactions between 5 and 9 Ångström are removed from the potential. Table 1

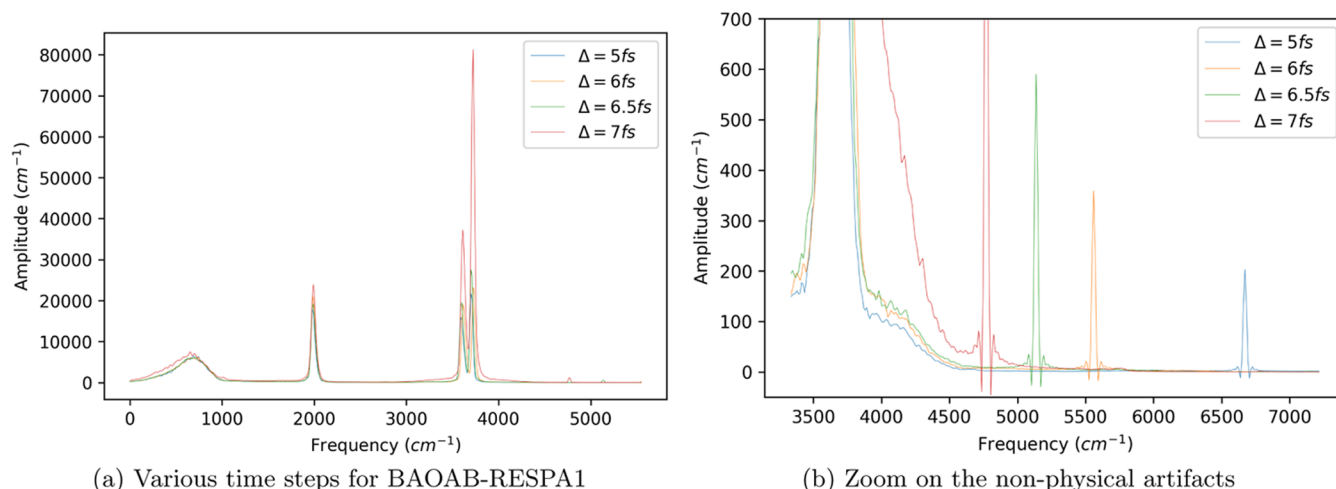


Figure 3. Velocity autocorrelation spectra of BAOAB-RESPA1 with various external time-steps, for a system of 500 water molecules. From left to right, the first bump corresponds to some many-body frequencies of the system, the first peak corresponds to intramolecular angle oscillations, the second and third peaks correspond to bond oscillations. All the peaks at the right of those are nonphysical artifacts due to the multi-time-step. On the zoomed graph, the main nonphysical peak moves to the left as the external time step increases, coupling with the bond-oscillation peak at the stability limit.

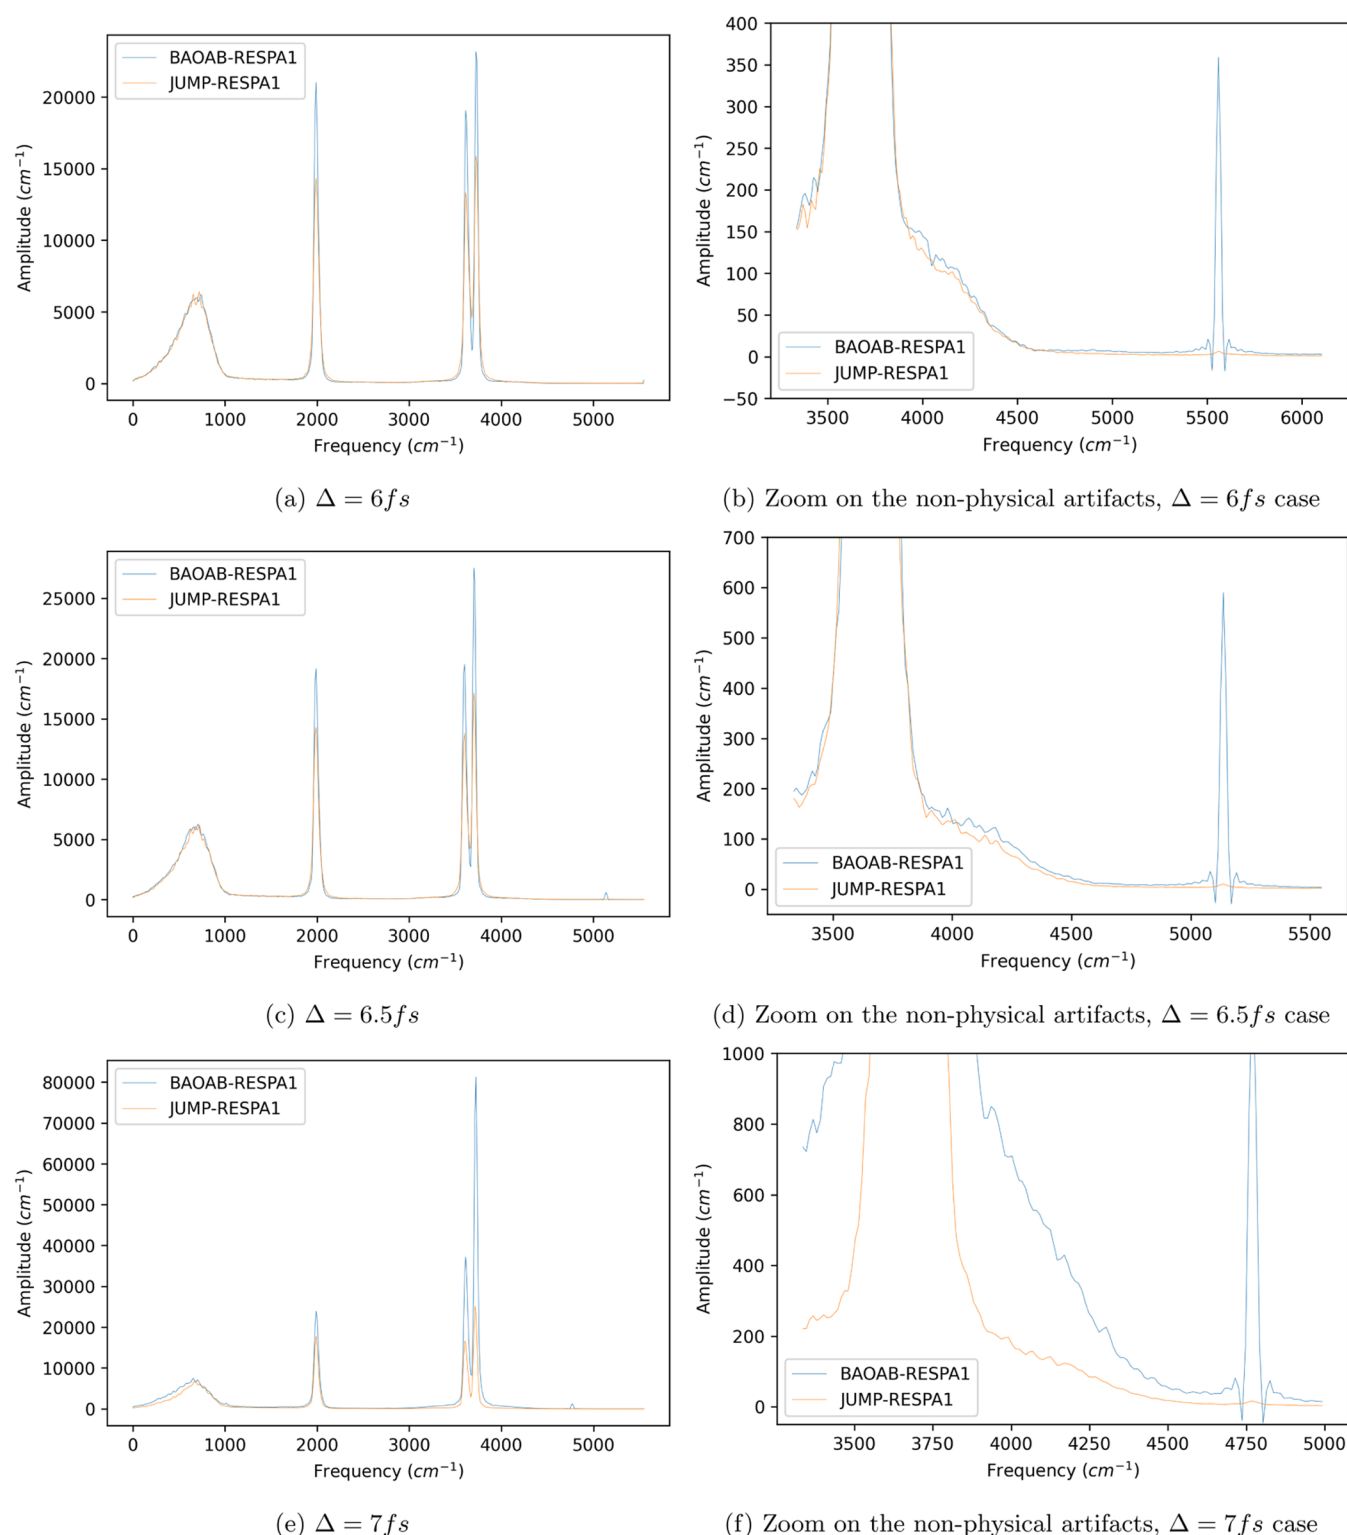


Figure 4. Velocity autocorrelation spectra of BAOAB-RESPA1 and JUMP-RESPA1 with various external time-steps. The nonphysical artifacts are largely suppressed in JUMP-RESPA1.

shows a comparison in hydration free energy of a water molecule, of a sodium ion Na^+ and of a potassium ion K^+ in a system of 216 water molecules simulated for 2 ns, using the Bennet Acceptance Ratio⁶⁷ (BAR) method. Note that the JUMP integrators are directly compatible with more advanced free energy methods.⁶⁸ The canonical measure is therefore similarly sampled by the JUMP integrators and the classical

ones. Figure 2 shows that the autodiffusion constants (computed using the Einstein formula) are well preserved, thanks to the proximity of the process with the Langevin dynamics, as explained in Section 2.1. Again, the tests are done on a system of 500 water molecules, simulated for 2 ns.

Resonance Effects. One of the well-known issues of multi-time-step methods are resonances effects between internal

Table 2. Speed Performances of the Algorithms, Expressed in Nanoseconds of Simulation per Day

number of atoms	CPU				GPU			
	1500	12,000	96,000	288,000	1500	12,000	96,000	288,000
BAOAB (1 fs)	3.50	3.95e-1	3.89e-2	8.15e-3	178.0	98.19	16.91	5.13
JUMP (1 fs)	4.74	4.94e-1	4.52e-2	8.74e-3	175.3	1.7.69	20.14	6.22
acceleration	35.5%	25.1%	16.2%	7.24%	−1.52%	9.68%	19.1%	21.3%
BAOAB-RESPA (1:2 fs)	5.68	6.29e-1	6.25e-2	1.46e-2	158.4	110.78	27.61	8.59
JUMP-RESPA (1:2 fs)	6.62	6.98e-1	6.24e-2	1.38e-2	154.7	115.93	30.58	9.82
acceleration	16.5%	11.0%	−0.16%	−5.48%	−2.34%	4.65%	10.8%	14.3%
BAOAB-RESPA1 (1:2:5 fs)	6.75	7.13e-1	6.73e-2	1.82e-2	276.6	141.3	34.69	11.93
JUMP-RESPA1 (1:2:6 fs)	8.21	8.30e-1	6.58e-2	1.81e-2	308.9	158.6	41.48	14.23
acceleration	21.6%	16.4%	−2.22%	−0.55%	11.7%	12.2%	19.6%	19.3%

Table 3. Time Spent in Each Routine

Nb of atoms	BAOAB on CPU			JUMP on CPU			
	vdW (%)	direct Coulomb (%)	SPME (%)	vdW (%)	direct Coulomb (%)	SPME (%)	jumps (%)
12,000	8.34	41.23	37.44	5.24	21.79	43.63	2.41
288,000	13.64	33.25	40.45	7.71	38.48	23.20	3.37
	BAOAB on GPU			JUMP on GPU			
	vdW (%)	direct Coulomb (%)	SPME (%)	vdW (%)	direct Coulomb (%)	SPME (%)	jumps (%)
12,000	25.5	34.0	17.7	26.2	21.9	18.6	8.0
280,000	8.5	32.3	52.0	6.6	46.7	28.0	7.4

Table 4. JUMP and PME-Grid Parameters

number of atoms	648	1500	12,000	96,000	288,000
box size (in Å)	18.643	24.662	49.232	98.5	142.27
default PME-grid	24 × 24 × 24	30 × 30 × 30	60 × 60 × 60	120 × 120 × 120	180 × 180 × 180
PME-grid in JUMP	24 × 24 × 24	30 × 30 × 30	60 × 60 × 60	96 × 96 × 96	128 × 128 × 128
α parameter in JUMP	0.554	0.554	0.554	0.42	0.3767
real-space cutoff JUMP (in Å)	7	7	7	9	10
short-range charge cutoff (in Å)	5	5	5	6	7
jump parameter ϵ_0	0.1	0.1	0.1	0.01	0.01

periodicities of the molecular system and the nonphysical periodicities created by the various time steps, that can induce numerical instabilities (such as abnormal fluctuations, or even explosion of the kinetic energy) and limit the largest time step that can be taken in the simulation. The reciprocal Coulomb part of the electrostatic interactions is partly responsible for those resonances (recall that this force is not purely long-range, since the erf function is not equal to zero around zero, and that the Ewald sum indeed implies all the atoms of the system). For this reason, increasing the real-space cutoff and replacing a part of the reciprocal electrostatic interactions by jumps, create much less oscillations in the system, thanks to the intrinsically random nature of velocity jumps, which allows to take a larger external time step in JUMP-RESPA1 than in BAOAB-RESPA1 and accelerate further the simulation.

These phenomena can be well observed on velocity autocorrelation spectra (computed from the MD trajectory at the level of the smallest time step using the implementation from ref 69), illustrated in Figures 3 and 4. On each graph, from left to right, the first band (below 1000 cm^{-1}) corresponds to molecule liberation and slow fluctuations in the hydrogen bond network, the first peak (around 2000 cm^{-1}) corresponds to intramolecular bending oscillations, the second and third peaks (around 3500 cm^{-1}) correspond to bond stretching motions. All the peaks at the right of those are nonphysical artifacts due to the multi-time-step. On the zoomed graphs, in BAOAB-RESPA1, the main nonphysical peak moves to the left when the external time step increases,

coming closer to the physical stretching peak, and even couples to it when the time step is too large. In JUMP-RESPA1, this phenomenon is largely suppressed, resulting in better numerical stability.

Performances of the Algorithm. In terms of simulation speed, the CPU version gives its best results on small systems (35.5% acceleration of JUMP with respect to BAOAB, and 21.6% of acceleration of JUMP-RESPA1 with respect to BAOAB-RESPA1 on the system of 1500 atoms), but this advantage disappears when looking at larger systems (with almost no difference between the algorithms when simulating 288000 atoms). On the highly parallel GPU implementation, however, it is the opposite: while there are almost no differences between JUMP and BAOAB on small systems (mostly because the GPU is not fully exploited on those, since the computational resources are not saturated, so tests on large molecular systems are more appropriate to this setting), the advantage appears on larger systems, and even is the best on the largest tested system (21.3% acceleration of JUMP with respect to BAOAB and 19.3% acceleration of JUMP-RESPA1 with respect to BAOAB-RESPA1 on the system of 288000 atoms). Table 2 shows the speed of the algorithms (in nanoseconds of simulation per day) on CPU and GPU implementations.

Those differences between GPU and CPU have different reasons. First, the neighbor lists are treated differently, see the end of this section for more details. Then, the proportion of time spent in the different routines of the integrator is different

on CPU and GPU. For instance, the parallel architecture of GPUs allows for a more efficient treatment of the van der Waals and direct electrostatic interactions, while the many-body Ewald sum takes almost 50% of the total computation time, especially on larger systems. As a consequence, on a large system, reducing the reciprocal part of electrostatics and treating a larger portion of the potential with jumps yields a significant computational advantage. For reference, Table 3 gives the proportion of time spent in each of the main routines of the CPU and GPU implementations, comparing two system sizes.

As explained in,²⁷ when increasing the real-space cutoff r_c , the Fourier space spacing can be multiplied by the same factor, which reduces the cost of its computation. It is particularly advantageous to do that on the large systems, where the interpolation of the structure factors of the Ewald sum is done on large grids (Table 4).

Simulation details

- Tinker-HP version: 1.2
- Force-field: OPLS-AA
- Water model: SPC/Fw
- switching parameter: 0.5 Å
- Friction: 1.0ps^{-1}
- Temperature: 300K.

Parameters for BAOAB, BAOAB-RESPA and BAOAB-RESPA1 simulations:

- Time step for BAOAB: $\delta = 1\text{fs}$.
- Time steps for BAOAB-RESPA: $\delta = 1\text{fs}$, $\Delta = 2\text{fs}$.
- Time steps for BAOAB-RESPA1: $\delta = 1\text{fs}$, $\kappa = 2\text{fs}$, $\Delta = 5\text{fs}$.
- Charge real-space cutoff: 7 Å.
- Ewald parameter: $\alpha = 0.5446$.
- Van de Waals long-range cutoff: 10 Å.

Parameters for JUMP, JUMP-RESPA, and JUMP-RESPA1 simulations:

- Time step for JUMP: $\delta = 1\text{fs}$.
- Time steps for JUMP-RESPA: $\delta = 1\text{fs}$, $\Delta = 2\text{fs}$.
- Time steps for JUMP-RESPA1: $\delta = 1\text{fs}$, $\kappa = 2\text{fs}$, $\Delta = 6.5\text{fs}$.
- van der Waals jumps only on CPU.
- Van de Waals long-range cutoff: 10 Å.
- Van de Waals short-range cutoff (if activated): 7 Å.
- Type of jump parameter: adaptive ϵ .

5. CONCLUSIONS AND PERSPECTIVES

To conclude, our work presents a new framework to construct velocity jumps (or JUMP) integrators for molecular dynamics simulations that can be parallelized on GPU and integrated into multi-time-step methods. It allows to accelerate simulations while preserving sampling precision and dynamical properties, and avoiding certain classical resonance issues. The JUMP approaches provide significant acceleration over their BAOAB, BAOAB-RESPA and BAOAB-RESPA1 counterparts. If the present implementation has been performed in the Tinker-HP package, the approach could be easily ported to other molecular dynamics software. Future work will deal with the extension of the approach to polarizable force fields^{19,20,70,71} and with its combination with advanced periodic boundary conditions treatment of electrostatics such as ANKH.⁷²

■ ASSOCIATED CONTENT

Supporting Information

The Supporting Information is available free of charge at <https://pubs.acs.org/doi/10.1021/acs.jctc.5c00023>.

Rigorous definitions on the continuous-time process and splitting schemes; proofs of the mathematical statements as well as technical details on the Tinker-HP implementation of the algorithms (PDF)

■ AUTHOR INFORMATION

Corresponding Authors

Pierre Monmarché — Sorbonne Université, CNRS, Paris 75005, France; Sorbonne Université, Université Paris Cité, CNRS, INRIA, Paris 75005, France; Institut Universitaire de France, Paris 75005, France; orcid.org/0000-0002-3356-2646; Email: pierre.monmarche@sorbonne-universite.fr

Jean-Philip Piquemal — Sorbonne Université, CNRS, Paris 75005, France; Advanced Research Department, Qubit Pharmaceuticals, Paris 75014, France; orcid.org/0000-0001-6615-9426; Email: jean-philip.piquemal@sorbonne-universite.fr

Authors

Nicolai Gouraud — Sorbonne Université, CNRS, Paris 75005, France; Sorbonne Université, Université Paris Cité, CNRS, INRIA, Paris 75005, France; Advanced Research Department, Qubit Pharmaceuticals, Paris 75014, France

Louis Lagardère — Sorbonne Université, CNRS, Paris 75005, France; Advanced Research Department, Qubit Pharmaceuticals, Paris 75014, France; orcid.org/0000-0002-7251-0910

Olivier Adjoua — Sorbonne Université, CNRS, Paris 75005, France

Thomas Plé — Sorbonne Université, CNRS, Paris 75005, France

Complete contact information is available at: <https://pubs.acs.org/doi/10.1021/acs.jctc.5c00023>

Notes

The authors declare the following competing financial interest(s): Louis Lagardère and Jean-Philip Piquemal are co-founders and shareholders of Qubit Pharmaceuticals.

■ ACKNOWLEDGMENTS

This work was made possible thanks to funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program (grant agreement No 810367), project EMC2 (JPP). Computations have been performed at IDRIS, GENCI (Jean Zay machine, France) on grant no A0130712052 (JPP).

■ REFERENCES

- (1) Allen, M. P.; Tildesley, D. J. *Computer Simulation of Liquids*; Oxford University Press, 2017.
- (2) Leimkuhler, B.; Matthews, C. *Molecular Dynamics*; Springer, 2015.
- (3) Goh, B. C.; Hadden, J. A.; Bernardi, R. C.; Singharoy, A.; McGreevy, R.; Rudack, T.; Cassidy, C. K.; Schulten, K. Computational methodologies for real-space structural refinement of large macromolecular complexes. *Annual Review of Biophysics. Annu. Rev. Biophys.* **2016**, *45* (1), 253–278.
- (4) Mortier, J.; Rakers, C.; Bermudez, M.; Murgueitio, M. S.; Riniker, S.; Wolber, G. The impact of molecular dynamics on drug

design: applications for the characterization of ligand–macromolecule complexes. *Drug Discovery Today* **2015**, *20* (6), 686–702.

(5) Walsh, T. R.; Knecht, M. R. Biointerface structural effects on the properties and applications of bioinspired peptide-based nanomaterials. *Chem. Rev.* **2017**, *117* (20), 12641–12704.

(6) Blazhynska, M.; Lagardère, L.; Liu, C.; Adjoua, O.; Ren, P.; Piquemal, J.-P. Water-glycan interactions drive the SARS-CoV-2 spike dynamics: insights into glycan-gate control and camouflage mechanisms. *Chem. Sci.* **2024**, *15*, 14177–14187.

(7) Freddolino, P.; Harrison, C.; Liu, Y.; Schulten, K. Challenges in protein folding simulations: Timescale, representation, and analysis. *Nat. Phys.* **2010**, *6*, 751–758.

(8) Jung, J.; Nishima, W.; Daniels, M.; Bascom, G.; Kobayashi, C.; Adedoyin, A.; Wall, M.; Lappala, A.; Phillips, D.; Fischer, W.; Tung, C.-S.; Schlick, T.; Sugita, Y.; Sanbonmatsu, K. Y. Scaling molecular dynamics beyond 100,000 processor cores for large-scale biophysical simulations. *J. Comput. Chem.* **2019**, *40* (21), 1919–1930.

(9) Shaw, D. E.; Dror, R. O.; Salmon, J. K.; Grossman, J. P.; Mackenzie, K. M.; Bank, J. A.; Young, C.; Deneroff, M. M.; Batson, B.; Bowers, K. J.; Chow, E.; Eastwood, M. P.; Ierardi, D. J.; Klepeis, J. L.; Kuskin, J. S.; Larson, R. H.; Lindorff-Larsen, K.; Maragakis, P.; Moraes, M. A.; Piana, S.; Shan, Y.; Towles, B. *Millisecond-scale Molecular Dynamics Simulations on Anton*; Proceedings of the Conference on High Performance Computing Networking, Storage and Analysis; ACM: New York, NY, USA, 2009; pp 39:1–39:11.

(10) Shibuta, Y.; Sakane, S.; Miyoshi, E.; Takaki, T.; Ohno, M. Micrometer-scale molecular dynamics simulation of microstructure formation linked with multi-phase-field simulation in same space scale. *Modelling Simul. Mater. Sci. Eng.* **2019**, *27* (5), No. 054002.

(11) Yu, I.; Mori, T.; Ando, T.; Harada, R.; Jung, J.; Sugita, Y.; Feig, M. Biomolecular interactions modulate macromolecular structure and dynamics in atomistic model of a bacterial cytoplasm. *eLife* **2016**, *5*, No. e19274.

(12) Huang, J.; Rauscher, S.; Nawrocki, G.; Ran, T.; Feig, M.; de Groot, B. L.; Grubmueller, H.; MacKerell, A. D., Jr CHARMM36m: an improved force field for folded and intrinsically disordered proteins. *Nat. Methods* **2017**, *14* (1), 71–73.

(13) Maier, J. A.; Martinez, C.; Kasavajhala, K.; Wickstrom, L.; Hauser, K. E.; Simmerling, C. ff14sb: Improving the Accuracy of Protein Side Chain and Backbone Parameters from ff99sb. *J. Chem. Theory Comput.* **2015**, *11* (8), 3696–3713.

(14) Robertson, M. J.; Tirado-Rives, J.; Jorgensen, W. L. Improved Peptide and Protein Torsional Energetics with the OPLS-AA Force Field. *J. Chem. Theory Comput.* **2015**, *11* (7), 3499–3509.

(15) Reif, M. M.; Hünenberger, P. H.; Oostenbrink, C. New Interaction Parameters for Charged Amino Acid Side Chains in the GROMOS Force Field. *J. Chem. Theory Comput.* **2012**, *8* (10), 3705–3723.

(16) Ren, P.; Ponder, J. W. Polarizable Atomic Multipole Water Model for Molecular Mechanics Simulation. *J. Phys. Chem. B* **2003**, *107* (24), S933–S947.

(17) Naseem-Khan, S.; Lagardère, L.; Narth, C.; Cisneros, G. A.; Ren, P.; Gresh, N.; Piquemal, J.-P. Development of the quantum-inspired sibfa many-body polarizable force field: Enabling condensed-phase molecular dynamics simulations. *J. Chem. Theory Comput.* **2022**, *18* (6), 3607–3621.

(18) Ponder, J. W.; Case, D. A. Force Fields for Protein Simulations. In *Protein Simulations*; Academic Press, 2003; Vol. 66, pp 27–85.

(19) Melcr, J.; Piquemal, Jean-Philip. Accurate biomolecular simulations account for electronic polarization. *Front. Mol. Biosci.* **2019**, *6*, No. 143.

(20) Jing, Z.; Liu, C.; Cheng, S. Y.; Qi, R.; Walker, B. D.; Piquemal, J.-P.; Ren, Pengyu. Polarizable force fields for biomolecular simulations: Recent advances and applications. *Annu. Rev. Biophys.* **2019**, *48* (1), 371–394.

(21) Adjoua, O.; Lagardère, L.; Jolly, L.-H.; Durocher, A.; Very, T.; Dupays, I.; Wang, Z.; Inizan, T. J.; Célerse, F.; Ren, P.; Ponder, J.; Piquemal, J.-P. Tinker-hp: Accelerating Molecular Dynamics Simulations of Large Complex Systems with Advanced Point Dipole

Polarizable Force Fields Using GPUs and Multi-GPU Systems. *J. Chem. Theory Comput.* **2021**, *17* (03), 2034–2053.

(22) Eastman, P.; Swails, J.; Chodera, J. D.; McGibbon, R. T.; Zhao, Y.; Beauchamp, K. A.; Wang, L.-P.; C. Simmonett, A.; Harrigan, M. P.; Stern, C. D.; et al. Openmm 7: Rapid development of high performance algorithms for molecular dynamics. *PLoS Comput. Biol.* **2017**, *13* (7), No. e1005659.

(23) Kutzner, C.; Páll, S.; Fechner, M.; Esztermann, A.; de Groot, B. L.; Grubmüller, H. Best bang for your buck: GPU nodes for GROMACS biomolecular simulations. 2015.

(24) Salomon-Ferrer, R.; Gotz, A. W.; Poole, D.; Le Grand, S.; Walker, R. C. Routine microsecond molecular dynamics simulations with AMBER on GPUs. 2. Explicit solvent particle mesh Ewald. *J. Chem. Theory Comput.* **2013**, *9* (9), 3878–3888.

(25) Thompson, A. P.; Aktulga, H. M.; Berger, R.; Bolintineanu, D. S.; Brown, W. M.; Crozier, P. S.; in 't Veld, P. J.; Kohlmeyer, A.; Moore, S. G.; Nguyen, T. D.; Shan, R.; Stevens, M. J.; Tranchida, J.; Trott, C.; Plimpton, S. J. LAMMPS - a flexible simulation tool for particle-based materials modeling at the atomic, meso, and continuum scales. *Comput. Phys. Commun.* **2022**, *271*, No. 108171.

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

(27) Abraham, M. J.; Murtola, T.; Schulz, R.; Pall, S.; Smith, J. C.; Hess, B.; Lindahl, E. GROMACS: High performance molecular simulations through multi-level parallelism from laptops to supercomputers. *SoftwareX* **2015**, *1–2*, 19–25.

(28) Salomon-Ferrer, R.; Case, D. A.; Walker, R. C. An overview of the Amber biomolecular simulation package. *WIREs Comput. Mol. Sci.* **2013**, *3* (2), 198–210.

(29) Lagardère, L.; Jolly, L.-H.; Lipparini, F.; Aviat, F.; Stamm, B.; Jing, Z. F.; Harger, M.; Torabifard, H.; Cisneros, G. A.; Schnieders, M. J.; Gresh, N.; Maday, Y.; Ren, P. Y.; Ponder, J. W.; Piquemal, J.-P. Tinker-hp: a massively parallel molecular dynamics package for multiscale simulations of large complex systems with advanced point dipole polarizable force fields. *Chem. Sci.* **2018**, *9*, 956–972.

(30) Hynninen, A.-P.; Crowley, M. F. New faster CHARMM molecular dynamics engine. *J. Comput. Chem.* **2014**, *35* (5), 406–413.

(31) Kobayashi, C.; Jung, J.; Matsunaga, Y.; Mori, T.; Ando, T.; Tamura, K.; Kamiya, M.; Sugita, Y. Genesis 1.1: A hybrid-parallel molecular dynamics simulator with enhanced sampling algorithms on multiple computational platforms. *J. Comput. Chem.* **2017**, *38* (25), 2193–2206.

(32) Plimpton, S. Fast Parallel Algorithms for Short-Range Molecular Dynamics. *J. Comput. Phys.* **1995**, *117* (1), 1–19.

(33) Gibson, D. A.; Carter, E. A. Time-reversible multiple time scale ab initio molecular dynamics. *J. Phys. Chem. A* **1993**, *97*, 13429–13434.

(34) Lagardère, L.; Aviat, F.; Piquemal, J.-P. Pushing the limits of Multiple-Timestep Strategies for Polarizable Point Dipole Molecular Dynamics. *J. Phys. Chem. Lett.* **2019**, *10* (10), 2593–2599.

(35) Tuckerman, M. E.; Berne, B. J.; Rossi, A. Molecular dynamics algorithm for multiple time scales: Systems with disparate masses. *J. Chem. Phys.* **1991**, *94*, 1465–1469, DOI: 10.1063/1.460004.

(36) Tuckerman, M.; Berne, B. J.; Martyna, G. J. Reversible multiple time scale molecular dynamics. *J. Chem. Phys.* **1992**, *97* (3), 1990–2001.

(37) Bou-Rabee, N. Time Integrators for Molecular Dynamics. *Entropy* **2014**, *16*, 138–162.

(38) Leimkuhler, B.; Matthews, C. Robust and efficient configurational molecular sampling via Langevin dynamics. *J. Chem. Phys.* **2013**, *138* (17), No. 174102.

(39) Leimkuhler, B.; Matthews, C.; Stoltz, G. The computation of averages from equilibrium and nonequilibrium Langevin molecular dynamics. *IMA J. Numer. Anal.* **2016**, *36* (1), 13–79.

(40) Matthews, C.; Leimkuhler, B. Rational Construction of Stochastic Numerical Methods for Molecular Sampling. *Appl. Math. Res. eXpress* **2012**, *2013* (1), 34–56.

- (41) Biesiadecki, J. J.; Skeel, R. D. Dangers of Multiple Time Step Methods. *J. Comput. Phys.* **1993**, *109* (2), 318–328.
- (42) Ma, Q.; Izaguirre, J.; Skeel, R. Verlet-i/r-respa/impulse is Limited by Nonlinear Instabilities. *SIAM J. Sci. Comput.* **2003**, *24* (6), 1951–1973.
- (43) Morrone, J. A.; Markland, T. E.; Ceriotti, M.; Berne, B. J. Efficient multiple time scale molecular dynamics: Using colored noise thermostats to stabilize resonances. *J. Chem. Phys.* **2011**, *134* (1), No. 014103.
- (44) Albaugh, A.; Tuckerman, M. E.; Head-Gordon, T. Combining Iteration-Free Polarization with Large Time Step Stochastic-Isokinetic Integration. *J. Chem. Theory Comput.* **2019**, *15* (4), 2195–2205.
- (45) Leimkuhler, B.; Margul, D. T.; Tuckerman, M. E. Stochastic, resonance-free multiple time-step algorithm for molecular dynamics with very large time steps. *Mol. Phys.* **2013**, *111* (22–23), 3579–3594.
- (46) Margul, D. T.; Tuckerman, M. E. A Stochastic, Resonance-Free Multiple Time-Step Algorithm for Polarizable Models That Permits Very Large Time Steps. *J. Chem. Theory Comput.* **2016**, *12* (5), 2170–2180.
- (47) Feenstra, K. A.; Hess, B.; Berendsen, H. J. C. Improving efficiency of large time-scale molecular dynamics simulations of hydrogen-rich systems. *J. Comput. Chem.* **1999**, *20* (8), 786–798.
- (48) Monmarché, P.; Weisman, J.; Lagardère, L.; Piquemal, J.-P. Velocity jump processes: An alternative to multi-timestep methods for faster and accurate molecular dynamics simulations. *J. Chem. Phys.* **2020**, *153* (2), No. 024101.
- (49) Gouraud, N.; Journal, L.; Monmarché, P. The velocity jump Langevin process and its splitting scheme: long time convergence and numerical accuracy. 2024. arXiv:2410.08846. arXiv.org e-Printarchive. <https://doi.org/10.48550/arXiv.2410.08846>.
- (50) Monmarché, P.; Rousset, M.; Zitt, P.-A. Exact targeting of Gibbs distributions using velocity-jump processes. *Stoch. Partial Differ. Equ.: Anal. Comput.* **2022**, *11*, 908–947, DOI: 10.1007/s40072-022-00247-9.
- (51) Monmarché, P. Piecewise deterministic simulated annealing. *ALEA* **2016**, *13* (1), 357–398.
- (52) Peters, E. A. J. F.; de With, G. Rejection-free Monte Carlo sampling for general potentials. *Phys. Rev. E* **2012**, *85*, No. 026703.
- (53) Vanetti, P.; Bouchard-Côté, A.; Deligiannidis, G.; Doucet, A. Piecewise Deterministic Markov Chain Monte Carlo. 2017. arXiv:1707.05296. arXiv.org e-Printarchive. <https://doi.org/10.48550/arXiv.1707.05296>.
- (54) Bertazzi, A.; Bierkens, J.; Dobson, P. Approximations of Piecewise Deterministic Markov Processes and their convergence properties. *Stoch. Process. Appl.* **2022**, *154*, 91–153.
- (55) Bertazzi, A.; Dobson, P.; Monmarché, P. Splitting schemes for second order approximations of piecewise-deterministic Markov processes. *arXiv e-prints* **2023**.
- (56) Bierkens, J.; Fearnhead, P.; Roberts, G. The Zig-Zag process and super-efficient sampling for Bayesian analysis of big data. *Ann. Stat.* **2019**, *47* (3), 1288–1320.
- (57) Corbella, A.; Spencer, S. E. F.; Roberts, G. O. Automatic Zig-Zag sampling in practice. *Stat. Comput.* **2022**, *32* (6), No. 107.
- (58) Pagani, F.; Chevallier, A.; Power, S.; House, T.; Cotter, S. Nuzz: Numerical Zig-Zag for general models. *Stat. Comput.* **2024**, *34* (1), No. 61.
- (59) Bouchard-Côté, A.; Vollmer, S. J.; Doucet, A. The bouncy particle sampler: a nonreversible rejection-free Markov chain Monte Carlo method. *J. Am. Stat. Assoc.* **2018**, *113* (S22), 855–867.
- (60) Lemaire, V.; Thiell, M.; Thomas, N. Exact Simulation of the Jump Times of a Class of Piecewise Deterministic Markov Processes. *J. Sci. Comput.* **2018**, *75* (3), 1776–1807.
- (61) Lewis, P. A. W.; Shedler, G. S. Simulation of nonhomogeneous poisson processes by thinning. *Nav. Res. Logist. Q.* **1979**, *26* (3), 403–413.
- (62) Wang, J.; Cieplak, P.; Kollman, P. A. How well does a restrained electrostatic potential (RESP) model perform in calculating conformational energies of organic and biological molecules? *J. Comput. Chem.* **2000**, *21* (12), 1049–1074.
- (63) Essmann, U.; Perera, L.; Berkowitz, M. L.; Darden, T.; Lee, H.; Pedersen, L. G. A smooth particle mesh Ewald method. *J. Chem. Phys.* **1995**, *103* (19), 8577–8593.
- (64) Morrone, J. A.; Zhou, R.; Berne, B. J. Molecular Dynamics with Multiple Time Scales: How to Avoid Pitfalls. *J. Chem. Theory Comput.* **2010**, *6* (6), 1798–1804.
- (65) Tinker-HP's github; Github. <https://github.com/TinkerTools/Tinker-HP>. Accessed: September 29, 2024.
- (66) Tinker-HP's Website; Tinker-HP. <https://tinker-hp.org/>. Accessed: September 29, 2024.
- (67) Bennett, C. H. Efficient estimation of free energy differences from Monte Carlo data. *J. Comput. Phys.* **1976**, *22* (2), 245–268.
- (68) Lagardère, L.; Maurin, L.; Adjoua, O.; Hage, K. E.; Monmarché, P.; Piquemal, J.-P.; Hénin, J. Lambda-abc: Simplified, portable, accurate, and cost-effective alchemical free-energy computation. *J. Chem. Theory Comput.* **2024**, *20*, 4481–4498, DOI: 10.1021/acs.jctc.3c01249.
- (69) Plé, T.; Mauger, N.; Adjoua, O.; Inizan, T. J.; Lagardère, L.; Huppert, S.; Piquemal, J.-P. Routine molecular dynamics simulations including nuclear quantum effects: From force fields to machine learning potentials. *J. Chem. Theory Comput.* **2023**, *19* (5), 1432–1445.
- (70) Gresh, N.; Cisneros, G. A.; Darden, T. A.; Piquemal, J.-P. Anisotropic, polarizable molecular mechanics studies of inter- and intramolecular interactions and ligand-macromolecule complexes. a bottom-up strategy. *J. Chem. Theory. Comput.* **2007**, *3* (6), 1960–1986.
- (71) Shi, Y.; Ren, P.; Schnieders, M.; Piquemal, J.-P. *Polarizable Force Fields for Biomolecular Modeling*; John Wiley & Sons, Ltd, 2015; pp 51–86.
- (72) Chollet, I.; Lagardère, L.; Piquemal, J.-P. Ankh: A generalized $\mathcal{O}(n)$ interpolated ewald strategy for molecular dynamics simulations. *J. Chem. Theory. Comput.* **2023**, *19* (10), 2887–2905.



CAS INSIGHTS™

EXPLORE THE INNOVATIONS SHAPING TOMORROW

Discover the latest scientific research and trends with CAS Insights. Subscribe for email updates on new articles, reports, and webinars at the intersection of science and innovation.

[Subscribe today](#)

CAS
A division of the
American Chemical Society