



HAL
open science

Necessary and sufficient symmetries in Event-Chain Monte Carlo with generalized flows and Application to hard dimers

Tristan Guyon, Arnaud Guillin, Manon Michel

► **To cite this version:**

Tristan Guyon, Arnaud Guillin, Manon Michel. Necessary and sufficient symmetries in Event-Chain Monte Carlo with generalized flows and Application to hard dimers. 2023. hal-04164032

HAL Id: hal-04164032

<https://hal.science/hal-04164032v1>

Preprint submitted on 18 Jul 2023

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.

Necessary and sufficient symmetries in Event-Chain Monte Carlo with generalized flows and Application to hard dimers

Tristan Guyon,^{*} Arnaud Guillin,[†] and Manon Michel[‡]
*Laboratoire de Mathématiques Blaise Pascal UMR 6620,
CNRS, Université Clermont-Auvergne, Aubière, France.*

Event-Chain Monte Carlo methods generate continuous-time and non-reversible Markov processes which often display important accelerations compared to their reversible counterparts. However their generalization to any system may appear less straightforward. In this work, we build on the recent analytical characterization of such methods as generating Piecewise Deterministic Markov Processes (PDMP) to clearly decipher the necessary symmetries the PDMP must obey from the sufficient ones which may prove to be too restrictive in a general setting. Thus, we derive a necessary rotational invariance of the probability flows and the minimum event rate, which identifies with the corresponding infinitesimal rejection rate. Such conditions always yield a correct ECMC scheme. We then generalize such results to the case of more general deterministic flows than the translational ones. In particular, we define two classes of interest of general flows, the *ideal* and *uniform-ideal* ones, which respectively suppresses or reduces the event rates. From there, we implement a complete non-reversible sampling of a systems of hard dimers, thanks to the introduction of rotational flows, which are uniform-ideal and shows a speed-up of up to ~ 3 compared to the state-of-the-art ECMC/Metropolis hybrid scheme.

I. INTRODUCTION

Since the introduction of the seminal Metropolis algorithm [1] first applied to hard-sphere systems, Markov-chain Monte Carlo methods have become an ubiquitous tool in computational physics [2]. They produce numerical evaluation of high-dimensional integrals, often arising from a stochastic description linked to a Boltzmann probability distribution, by a discrete sum over random configuration samples. Such samples are outputted along a Markov process which explores the system configuration space while leaving the target Boltzmann distribution invariant. The quality of such numerical approximation is directly linked to the capacity of the Markov process to produce a sequence of samples as uncorrelated as possible [3]. The Metropolis algorithm generates reversible Markov processes, which rely on rejections to target the correct invariant distribution. It then typically displays a diffusive dynamics, most often leading to expensive correlation times [4]. This situation can be further worsened in presence of critical slowing down phenomena at phase transitions [5], where the correlation length diverges and spans the full system. Therefore, in order to alleviate finite-size effects, a large community effort is devoted to the developments of MCMC methods which exhibit better performance and scalability with system sizes.

While the efficient clusters methods [6, 7] have been derived for lattice spin systems thanks to the spin-flip symmetry they present, the lack of a natural involution in continuous particle systems has motivated the development of non-reversible MCMC methods, known

as Event-Chain Monte Carlo (ECMC) [8, 9]. ECMC generates persistent moves, where singles particles are sequentially updated along ballistic trajectories forming up a *chain* whose sequence is controlled by some jump process, the *events*, reacting to the energy changes along the deterministic flow by resampling its parameters by some Markov kernel. The scheme then is rejection-free and relies on a control by the events of the ballistic exploration to ensure the correct invariant distribution. The observed accelerations have motivated the applications of such methods to a large variety of systems, as polymers [10] or continuous spins [11, 12]. It led to further generalization, mainly revolving around the exploitation of more global symmetries at the events, in order to maintain the necessary control on the ballistic exploration, while avoiding backtracking as much as possible. Thus, starting with the initial pairwise symmetry [9], events can now exploit translational [13] and rotational [14] symmetry or even the addition of some artificial kinetic energy [15].

While improving the control of the ballistic flow at the events has been under much focus, the development of different ballistic flows other than the standard translational updates has received smaller attention, see for example [16, 17]. Indeed, more involved flow schemes may lead to an increased difficulty in computing the event time in complex systems. However, the potential accelerations could counterbalance this effect. More importantly, any ECMC application to anisotropic particles, as dimers or more generally molecules, requires the introduction of sequences of rotations, in order to thermalize all degrees of freedom. Because of the lack of rotational-flow schemes, ECMC has thus up to now only been applied to tethered-type [13, 18] or elastic [13] interactions in anisotropic particles, so that translations can still thermalize the systems, or has been coupled to some Metropolis scheme to generate rotational moves [19]. In

^{*} tristan.guyon@cnrs.fr

[†] arnaud.guillin@uca.fr

[‡] manon.michel@cnrs.fr

addition to simulating all anisotropic models, devising non-reversible general flows may further lead to accelerations, as [19] observes that the more anisotropic the particles in a system are, the more important the observed accelerations is when introducing non reversibility. Finally, the design of generalized flows more generally raises the question around the understanding of the necessary fundamental conditions in ECMC versus the convenient sufficient choices, which is reminiscent of the global balance vs detailed balance question.

In this work, we directly address this question, which allows us to explicitly generalize the ECMC methods to more general ballistic flows. To do so, we build on the characterization of ECMC as Piecewise deterministic Markov processes (PDMP) [20] to decipher precisely the necessary conditions from the sufficient choices. We do so first for the standard case of translational flows and we establish the necessary rotational invariance of the probability flows in the general case of local events and the fundamental minimum event rate, which comes down to the corresponding infinitesimal Metropolis rejection rate. This reveals that one of the state-of-the-art ECMC schemes, known as the Forward variant or more precisely the direct event kernel [14], actually does not require any further restrictive conditions to be implemented, as was previously assumed. We then consider the case of general flows, where we generalize the necessary conditions of rotational invariance and minimum event rate. It then leads us to define *ideal* and *uniform-ideal* flows, which respectively suppress the minimum event rate in general or only for uniform stationary distributions. We then find that the Markov kernels used in the translational case can be directly adapted to any generalized flow. While *ideal* flows may be out of reach in terms of algorithmic implementations, *uniform-ideal* flows can however be more easily implemented, as they include for instance the translational or rotational flows. We then design and numerically benchmark such rotational flows in both isotropic and anisotropic systems, i.e. bidimensional systems of hard spheres, fundamental MCMC testbeds, and of hard dimers. In hard-disk systems, it appears that rotational flows retain good irreducibility property independently of the refreshment rate. In hard-dimer systems, the symmetry of the dimers itself interestingly imposes strong conditions on the possible deterministic flows. Introducing completely non-reversible rotations leads at the considered system size of 32 to a speed-up factor of up to 3 in the densest considered density ($\rho = 0.7$) compared to the state-of-the-art ECMC scheme [19].

We start by introducing in Section II the sampling of particle systems by a standard Metropolis algorithm. We then study in depth in Section III the ECMC method with standard translational flow, and such directly in a PDMP framework which allows to derive the necessary requirements as rotational invariance of the probability flows or the minimum event rate. We also illustrate the standard translational ECMC on isotropic particle sys-

tems. We then present in Section IV how to generalize such study to more general deterministic flows and discuss the design of rotational flows for hard disks and for anisotropic bidimensional dimers. Finally, we study the impact of the flow nature on the numerical performances in a system of bidimensional hard spheres and dimers in Section VI.

II. METROPOLIS SAMPLING FOR PARTICLE SYSTEMS

We consider a system composed of N isotropic particles in a d -dimensional periodic box of length L . A system configuration $\mathbf{x} \in \mathcal{S} = (\mathbb{R}/(L\mathbb{Z}))^{dN}$ is entirely described by all the particle positions x_i , i.e. $\mathbf{x} = (\mathbf{x}_i)_{1 \leq i \leq N}$. We introduce the notation $\mathbf{x}_{\setminus i} = ((\mathbf{x}_j)_{j < i}, (\mathbf{x}_j)_{j > i})$ for the configuration of all particle positions but the i -th one. The particles all interact in a repulsive pairwise manner following the potential,

$$U(\mathbf{x}) = \sum_{i < j} u(r(\mathbf{x}_i, \mathbf{x}_j)), \quad (1)$$

with r the periodic distance on the torus. Furthermore, the system may only admit a subset $\Omega \subset \mathcal{S}$ of valid configurations, as is the case when dealing with particles with hard cores of diameter σ that cannot overlap.

Noting β the inverse temperature, the system steady state then follows the Boltzmann measure,

$$\pi(\mathbf{x}) \propto \mathbb{1}_{\Omega}(\mathbf{x}) \exp(-\beta U(\mathbf{x})), \quad (2)$$

which we extend by continuity on $\partial\Omega$. In the following, we will set $\beta = 1$ without loss of generality. A typical Metropolis algorithm sampling from π consists in a proposal distribution q , verifying $\int_{\mathcal{S}} q(\mathbf{x}'|\mathbf{x}) d\mathbf{x}' = 1$ and commonly chosen to be a uniform increment of the position \mathbf{x}_i of a random single sphere i in \mathbf{x} , i.e.

$$q(\mathbf{x}'|\mathbf{x}) = \frac{1}{Nh^d} \sum_{i=1}^N \mathbb{1}_{\mathbf{x}_{\setminus i}}(\mathbf{x}'_{\setminus i}) \mathbb{1}_{C_1} \left(\frac{\mathbf{x}_i - \mathbf{x}'_i}{h} \right) \quad (3)$$

with $h \in]0, 1]$ some step amplitude and C_1 the centered unit hypercube of \mathbb{R}^d and, in an acceptance rate,

$$a(\mathbf{x}'|\mathbf{x}) = \min \left(1, \frac{\pi(\mathbf{x}')}{\pi(\mathbf{x})} \right). \quad (4)$$

The Markov chain generated by this Metropolis algorithm then follows the following kernel,

$$K(\mathbf{x}, \mathbf{x}') = q(\mathbf{x}'|\mathbf{x}) a(\mathbf{x}'|\mathbf{x}) + \left(1 - \int_{y \in \mathcal{S}} q(\mathbf{y}|\mathbf{x}) a(\mathbf{y}|\mathbf{x}) d\mathbf{y} \right) \mathbb{1}_{\{\mathbf{x}\}}(\mathbf{x}'), \quad (5)$$

which leaves π invariant, as it satisfies the detailed balance,

$$\pi(d\mathbf{x}') K(\mathbf{x}', d\mathbf{x}) = \pi(d\mathbf{x}) K(\mathbf{x}, d\mathbf{x}'). \quad (6)$$

It is however the global balance,

$$\int_{\mathbf{x}' \in \mathcal{S}} \pi(d\mathbf{x}') K(\mathbf{x}', d\mathbf{x}) = \int_{\mathbf{x}' \in \mathcal{S}} \pi(d\mathbf{x}) K(\mathbf{x}, d\mathbf{x}'), \quad (7)$$

which is the necessary condition for the invariance of π . Different methods have then been developed in order to gain efficiency by generating non-reversible processes which only obeys (7), in particular the Event-Chain Monte Carlo (ECMC) method [8, 9]. However, the generalization of such methods to any systems, for instance to the sampling of anisotropic particles requiring rotations, is not straightforward, whereas the Metropolis algorithm can be adapted quite easily by considering the correct potential U and by proposing steps that ensure irreducibility, as adding rotations for anisotropic particles. Aiming at devising such flow generalization for ECMC, we first present in the next section the standard ECMC method based on translations through their analytical characterization [20], which makes possible to precisely separate the necessary symmetries from the imposed sufficient ones, similarly to the difference between global and detailed balances. From there, we discuss how to introduce generalized flows and especially explain how to generate rotations.

III. EVENT-CHAIN MONTE CARLO WITH TRANSLATIONAL FLOW

Contrary to the reversible and discrete-time Metropolis scheme, Event-Chain Monte Carlo (ECMC) [9] generates a continuous-time and non-reversible Markov process. Considered as the infinitesimal limit of a discrete-time scheme, it then breaks detailed balance while still satisfying the global one. To do so, a lifting variable $\mathbf{v} \in \mathcal{V}$ following some chosen distribution μ is introduced so that the state space \mathcal{S} is extended to $\mathcal{S} \times \mathcal{V}$ and the extended state (\mathbf{x}, \mathbf{v}) now follows the product measure $\pi \otimes \mu$. The generated non-reversible process is then composed of ballistic updates of \mathbf{x} set by \mathbf{v} , which is updated by a Markov kernel at domain boundaries (e.g. when two hard cores collides in a particle system) or at the events of a Poisson process whose rate depends on the infinitesimal potential U increment. Thus, in particle systems, the generated process commonly comes down to updating a sphere position along some direction, until an event stemming from a pairwise interaction or hardcore collision with another sphere arises and this latter sphere is then the one being updated, in a billiard-like fashion as can be seen on Fig. 1. Then, on the extended state space $\mathcal{S} \times \mathcal{V}$, the former rejections in a Metropolis-scheme setting are now transformed into an update of \mathbf{v} , making the scheme rejection-free.

ECMC was initially built and justified through the infinitesimal limit of a discrete-time lifted Markov chain [9]. The process $(\mathbf{x}_t, \mathbf{v}_t)$ however identifies with a Piecewise deterministic Markov process [21, 22], which allows for a direct and efficient formalization directly at the

continuous-space and -time level through its infinitesimal generator and boundary conditions. In particular, the invariance condition is straightforward. Following the PDMP characterization of ECMC detailed in [20], we first explicit in terms of PDMP the ECMC process in a general setting. We then precisely characterize the necessary symmetries stemming from the invariance and the more restrictive but explicit possible sufficient symmetries. We finally illustrate it on the example of particle systems.

A. PDMP characterization

Generator and boundary kernel. First, the generated process is characterized in the bulk, i.e. for $(\mathbf{x}, \mathbf{v}) \in \Omega \times \mathcal{V}$, by its generator \mathcal{A} , coding for the infinitesimal changes through the process on any observable f (i.e. $\mathcal{A}f = \lim_{t \rightarrow 0} \frac{1}{t} E_{\mathbf{x}, \mathbf{v}} [f(\mathbf{x}(t), \mathbf{v}(t)) - f(\mathbf{x}, \mathbf{v})]$). In a general setting, the generator of the ECMC process, writes itself,

$$\mathcal{A} = \underbrace{\langle \phi((\mathbf{x}, \mathbf{v})), \nabla \rangle}_{\text{Transport}} + \underbrace{\lambda(\mathbf{x}, \mathbf{v})}_{\text{rate}} \underbrace{\left(\underbrace{Q((\mathbf{x}, \mathbf{v}), \cdot)}_{\text{jump}} - \text{Id} \right)}_{\text{Events}}, \quad (8)$$

where the deterministic flow ϕ is such that $(\dot{\mathbf{x}}(t), \dot{\mathbf{v}}(t)) = \phi(\mathbf{x}(t), \mathbf{v}(t)) = (\phi_x(\mathbf{x}(t), \mathbf{v}(t)), \phi_v(\mathbf{x}(t), \mathbf{v}(t)))$ during the ballistic phase. The rate $\lambda(\mathbf{x}, \mathbf{v})$ rules the event times, at which the flow is controlled through the update of lifting variable \mathbf{v} by the Markov kernel Q to another state $\mathbf{v}' \in \mathcal{V}$. First, we consider the common case where the flow $\phi(\mathbf{x}, \mathbf{v}) = \phi^T(\mathbf{x}, \mathbf{v}) = (\phi_x(\mathbf{v}), 0)$ identifies with a homogeneous translation of \mathbf{x} uniquely set by a fixed \mathbf{v} . We discuss the impact of more general transports on the invariance conditions in the next section.

In presence of hard validity constraints on Ω , as hardcore interactions, the ECMC characterization is complete once added a boundary Markov kernel Q_b , which updates a state $(\mathbf{x}, \mathbf{v}) \in \partial\Omega \times \mathcal{V}$ that is exiting the bulk ($\langle n(\mathbf{x}), \phi_x(\mathbf{v}) \rangle > 0$, $n(\mathbf{x})$ being the local normal on the boundary, such states form up the set Γ^+) to an entering one ($\langle n(\mathbf{x}), \phi_x(\mathbf{v}') \rangle < 0$, forming up the set Γ^-).

Refreshment and irreducibility. Irreducibility may require the introduction of a refreshment mechanism, which updates the lifting variable \mathbf{v} according to its invariant distribution. This refreshment can be introduced at an exponential homogeneous rate, then appearing directly in the generator, or as a variety of different schemes, including the commonly used fixed-time refreshment, as allowed by a formalization of the refreshment as a boundary effect [20]. As the refreshment term self-cancels in the invariance computation, we omit it in the following.




FIG. 1. Representation of a generalized flow in an event-chain sampling of hard spheres (**left**) and of different explicit uniform-ideal flows for spheres (**middle**), namely translations and rotations, and for dimers (**right**), namely bisector rotations and their hybrid variant combined with translations, Both the deterministic straight kernel and the stochastic direct kernel are illustrated.

Outputting samples. As the generated process is continuous-time, observables of interest must be integrated along the full trajectory $(\mathbf{x}_t, \mathbf{v}_t)$ or averaged over an unbiased collection of samples $(\mathbf{x}_{\tau_n}, \mathbf{v}_{\tau_n})_{\tau_n}$ from the full trajectory. As the refreshment process, τ_n can follow an exponential law or identify with some fixed value. In particular, if the vectors $\phi_x(\mathbf{v})$ are of same norm, it is equivalent to output samples after a fixed length has been traveled over, which is the method used in most works. If $\phi_x(\mathbf{v})$ can vary in norm, one should be careful to output samples at a fixed time or at a fixed length but renormalized by the norm of $\phi_x(\mathbf{v})$.

B. Necessary and sufficient symmetries for invariance

Invariance. The PDMP framework allows for a direct derivation of the invariance of the target π distribution,

$$\int_{\Omega \times \mathcal{V}} \mathcal{A}f(\mathbf{x}, \mathbf{v}) d\pi(\mathbf{x}) d\mu(\mathbf{v}) = 0, \quad (9)$$

for test functions f satisfying the boundary condition for $(\mathbf{x}, \mathbf{v}') \in \Gamma^-$,

$$\int_{\mathcal{V}} \mathbb{1}_{\Gamma^+}((\mathbf{x}, \mathbf{v})) Q_b((\mathbf{x}, \mathbf{v}), d\mathbf{v}') f(\mathbf{x}, \mathbf{v}) d\mu(\mathbf{v}) = f(\mathbf{x}, \mathbf{v}').$$

An integration by parts on the transport term shows the necessary compensation between the transport term and the event one, giving in the fixed translation case,

$$\begin{aligned} & (\langle \phi_x(\mathbf{v}), \nabla \ln \pi(\mathbf{x}) \rangle + \lambda(\mathbf{x}, \mathbf{v})) \mu(\mathbf{v}) \\ &= \int_{\mathcal{V}} \lambda(\mathbf{x}, \mathbf{v}') Q((\mathbf{x}, \mathbf{v}'), \mathbf{v}) d\mu(\mathbf{v}'), \end{aligned} \quad (10)$$

and the transport cancellation by the boundary kernel at

the boundary, for $(\mathbf{x}, \mathbf{v}) \in \partial\Omega \times \mathcal{V}$,

$$\begin{aligned} & \langle \phi_x(\mathbf{v}), n(\mathbf{x}) \rangle_- \mu(\mathbf{v}) \\ &= \int_{\mathcal{V}} Q_b((\mathbf{x}, \mathbf{v}'), \mathbf{v}) \langle \phi_x(\mathbf{v}'), n(\mathbf{x}) \rangle_+ d\mu(\mathbf{v}'). \end{aligned} \quad (11)$$

Similar to the global balance for discrete-time Markov chains, both conditions (10) and (11) are the only necessary and sufficient ones to ensure the correct invariance. As the firstly derived discrete-time MCMC schemes were detailed-balanced and obeyed stricter sufficient conditions, the invariance conditions in actual MCMC schemes based on PDMP actually satisfy more restrictive conditions. We now more precisely derive the necessary conditions stemming from (10) and (11).

Necessary symmetries. First, a direct necessary condition on the choice of μ and ϕ is the one of the conservation of the probability flows, obtained by an integration over \mathbf{v} in (10) or (11),

$$\begin{cases} \int_{\mathcal{V}} \langle \phi_x(\mathbf{v}), \nabla \ln \pi(\mathbf{x}) \rangle_+ d\mu(\mathbf{v}) = 0 \\ \int_{\mathcal{V}} \langle \phi_x(\mathbf{v}), n(\mathbf{x}) \rangle_+ d\mu(\mathbf{v}) = 0 \end{cases}, \quad (12)$$

i.e.,

$$\begin{cases} \int_{\mathcal{V}} \langle \phi_x(\mathbf{v}), \nabla \ln \pi(\mathbf{x}) \rangle_+ d\mu(\mathbf{v}) \\ \quad = \int_{\mathcal{V}} \langle \phi_x(\mathbf{v}), \nabla \ln \pi(\mathbf{x}) \rangle_- d\mu(\mathbf{v}) \\ \int_{\mathcal{V}} \langle \phi_x(\mathbf{v}), n(\mathbf{x}) \rangle_+ d\mu(\mathbf{v}) \\ \quad = \int_{\mathcal{V}} \langle \phi_x(\mathbf{v}), n(\mathbf{x}) \rangle_- d\mu(\mathbf{v}) \end{cases}, \quad (13)$$

Thus, the distribution and the flow set by the lifting variable \mathbf{v} must ensure that there is a balance between probability flows increasing or decreasing along $-\nabla \ln \pi$, i.e. the gradient of potential U , or along $n(\mathbf{x})$ on the boundary. This necessary flow symmetry (13) underlines that, when the time reversibility is broken, it is replaced by another key symmetry ensured by ϕ and μ and whose

origin is deeply rooted in the conservation of probability flow itself.

Furthermore, such symmetry plays a central role in the transport-event compensation. First, as the right-hand term of (10) is positive, the rate λ must obey the following condition for all $(\mathbf{x}, \mathbf{v}) \in \Omega \times \mathcal{V}$,

$$\lambda(\mathbf{x}, \mathbf{v}) + \langle \phi_x(\mathbf{v}), \nabla \ln \pi(\mathbf{x}) \rangle \geq 0, \quad (14)$$

so that the choice realizing the smallest event rate possible is,

$$\lambda_M(\mathbf{x}, \mathbf{v}) = \langle \phi_x(\mathbf{v}), -\nabla \ln \pi(\mathbf{x}) \rangle_+. \quad (15)$$

It corresponds to the rejection rate of the equivalent infinitesimal lifted Metropolis scheme. When constructing ECMC as such infinitesimal limit, this choice seems natural [9] but (14) further shows that it is indeed the minimal possible rate.

Considering now the conditions on the event Markov kernel Q , we obtain by setting λ to λ_M in (10),

$$\begin{aligned} \langle \phi_x(\mathbf{v}), -\nabla \ln \pi(\mathbf{x}) \rangle_- \mu(\mathbf{v}) = \\ \int_{\mathcal{V}} d\mathbf{v}' \mu(\mathbf{v}') \langle \phi_x(\mathbf{v}'), -\nabla \ln \pi(\mathbf{x}) \rangle_+ Q((\mathbf{x}, \mathbf{v}'), \mathbf{v}). \end{aligned} \quad (16)$$

Thus, the kernel Q can be set to any kernel which leave invariant, up to a *flip*, the distribution $\langle \phi_x(\mathbf{v}), -\nabla \ln \pi(\mathbf{x}) \rangle_+ \mu(d\mathbf{v})$. A valid explicit choice then is the direct sampling,

$$Q_{\text{Dir}}((\mathbf{x}, \mathbf{v}'), d\mathbf{v}) = \frac{\langle \phi_x(\mathbf{v}'), -\nabla \ln \pi(\mathbf{x}) \rangle_- \mu(d\mathbf{v})}{\int \langle \phi_x(\mathbf{v}'), -\nabla \ln \pi(\mathbf{x}) \rangle_+ \mu(d\mathbf{v})}, \quad (17)$$

which correctly sums up to 1 thanks to the necessary flow symmetry (13) and does not depend on the norm of the local gradient $\nabla \ln \pi(\mathbf{x})$. Such kernel can be understood as a direct pick from the *event* distribution $(\langle \phi_x(\mathbf{v}), -\nabla \ln \pi(\mathbf{x}) \rangle_+ \mu(d\mathbf{v}))$ combined with some *flip*. Such derivation is similar regarding the boundary condition (11) and Q^b can be chosen as Q once updated the reference vector from $-\nabla \ln \pi(\mathbf{x})$ to $n(\mathbf{x}, \mathbf{v})$.

This type of kernel was already introduced in the Forward event-chain generalization [14] by explicitly assuming that $\phi_x(\mathbf{v})\mu(\mathbf{v})$ is rotationally-invariant in order to derive (13). As the symmetry of probability flows around the gradient (13) actually is as necessary as the conservation condition itself, we show here that such direct sampling is always possible without further assumption. However, the advantage of imposing a rotationally-invariant property lies in an easier and general implementation of the kernel (17), as the dependence on \mathbf{x} only impacts the decomposition of \mathbf{v} into a parallel component along $\nabla \ln \pi(\mathbf{x})$, and an orthogonal one, whose values can be resampled independently from \mathbf{x} . More generally, restricting μ to some choice allowing to easily define the *flip* mentioned earlier allows some explicit construction we now address.

Sufficient symmetries. The need to derive explicit and general choices for λ , Q and Q_b has indeed led to the development of schemes satisfying stricter sufficient conditions but allowing easier numerical implementation, as the choice of the detailed balance has historically allowed for discrete-time Markov chains. Actually, the choice of imposing a product measure $\pi \otimes \mu$ is already a stricter choice than the necessary condition that π is the marginal of the extended measure, but it definitely simplifies the task of designing a correct scheme. We now discuss the most commonly used schemes.

First, we more generally decompose λ as,

$$\lambda(\mathbf{x}, \mathbf{v}) = \alpha(\mathbf{x}, \mathbf{v}) + \langle \phi_x(\mathbf{v}), -\nabla \ln \pi(\mathbf{x}) \rangle_+,$$

with $\alpha(\mathbf{x}, \mathbf{v}) \geq 0$ some excess rate. The condition (10) is now decomposed as,

$$\begin{aligned} \langle \phi_x(\mathbf{v}), -\nabla \ln \pi(\mathbf{x}) \rangle_- \mu(\mathbf{v}) = \\ \int_{\mathcal{V}} d\mathbf{v}' \mu(\mathbf{v}') \langle \phi_x(\mathbf{v}'), -\nabla \ln \pi(\mathbf{x}) \rangle_+ Q_\phi((\mathbf{x}, \mathbf{v}'), \mathbf{v}), \end{aligned} \quad (18)$$

and,

$$\alpha(\mathbf{x}, \mathbf{v})\mu(\mathbf{v}) = \int_{\mathcal{V}} d\mathbf{v}' \mu(\mathbf{v}') \alpha(\mathbf{x}, \mathbf{v}') Q_\alpha((\mathbf{x}, \mathbf{v}'), \mathbf{v}) \quad (19)$$

with the Markov kernels Q_ϕ and Q_α so that,

$$\begin{aligned} Q((\mathbf{x}, \mathbf{v}'), \mathbf{v}) = \frac{\langle \phi_x(\mathbf{v}'), -\nabla \ln \pi(\mathbf{x}) \rangle_+ Q_\phi((\mathbf{x}, \mathbf{v}'), \mathbf{v})}{\lambda(\mathbf{x}, \mathbf{v}')} \\ + \frac{\alpha(\mathbf{x}, \mathbf{v}')}{\lambda(\mathbf{x}, \mathbf{v}')} Q_\alpha((\mathbf{x}, \mathbf{v}'), \mathbf{v}). \end{aligned} \quad (20)$$

The kernel Q_ϕ can be set to any kernel obeying (16), as Q_{Dir} (17). The Markov kernel Q_α , stemming from the unnecessary excess rate α , can be any kernel which leaves $\alpha(\mathbf{x}, \mathbf{v})\mu(\mathbf{v})$ invariant, as a direct or Metropolis-like sampling following $\alpha(\mathbf{x}, \mathbf{v})\mu(\mathbf{v})$. Common choices for α are motivated by an easier computation of the event times, e.g. by the thinning of the actual minimal rate, $Q_\alpha((\mathbf{x}, \mathbf{v}), d\mathbf{v}')$ then identifying with $\delta(\mathbf{v} - \mathbf{v}')d\mathbf{v}'$ (no real event is actually sampled), or by the factorization of the minimal rate along the components of the gradient $-\nabla \ln \pi$, e.g. along interaction terms, which may allow for some inversion sampling of the event times.

Interestingly, in the case of factorized rates, writing,

$$\begin{aligned} \alpha(\mathbf{x}, \mathbf{v}) = \sum_i \langle f_i(\mathbf{x}, \mathbf{v}) \rangle_+ - \langle \sum_i f_i(\mathbf{x}, \mathbf{v}) \rangle_+ \\ = \sum_i \langle f_i(\mathbf{x}, \mathbf{v}) \rangle_- - \langle \sum_i f_i(\mathbf{x}, \mathbf{v}) \rangle_-, \end{aligned}$$

with $\{f_i\}$ the factors so that

$$\sum_i f_i(\mathbf{x}, \mathbf{v}) = \langle \phi_x(\mathbf{v}'), -\nabla \ln \pi(\mathbf{x}) \rangle,$$

we obtain by setting Q_ϕ and Q_α to the same kernel Q_{Fact} , as usually done, the condition,

$$\sum_i \langle f_i(\mathbf{x}, \mathbf{v}) \rangle_- \mu(\mathbf{v}) = \int_{\mathcal{V}} d\mathbf{v}' \mu(\mathbf{v}') \sum_i \langle f_i(\mathbf{x}, \mathbf{v}') \rangle_+ Q_{\text{Fact}}((\mathbf{x}, \mathbf{v}'), \mathbf{v}), \quad (21)$$

which necessarily requires the factorized symmetry for flow conservation,

$$\sum_i \int_{\mathcal{V}} \mu(d\mathbf{v}) \langle f_i(\mathbf{x}, \mathbf{v}) \rangle_- = \sum_i \int_{\mathcal{V}} \mu(d\mathbf{v}) \langle f_i(\mathbf{x}, \mathbf{v}) \rangle_+, \quad (22)$$

which comes down to (13), so that it is always possible to set Q_{Fact} to the direct variant

$$Q_{\text{Fact,Dir}}^{\text{Glob}}((\mathbf{x}, \mathbf{v}'), d\mathbf{v}) = \frac{\sum_i \langle f_i(\mathbf{x}, \mathbf{v}) \rangle_- \mu(d\mathbf{v})}{\int \sum_i \langle f_i(\mathbf{x}, \mathbf{v}) \rangle_+ \mu(d\mathbf{v})}. \quad (23)$$

However, it may prove interesting to consider ϕ and μ so that the condition (22) is achieved in a more restrictive manner thanks to some detailed symmetry, as,

$$\forall i, \int_{\mathcal{V}} \mu(d\mathbf{v}) \langle f_i(\mathbf{x}, \mathbf{v}) \rangle_- = \int_{\mathcal{V}} \mu(d\mathbf{v}) \langle f_i(\mathbf{x}, \mathbf{v}) \rangle_+. \quad (24)$$

This symmetry allows to treat every factor independently and to consider the following kernel,

$$Q_{\text{Fact,Dir}}^{\text{Det}}((\mathbf{x}, \mathbf{v}'), d\mathbf{v}) = \frac{\sum_i \langle f_i(\mathbf{x}, \mathbf{v}) \rangle_+ \mu(d\mathbf{v}) Q_i}{\sum_i \int \langle f_i(\mathbf{x}, \mathbf{v}) \rangle_+ \mu(d\mathbf{v})} \quad (25)$$

with, $Q_i((\mathbf{x}, \mathbf{v}'), d\mathbf{v}) = \frac{\langle f_i(\mathbf{x}, \mathbf{v}) \rangle_- \mu(d\mathbf{v})}{\int \langle f_i(\mathbf{x}, \mathbf{v}) \rangle_+ \mu(d\mathbf{v})}$.

Drawing on the lines of more restrictive and detailed symmetry for the flow conservation as in (24) and of a particular role played by some *flip*, most explicit schemes, apart the direct kernel, are building on some *flip* mapping $F_x : \mathcal{V} \rightarrow \mathcal{V}$ present in flow ϕ and μ , as,

$$\begin{cases} \langle \phi_x(\mathbf{v}), -\nabla \ln \pi(\mathbf{x}) \rangle = -\langle \phi_x(F_x(\mathbf{v})), -\nabla \ln \pi(\mathbf{x}) \rangle \\ \mu(\mathbf{v}) = \mu(F_x^{-1}(\mathbf{v})) \end{cases} \quad (26)$$

Such symmetry is sufficient to meet the flow conservation condition (12) in a detailed manner. To get such mapping, μ is typically chosen uniform or Gaussian so that $\phi_x(\mathbf{v})$ is rotationally invariant and F_x typically codes at the level of $\phi_x(\mathbf{v})$ for a full flip, a reflection across the potential gradient $-\nabla \ln \pi$ or some other particular symmetry exploiting the ones of $\nabla \ln \pi$ as the pairwise mirror symmetry or translational invariance in particle systems. In particular, it is the existence and exploitation of such underlying symmetry which has allowed to design deterministic Markov kernels of the type $\delta(F_x(\mathbf{v}) - \mathbf{v}') d\mathbf{v}'$, in comparison to the direct kernels previously discussed.

Further generalizations. The derived necessary flow symmetry (12) can actually be alleviated when considering more general processes. From a general perspective, it is most of the times not possible to obtain without rejections direct samples from π , except up to some known symmetries of π as for instance done in overrelaxation moves in spin systems. Therefore, focus has been on developing Markov kernel Q only updating the lifting variable \mathbf{v} . We here consider more general kernels potentially proposing also updates of the physical \mathbf{x} variable.

First, the invariance condition then yields, in the bulk,

$$\begin{aligned} & (\langle \phi_x(\mathbf{v}), \nabla \ln \pi(\mathbf{x}) \rangle + \lambda(\mathbf{x}, \mathbf{v})) \mu(\mathbf{v}) \pi(\mathbf{x}) \\ &= \int_{\Omega \times \mathcal{V}} \lambda(\mathbf{x}', \mathbf{v}') Q((\mathbf{x}', \mathbf{v}'), (\mathbf{x}, \mathbf{v})) d\pi(\mathbf{x}') d\mu(\mathbf{v}'), \end{aligned} \quad (27)$$

and at the boundary,

$$\begin{aligned} & \langle \phi_x(\mathbf{v}), n(\mathbf{x}) \rangle_- \mu(\mathbf{v}) \pi(\mathbf{x}) \\ &= \int_{\partial\Omega \times \mathcal{V}} Q_b((\mathbf{x}', \mathbf{v}'), (\mathbf{x}, \mathbf{v})) \langle \phi_x(\mathbf{v}'), n(\mathbf{x}') \rangle_+ d\pi(\mathbf{x}') d\mu(\mathbf{v}'). \end{aligned} \quad (28)$$

So that the conservation of the probability flows actually require, by integrating over (\mathbf{x}, \mathbf{v}) either in (27) or (28),

$$\int_{\partial\Omega \times \mathcal{V}} \langle \phi_x(\mathbf{v}), n(\mathbf{x}) \rangle d\pi(\mathbf{x}) d\mu(\mathbf{v}) = 0. \quad (29)$$

This condition is a global one, compared to the local (12), and is non-restrictive in the case where $\partial\Omega = \emptyset$ or when π already admits some symmetry leading to, for any $\mathbf{e} \in \mathcal{S}$,

$$\int_{\partial\Omega} \langle \mathbf{e}, n(\mathbf{x}) \rangle d\pi(\mathbf{x}) = 0,$$

as is the case of the particle systems presented in the previous section, which presents the pairwise mirror symmetry ($\nabla_{x_i} r(\mathbf{x}_i, \mathbf{x}_j) = -\nabla_{x_j} r(\mathbf{x}_i, \mathbf{x}_j)$). Otherwise, the conservation condition imposes on the choice of the flow ϕ and distribution μ a global symmetry of the exit ($\langle \pi(\mathbf{x}) \mu(\mathbf{v}) \phi_x(\mathbf{v}), n(\mathbf{x}) \rangle_+$) and entering probability flow ($\langle \pi(\mathbf{x}) \mu(\mathbf{v}) \phi_x(\mathbf{v}), n(\mathbf{x}) \rangle_-$) along the boundary. Compared to reversible schemes, this condition also directly stems from the non-reversible continuous-time ballistic component of the process, which pushes it up to the boundary $\partial\Omega$.

Eventually, the conservation or equivalently symmetry condition (29) could furthermore be suppressed if the Markov kernel Q (resp. Q_b) was even more general and for instance allowed to propose jumps from the bulk (resp. the boundary) to the boundary (resp. the bulk).

C. Translational ECMC for isotropic particle systems

We now present commonly-used ECMC schemes in particle systems, as illustrated in Fig. 1. They are

characterized by:

Transport. The lifting variable \mathbf{v} identifies with a tuple (e, i) of a vector \mathbf{e} in \mathbb{R}^d and the label $i \in \llbracket 1, N \rrbracket$ of the updated sphere along said vector. Standard choices for μ is a product measure $\mu_e \otimes \mu_n$, with μ_e the uniform distribution over $[\mathbf{u}_1, \dots, \mathbf{u}_d]$ (moves along the canonical basis of \mathbb{R}^d , as done in [8, 9]) or over \mathbb{S}^{d-1} (\mathbf{e} is some unit vector, as done in [13, 14]) or a Normal distribution over \mathbb{R}^d (as used in [15]) and μ_n the uniform distribution over $\llbracket 1, N \rrbracket$.

The deterministic flow ϕ identifies with the translational flow ϕ^T which updates (\mathbf{x}, \mathbf{v}) by a translation of the i -th sphere along \mathbf{e} , i.e.

$$\phi(\mathbf{x}, (e, i)) = ((0, \dots, 0, \underbrace{\mathbf{e}}_{i\text{-th}}, 0, \dots, 0), (0, 0)). \quad (30)$$

The conservation condition (12) is met, in particular thanks to the pairwise symmetry in the case of μ_e the uniform distribution over the canonical basis.

Event rate. The rates are factorized along the pairwise interactions, i.e.,

$$\lambda(\mathbf{x}, \mathbf{v})Q((\mathbf{x}, \mathbf{v}), \cdot) = \sum_{j \neq i} \lambda_{ij}(\mathbf{x}, \mathbf{v})Q_j((\mathbf{x}, \mathbf{v}), \cdot), \quad (31)$$

with,

$$\begin{aligned} \lambda_{ij}(\mathbf{x}, (e, i)) &= \frac{\beta}{2} \langle \nabla u(r(\mathbf{x}_i, \mathbf{x}_j)), \mathbf{e} \rangle_+ \\ &= \frac{\beta}{2} |u'(r(\mathbf{x}_i, \mathbf{x}_j))| \langle \mathbf{n}_{ji}, \mathbf{e} \rangle_+, \end{aligned} \quad (32)$$

with the pairwise normalized gradient $\mathbf{n}_{ij}(\mathbf{x}) = 2\nabla_{\mathbf{x}_i} r(\mathbf{x}_i, \mathbf{x}_j) = -\mathbf{n}_{ji}(\mathbf{x})$. Such superposition allows to treat every interaction independently. It can also ease the sampling of the Poisson process, either by inversion sampling [9] or by thinning [23], as mentioned previously. The factorized flow conservation condition (24) is also met, making it possible to consider a kernel Q similar to (25).

Markov kernel. Thanks to the existence of a mapping $F_{\mathbf{x}}$ as in (26) consisting in exchanging particle labels, it is possible to design deterministic kernels for $(Q_j)_{1 \leq j \leq N}$, as the straight Q^S or reflection kernel Q^R [8, 9],

$$\begin{cases} Q_j^S((\mathbf{x}, (e, i)), (e', i')) = \mathbb{1}_{\{j\}}(i') \mathbb{1}_{\{e\}}(e') \\ Q_j^R((\mathbf{x}, (e, i)), (e', i')) = \mathbb{1}_{\{j\}}(i') \mathbb{1}_{\{R_x(e)\}}(e'), \end{cases} \quad (33)$$

with $R_x(e) = -e + 2\langle \mathbf{n}_{ij}(\mathbf{x}), e \rangle \mathbf{n}_{ij}(\mathbf{x})$. At an event involving the moving sphere i and fixed one j , such kernels updates the moving sphere from i to j , either by keeping the same direction \mathbf{e} or reflecting it.

Now, in the case of a rotationally-invariant $\mu_e(\cdot)$, as the uniform distribution over \mathbb{S}^{d-1} or some Normal distribution over \mathbb{R}^d , the direct kernel variant (25) is actually

explicit [14]. By expressing \mathbf{e} in spherical coordinates with $\mathbf{n}_{ij}(\mathbf{x})$ as a polar axis ($\langle \mathbf{n}_{ij}(\mathbf{x}), \mathbf{e} \rangle = r \cos \theta_1$),

$$\mu_e(d\mathbf{e}) \propto r^{d-1} \mu_{e,r}(dr) \sin^{d-2}(\theta_1) \dots \sin(\theta_{d-2}) d\theta_1 \dots d\theta_{d-1},$$

it yields the following event distribution,

$$\int \mathbb{1}_A(\mathbf{e}) \langle \mathbf{e}, \mathbf{n}_{ji} \rangle_+ \mu_e(d\mathbf{e}) \propto \int_0^1 b^{d-2} db \int_{\mathbb{R}^+} r^{d-1} \mu_{e,r}(dr) \mathbb{1}_A\left(\left(r\sqrt{1-b^2}\mathbf{n}_{ij}(\mathbf{x}) + rb \frac{\mathbf{e} - \langle \mathbf{n}_{ij}(\mathbf{x}), \mathbf{e} \rangle \mathbf{n}_{ij}(\mathbf{x})}{\|\mathbf{e} - \langle \mathbf{n}_{ij}(\mathbf{x}), \mathbf{e} \rangle\|}, i\right)\right).$$

Thus, in the case of the uniform distribution over the hypersphere, a pair (\mathbf{e}, i) where \mathbf{e} is uniformly oriented in the hyperplane $\text{Span}\{\mathbf{n}_{ij}(\mathbf{x})\}^\perp$ and admits $\arcsin(b)$ as a polar angle to $\mathbf{n}_{ji}(\mathbf{x})$, follows the event distribution. In the case of a Gaussian distribution, $\langle \mathbf{e}, \mathbf{n}_{ji}(\mathbf{x}) \rangle$ should follow a χ -2 law and the other components of \mathbf{e} the usual Gaussian one.

In the following, we will consider the uniform distribution over \mathcal{S}^{d-1} and the corresponding following direct kernel, sampling from the event distribution combined with the label flip,

$$\begin{aligned} Q_j^{\text{Dir}}((\mathbf{x}, (e, i)), A) &= \int_0^1 b^{d-2} db \\ &\times \mathbb{1}_A\left(\left(\sqrt{1-b^2}\mathbf{n}_{ji}(\mathbf{x}) + b \frac{\mathbf{e} - \langle \mathbf{n}_{ij}(\mathbf{x}), \mathbf{e} \rangle \mathbf{n}_{ij}(\mathbf{x})}{\|\mathbf{e} - \langle \mathbf{n}_{ij}(\mathbf{x}), \mathbf{e} \rangle\|}, j\right)\right), \end{aligned} \quad (34)$$

which transfers the ballistic move from the sphere i to the j -th one, keeps the same direction in the hyperplane orthogonal to $\mathbf{n}_{ij}(\mathbf{x})$ but, compared to the deterministic kernels, update the polar angle directly from its steady-state distribution, via a sampling of $b \in [0, 1]$, typically by inversion sampling, as,

$$b = \nu^{1/(d-1)}, \text{ with } \nu \sim \text{Unif}(0, 1).$$

Finally, out of completeness, in the case of the uniform distribution over the canonical basis, a direct sampling is enforced by,

$$\begin{aligned} Q_j^{\text{Can}}((\mathbf{x}, (e, i)), d\mathbf{e}') &= \mathbb{1}_{\{u_k\}_k}(\mathbf{e}') \\ &\times \frac{\langle \mathbf{n}_{ji}(\mathbf{x}), \mathbf{e}' \rangle_+}{\sum_{k=1}^d |\langle \mathbf{n}_{ij}(\mathbf{x}), \mathbf{u}_k \rangle|} d\mathbf{e}', \end{aligned} \quad (35)$$

which can be sampled by inversion sampling.

Boundary kernel. A boundary kernel Q_b is necessary to take into account the hardcore interactions. As detailed in [20] for particle system and more generally in the previous subsection, the boundary kernel can be chosen to identify with Q . Simply put, hardcore interactions can be understood as soft ones but with diverging $\lambda_{ij} \rightarrow \infty$.

IV. GENERALIZED FLOW IN ECMC

A. Necessary and sufficient symmetries for invariance

Since its first developments, the translational flow ϕ^T has always been the one implemented in ECMC sampling. However, the PDMP formalism allows for more general flow: if the ODE $\frac{d(\mathbf{x}_t, \mathbf{v}_t)}{dt} = \phi((\mathbf{x}_t, \mathbf{v}_t))$ with differentiable drift ϕ defines a càdlàg (in time) $(\phi_t)_{t \geq 0}$ deterministic flow satisfying the semigroup property ($\phi_{t+s} = \phi_t \circ \phi_s$), then the drift part of the generator (8) is well defined. It may however impact the conservation condition and the balance required for invariance between the transport and event/boundary contributions.

Invariance. Indeed, in the general case, the integration by part on the transport formally generates additional terms as the divergence term $\nabla \cdot \phi$ and derivative $\langle \phi, \nabla \mu(\mathbf{v}) \rangle$, see for example [16] and as we derive and extend the invariance condition to constrained domain as imposed by hardcore interactions below. In the following, with a slight abuse of notation when denoting $\nabla_v f(\mathbf{x}, \mathbf{v})$, the gradient operator concerns only the continuous part of \mathbf{v} , which is the only part of the lifted variable that could be subject to the flow $\phi(\mathbf{x}, \mathbf{v})$. However the flow ϕ may depend both on the continuous and discrete parts of the lifting variable \mathbf{v} , e.g. when \mathbf{v} denotes the velocity and label of the active particle. By integration by parts,

$$\begin{aligned} \int_{\Omega \times \mathcal{V}} \langle \phi(\mathbf{x}, \mathbf{v}), \nabla f(\mathbf{x}, \mathbf{v}) \rangle d\pi(\mathbf{x}) d\mu(\mathbf{v}) = \\ \int_{\partial(\Omega \times \mathcal{V})} d\pi(\mathbf{x}) d\mu(\mathbf{v}) f(\mathbf{x}, \mathbf{v}) \langle n(\mathbf{x}, \mathbf{v}), \phi(\mathbf{x}, \mathbf{v}) \rangle \\ - \int_{\Omega \times \mathcal{V}} d\pi(\mathbf{x}) d\mu(\mathbf{v}) f(\mathbf{x}, \mathbf{v}) \left[\langle \nabla, \phi(\mathbf{x}, \mathbf{v}) \rangle \right. \\ \left. + \langle \phi(\mathbf{x}, \mathbf{v}), \nabla \ln(\pi(\mathbf{x})\mu(\mathbf{v})) \rangle \right]. \end{aligned} \quad (36)$$

As $\phi(\mathbf{x}, \mathbf{v}) = (\phi_x(\mathbf{x}, \mathbf{v}), \phi_v(\mathbf{x}, \mathbf{v}))$, it leads to the invariance conditions, for $(\mathbf{x}, \mathbf{v}) \in \Omega \times \mathcal{V}$,

$$\begin{aligned} \left(\langle \nabla_x, \phi_x(\mathbf{x}, \mathbf{v}) \rangle + \langle \nabla_v, \phi_v(\mathbf{x}, \mathbf{v}) \rangle \right. \\ \left. + \langle \phi_x(\mathbf{x}, \mathbf{v}), \nabla \ln \pi(\mathbf{x}) \rangle + \langle \phi_v(\mathbf{x}, \mathbf{v}), \nabla \ln \mu(\mathbf{v}) \rangle \right. \\ \left. + \tilde{\lambda}(\mathbf{x}, \mathbf{v}) \right) \mu(\mathbf{v}) = \int_{\Omega \times \mathcal{V}} \tilde{\lambda}(\mathbf{x}, \mathbf{v}') \tilde{Q}((\mathbf{x}, \mathbf{v}'), \mathbf{v}) d\mu(\mathbf{v}') \end{aligned} \quad (37)$$

and, for $(\mathbf{x}, \mathbf{v}) \in \partial(\Omega \times \mathcal{V})$,

$$\begin{aligned} \langle n(\mathbf{x}, \mathbf{v}), \phi(\mathbf{x}, \mathbf{v}) \rangle - \mu(\mathbf{v}) = \\ \int_{\partial(\Omega \times \mathcal{V})} \tilde{Q}_b((\mathbf{x}, \mathbf{v}'), \mathbf{v}) \langle n(\mathbf{x}, \mathbf{v}'), \phi(\mathbf{x}, \mathbf{v}') \rangle_+ d\mu(\mathbf{v}'), \end{aligned} \quad (38)$$

Necessary symmetries. Thus, considering a more general flow ϕ than the translation ϕ^T a priori imposes different necessary conditions as the ones derived in Sec. III,

a priori leading to different rates $\tilde{\lambda}$ and event kernel \tilde{Q} from the ones derived in Section III B, in order to maintain the correct target distribution $\pi \otimes \mu$ invariant. The condition on the boundary (38) is however not impacted, e.g. the boundary kernel \tilde{Q}_b obey the condition (11) but for a generalized flow and can be set to a valid choice Q from the translational case.

First, looking at the conservation condition as obtained in (12), it is now updated to,

$$\begin{cases} \int_{\mathcal{V}} \left(\langle \nabla, \phi(\mathbf{x}, \mathbf{v}) \rangle + \langle \phi_x(\mathbf{x}, \mathbf{v}), \nabla \ln \pi(\mathbf{x}) \rangle \right. \\ \left. + \langle \phi_v(\mathbf{x}, \mathbf{v}), \nabla \ln \mu(\mathbf{v}) \rangle \right) d\mu(\mathbf{v}) = 0, \\ \int_{\mathcal{V}} \langle \phi(\mathbf{x}, \mathbf{v}), n(\mathbf{x}, \mathbf{v}) \rangle d\mu(\mathbf{v}) = 0 \end{cases} \quad (39)$$

and the positivity condition imposing the minimal value of the rates $\tilde{\lambda}$, as obtained in (14), is now updated to,

$$\begin{aligned} \tilde{\lambda}_M(\mathbf{x}, \mathbf{v}) = \left[\langle \phi_x(\mathbf{v}), -\nabla \ln \pi(\mathbf{x}) \rangle - \langle \nabla, \phi(\mathbf{x}, \mathbf{v}) \rangle \right. \\ \left. + \langle \phi_v(\mathbf{x}, \mathbf{v}), -\nabla \ln \mu(\mathbf{v}) \rangle \right]_+. \end{aligned} \quad (40)$$

As for the translational case, the necessary conditions are imposing the same structure for the choice of rates $\tilde{\lambda}$ and kernel \tilde{Q} , as detailed in (18), (19) and (20) but updated to the minimal-rate condition (40). As done for the translational case, we now discuss more restrictive but explicit classes of valid flows, which obey some further symmetries than the necessary ones.

Ideal flows. Interestingly, it is possible to decrease the event rate minimal value to 0, i.e. to design a flow such that no events are needed to target the correct invariance distribution. A null minimal rate is equivalent to,

$$\begin{aligned} \langle \nabla, \phi(\mathbf{x}, \mathbf{v}) \rangle \geq \\ \langle \phi_x(\mathbf{v}), -\nabla \ln \pi(\mathbf{x}) \rangle + \langle \phi_v(\mathbf{x}, \mathbf{v}), -\nabla \ln \mu(\mathbf{v}) \rangle. \end{aligned}$$

Once combined with the conservation condition (39), the inequality is necessarily tight, so that,

$$\begin{aligned} \langle \nabla, \phi(\mathbf{x}, \mathbf{v}) \rangle = \\ \langle \phi_x(\mathbf{v}), -\nabla \ln \pi(\mathbf{x}) \rangle + \langle \phi_v(\mathbf{x}, \mathbf{v}), -\nabla \ln \mu(\mathbf{v}) \rangle, \end{aligned} \quad (41)$$

which can be summarized as,

$$\langle \nabla, \phi \rangle = \langle \nabla \mathcal{H}, \phi \rangle. \quad (42)$$

with $\mathcal{H}(\mathbf{x}, \mathbf{v}) = -(\ln \pi(\mathbf{x}) + \ln \mu(\mathbf{v}))$, the Hamiltonian of the extended system. It is analogous to a continuity equation and the incompressibility nature of the probability current $e^{-\mathcal{H}} \phi$, leaving the Boltzmann distribution stationary. Any *ideal* flow then presents a compensation between an increase/decrease of the energy of the extended state along $\nabla \mathcal{H}$ and an expansion/compression of the infinitesimal volume element, in order to preserve the total probability to be in this volume element.

A particular case then is the one of ideal solenoidal flows which are volume-preserving on top of being

probability-preserving. They then obey the following condition,

$$\langle \nabla, \phi \rangle = 0 = \langle \nabla \mathcal{H}, \phi \rangle. \quad (43)$$

Naturally, setting ϕ as the Hamiltonian flow $\phi^H = (\nabla_v \mathcal{H}, -\nabla_x \mathcal{H})$ appears as a sufficient choice to ensure the correct Boltzmann invariant distribution with a null event rate, but its implementation is not tractable in almost all cases of interest. More generally, Hamiltonian flows belong to a larger class of solenoidal flows, which can be characterized as,

$$\phi(\mathbf{x}, \mathbf{v}) = \mathbf{A} \nabla \mathcal{H}, \quad (44)$$

with \mathbf{A} some skew-symmetric matrix.

Uniform-ideal flows. In the case of a non-ideal flow, sampling the Poisson process with minimal event rate (40) may imply more involved computations than in the translational case (14). Computations may be eased by considering different degree of superposition, e.g.,

$$\begin{aligned} (a) & \langle \nabla(\ln(\pi(\mathbf{x})\mu(\mathbf{v})) + \nabla, -\phi(\mathbf{x}, \mathbf{v}))_+ \\ (b) & \langle \nabla(\ln(\pi(\mathbf{x})\mu(\mathbf{v})), -\phi(\mathbf{x}, \mathbf{v}))_+ + \langle \nabla, -\phi(\mathbf{x}, \mathbf{v}) \rangle_+ \\ (c) & \langle \nabla \ln \pi(\mathbf{x}), -\phi(\mathbf{x}, \mathbf{v}) \rangle_+ + \langle \nabla \ln \mu(\mathbf{v}), -\phi(\mathbf{x}, \mathbf{v}) \rangle_+ \\ & + \langle \nabla, -\phi(\mathbf{x}, \mathbf{v}) \rangle_+ \\ (d) & \dots, \end{aligned} \quad (45)$$

and implementing a thinning procedure on the different terms.

While these different factorization schemes (45) may indeed ease the numerical implementations, it may be more advantageous to introduce a generalized flow to decrease the averaged number of events, down to the ideal case. However, any flow choice involving a precise knowledge of $\nabla \ln \pi$ will likely be intractable, except in some particular simple cases. Then, a trade-off lies in designing flows which do not increase the event rate past its physical component (14), which is null for instance in presence of only hardcore interactions. We then call such flows *uniform-ideal* and they must obey the following condition for any (\mathbf{x}, \mathbf{v}) ,

$$\langle \nabla, \phi(\mathbf{x}, \mathbf{v}) \rangle + \langle \nabla_v \ln \mu(\mathbf{v}), \phi_v(\mathbf{x}, \mathbf{v}) \rangle = 0. \quad (46)$$

It indeed identifies with the condition (41) for uniform π as in hardcore models, making such flow ideal in this case. Any ECMC with such generalized flow then comes down to the well known case of translational flow, as the conditions of flow conservation (39) and minimal event rate (40) now identify with the translational ones, respectively (12) and (14), so that one can respectively set $\tilde{\lambda}$ and \tilde{Q} to λ and Q as obtained in the translational case.

Similarly to the ideal case, solenoidal flows as,

$$\phi(\mathbf{x}, \mathbf{v}) = \mathbf{A}(\nabla_x f(\mathbf{x}), \nabla_v \ln \mu(\mathbf{v})), \quad (47)$$

with \mathbf{A} some skew-symmetric matrix and $f : \mathcal{S} \rightarrow \mathcal{S}$ some function obeys (46). It identifies with a rotation in the case of μ Gaussian and $f(\mathbf{x}) = \mathbf{x}^2$. In the case of a uniform distribution μ , any solenoidal flow is valid, which includes rotational flows.

Hybrid flows. The PDMP framework allows for a great flexibility. Thus, a direct generalization consists in implementing sequentially different flows, as a translational flow and a rotational one, at the price of an additional lifting variable indicating the considered flow (and corresponding event rate $\tilde{\lambda}$ and Markov kernel \tilde{Q} if necessary). This lifting variable, and hence the implemented flow, can then be resampled at the events or refreshments.

V. ECMC WITH ROTATIONAL FLOW

As rotational flows constitute a large class of valid uniform-ideal flows, we describe in the following their implementations in hard-sphere and hard-dimer systems. The generalization to soft interactions is straightforward, as rotational flows, being uniform-ideal, can directly be implemented in the presence of interactions or an external potential by replacing the null event rate λ to $\langle \phi(\mathbf{x}, \mathbf{v}), -\nabla \ln \pi(\mathbf{x}) \rangle_+$.

A. Isotropic particle systems: Sphere systems

Instead of the translational flow as described in Section III C parameterized by a lifting variable of the form $(\mathbf{e}, i) \in \mathbb{S}^1 \times \{1, \dots, N\}$, we now consider rotational ones, parameterized by the lifting variable $\mathbf{v} = (\mathbf{e}, l, i) \in \mathbb{S}^1 \times \mathbb{R} \times \{1, \dots, N\}$. As for translations, \mathbf{e} corresponds to the infinitesimal update applied to \mathbf{x}_i but which adds up to a rotation around a center located at a distance $|l|$ from \mathbf{x}_i . As rotational flows are uniform-ideal, we can use the choice for λ and Q_b described in Section III C. Therefore we set,

$$\left\{ \begin{array}{l} \phi_x(\mathbf{v}) = (0, \dots, \underbrace{\mathbf{e}}_{i\text{-th}}, \dots, 0) \\ \phi_v(\mathbf{x}, \mathbf{v}) = (\frac{1}{l} A \mathbf{e}, 0, 0) \\ \lambda(\mathbf{x}, \boldsymbol{\theta}, \mathbf{v}) = 0 \\ \mu(\mathbf{e}, l, i) = \frac{1}{2\pi} \mathbb{1}_{\mathbb{S}^1}(\mathbf{e}) \nu(l) \frac{1}{N} \mathbb{1}_{\{1, \dots, N\}}(i) \end{array} \right., \quad (48)$$

with the infinitesimal rotation matrix $A = \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix}$. Such flow is indeed uniform-ideal, as $\langle \nabla, \phi \rangle = 0$. It is noteworthy that there is no constraint on the distribution of the auxiliary variable l , setting the distance between the rotation center and \mathbf{x}_i but also the anticlockwise or clockwise nature of the rotation. Thus, one can set $\nu(l)$ so that the generated rotations are always clockwise and at a fixed distance $|l|$ or so that $\nu(l)$ is some Gaussian distribution. Such flexibility offers then many possibilities to optimize the flow to a given problem.

Regarding the kernel Q_b , as was previously mentioned, it can be set to a deterministic kernel (33) or to the direct one (34). As the flow ϕ_x is directly parameterized in terms of an infinitesimal vector \mathbf{e} , the implementation is straightforward as it comes down to resampling \mathbf{e} . Here again the choice of l does not have an impact.

Finally, we study the numerical performances of implementing non-reversible rotations in sphere systems in Section VI A.

B. Anisotropic particle systems: Dimer systems

The hard-dimer system, modeled closely after hard spheres, provides a simple setting to study the rich behavior and the phase transitions of anisotropic particles [24, 25]. Such particles then require some degree of rotation and we now devise a ECMC schemes generating them in a completely non-reversible manner.

Dimer configurations. We now consider a system of N hardcore dimers of radius σ in a periodic L -box. Their configurations are described by the following state variable, $(\mathbf{x}, \boldsymbol{\theta}) \in (\mathbb{R}/(L\mathbb{Z}))^{2N} \times (\mathbb{R}/(\pi\mathbb{Z}))^N$, where \mathbf{x} is the sequence of dimer centers and $\boldsymbol{\theta}$ the one of angles. We furthermore introduce the following notations for the position of the two monomers constituting one dimer,

$$\mathbf{x}_i^+ = \mathbf{x}_i + \sigma \boldsymbol{\delta}_{\parallel}(\theta_i), \mathbf{x}_i^- = \mathbf{x}_i - \sigma \boldsymbol{\delta}_{\parallel}(\theta_i).$$

with $\boldsymbol{\delta}_{\parallel}(\theta) = (\cos(\theta), \sin(\theta)) \in \mathbb{S}^1$ for $\theta \in [0, \pi]$. The valid configurations forming up the set Ω satisfy the following constraint for every pair of dimers,

$$d(\mathbf{x}_i^{\pm}, \mathbf{x}_j^{\pm}) > d_{\text{pair}} = 2\sigma, \quad (49)$$

Now, similarly to the sphere case, in order to devise an ECMC sampling, we first extend the configurations to $(\mathbf{x}, \boldsymbol{\theta}, \mathbf{v})$ with $\mathbf{v} = (\mathbf{e}_+, l_{\perp}, i) \in \mathbb{S}^1 \times \mathbb{R} \times \{1, \dots, N\}$ where i is the label of the updated dimer, \mathbf{e}_+ the update vector of the $+$ -monomer and l_{\perp} is the variable setting, respectively to \mathbf{e}_+ and $\mathbf{x}_i^+, \mathbf{x}_i^-$, the update vector \mathbf{e}_- of the $-$ -monomer, so that (49) is obeyed, i.e.,

$$\mathbf{e}^- = \mathbf{e}^+ + l_{\perp} \boldsymbol{\delta}_{\perp}(\theta_i), \text{ with } \boldsymbol{\delta}_{\perp}(\theta) = (-\sin \theta, \cos \theta). \quad (50)$$

It is naturally possible to devise a scheme directly at the dimer level, updating \mathbf{x}_i and θ_i . However, introducing the infinitesimal update of each monomer allows to address the dimer constraint more easily and to later directly use the kernels devised in the translational case in Section III C. Furthermore, contrary to sphere systems, the dimer constraint imposes a correlated update of each of its monomer and it leads to requirements on the choice of l_{\perp} . For instance, setting l_{\perp} to some non-zero fixed value cannot give a valid scheme, as each dimer can only rotate in a single manner, as enforced by (50), and flows are not conserved and (39) not satisfied, as there is no label flip symmetry. The scheme could be correct at the

cost of introducing both clockwise and anticlockwise rotations and proposing, at collision, to backtrack the rotation of the updated dimer. Therefore we now derive the required distribution of l in order to ensure a completely non-reversible process.

As it is necessary for ergodicity to rotate the dimers, the monomer update vectors \mathbf{e}_+ and \mathbf{e}_- should cover the whole hypersphere \mathbb{S}^1 and therefore we choose for \mathbf{e}_+ , and by symmetry for \mathbf{e}_- , the uniform distribution on \mathbb{S}_1 , another possible choice being a Gaussian distribution for instance. Doing so, it constrains the possible choice of l_{\perp} to 0 (i.e. translation) or to a fixed value depending on θ_i , which we obtain, parameterizing for the i -th dimer,

$$\mathbf{e}_+ = (\cos(\theta_i + \alpha), \sin(\theta_i + \alpha)), \alpha \in [0, 2\pi], \quad (51)$$

as,

$$\mathbf{e}_- = (\cos(\theta_i - \alpha), \sin(\theta_i - \alpha)), l_{\perp} = 2 \sin \alpha. \quad (52)$$

Thus, any normed choice $(\mathbf{e}_+, \mathbf{e}_-)$ corresponds to a rotation of the dimer around a center placed on the dimer bisector, hence the derivation of what we call *bisector rotational flow*, as now detailed and illustrated on Fig. 1.

Bisector rotational flow. Out of simplicity, we now set $\mathbf{v} = (\alpha, i) \in [0, 2\pi] \times \{1, \dots, N\}$, so that the ECMC sampling by bisector rotational flow is completely characterized by,

$$\left\{ \begin{array}{l} \phi_x(\boldsymbol{\theta}, \mathbf{v}) = (0, \dots, \underbrace{\cos \alpha \boldsymbol{\delta}_{\parallel}(\theta_i)}_{i\text{-th}}, \dots, 0) \\ \phi_{\theta}(\mathbf{v}) = (0, \dots, \underbrace{\sin \alpha / \sigma}_{i\text{-th}}, \dots, 0) \\ \phi_v(\mathbf{x}, \mathbf{v}) = 0 \\ \lambda(\mathbf{x}, \boldsymbol{\theta}, \mathbf{v}) = 0 \\ \mu(\alpha, i) = \frac{1}{2\pi} \mathbb{1}_{[0, 2\pi]}(\alpha) \frac{1}{N} \mathbb{1}_{\{1, \dots, N\}}(i) \end{array} \right., \quad (53)$$

and any choice of Q_b corresponding to a valid one for sphere systems, up to the mapping of α to \mathbf{e}_+ in (51) or \mathbf{e}_- in (52), depending on which monomer is involved in a collision, and then mapping back from the resampled \mathbf{e}'_{\pm} to α' .

From an algorithmic point of view and collision computations, it may however be easier to work with a parametrization directly based on the rotation center. We then set $\mathbf{v} = (l, \gamma, i) \in \mathbb{R} \times \{-1, 1\} \times \{1, \dots, N\}$ to fix the rotation center $(\mathbf{x}_i + l \boldsymbol{\delta}_{\perp}(\theta_i))$ and sign ($\gamma = +1, -1$

for respectively anticlockwise, clockwise), so that,

$$\left\{ \begin{array}{l} \phi_x(\boldsymbol{\theta}, \mathbf{v}) = (0, \dots, \underbrace{\frac{\gamma l}{\sqrt{l^2 + \sigma^2}} \delta_{\parallel}(\theta_i)}_{i\text{-th}}, \dots, 0) \\ \phi_{\theta}(\mathbf{v}) = (0, \dots, \underbrace{\frac{\gamma}{\sqrt{l^2 + \sigma^2}}}_{i\text{-th}}, \dots, 0) \\ \phi_v(\mathbf{x}, \mathbf{v}) = 0 \\ \lambda(\mathbf{x}, \boldsymbol{\theta}, \mathbf{v}) = 0 \\ \mu(l, \gamma, i) = \frac{1}{\pi} \frac{\sigma}{l^2 + \sigma^2} \frac{1}{2} \mathbb{1}_{\{-1, 1\}}(\gamma) \frac{1}{N} \mathbb{1}_{\{1, \dots, N\}}(i) \end{array} \right. , \quad (54)$$

leading to the following mapping to \mathbf{e}_{\pm} for the i -th dimer,

$$\left\{ \begin{array}{l} \mathbf{e}_+ = \frac{\gamma}{\sqrt{l^2 + \sigma^2}} (l \delta_{\parallel}(\theta_i) + \sigma \delta_{\perp}(\theta_i)) \\ \mathbf{e}_- = \frac{\gamma}{\sqrt{l^2 + \sigma^2}} (l \delta_{\parallel}(\theta_i) - \sigma \delta_{\perp}(\theta_i)) \end{array} \right. , \quad (55)$$

which can be used to implement Q_b , as previously mentioned. For both parametrization, checking that the flow is uniform-ideal ($\langle \nabla, \phi \rangle = 0$) is straightforward.

As can be seen from (54), the correct distribution of rotation center is not trivial and especially not uniform in space. Indeed, the distance l between the dimer and rotation centers actually follows a Cauchy distribution, so that high values for l are not rare and should be treated numerically with care. In particular, the situation $l = 0$ correspond to self-rotations, while the value $l \rightarrow \pm\infty$ codes for translations. Thus all possible single-dimer rotations with normed monomer velocities can indeed be generated.

We study the numerical performances of implementing non-reversible rotations in dimer systems in Section VI B.

VI. NUMERICAL EXPERIMENTS

A. Rotational flows in hard-sphere systems

We first study rotational flows for isotropic hard spheres, as presented in Section IV. We consider a system of N spheres of radius σ in a periodic box of size L at the vicinity of the liquid-hexatic phase transition [26], setting $\rho = N\pi\sigma^2/L^2 = 0.708$. In order to estimate the relaxation time, we consider the decorrelation of the a priori slowest observable, here being Ψ_6 the global orientational parameter [27]. It is defined as the average over each sphere j of the average of the angles of the bonds ϕ_{jk} with its n_j nearest neighboring particles, i.e.,

$$\Psi_6 = \frac{1}{N} \sum_{j=1}^N \frac{1}{n_j} \sum_{k=1}^{n_j} \exp(6i\phi_{jk}). \quad (56)$$

We compare two translational schemes, respectively along $+x, y$ with a straight kernel (StraightTXY) and along all directions in \mathbb{S}^1 with a direct kernel (DirectT), and two rotational schemes generating all clockwise rotations of radius $\ell = 4\sigma$, respectively with a straight kernel

(StraightR) and a direct one (DirectR). At fixed-time refreshments, all of the lifting variables are drawn again according to their invariant measure. As all schemes have constant velocity, it corresponds to a fixed refreshment distance d_{ref} .

First, it clearly appears on Fig. 2 that the StraightTXY scheme requires some fine-tuning in order to set d_{ref} to an optimal value, as already observed [8, 18]. On the contrary, all other schemes (DirectT, StraightR, and DirectR) show better decorrelation performance, as the refreshment time increases. This is a particularly interesting behavior for the StraightR scheme, as it becomes a purely deterministic process as $d_{\text{ref}} \rightarrow \infty$ but apparently still yields an ergodic and efficient scheme. Then, while the impact of the refreshment time greatly depends on the considered scheme, the scaling of the integrated autocorrelation time of Ψ_6 with the number of particles N appears similar for all four schemes, as shown in Fig. 3.

Finally, regarding the schemes generating rotational flows, we found that no gain in performance was achieved by introducing switches at events between clockwise and anti-clockwise rotations. Also, the choice of the rotation radius ℓ appears not critical, as can be seen on Fig. 4, where only the StraightR scheme shows some ℓ -dependence and a slightly better performance than all the other fine-tuned schemes for $\ell \sim L$.

B. Bisector rotational flows in hard-dimer systems

We now study the efficiency of rotational flows in the anisotropic case of the hard-dimer system, as described in section VB. The system consists in $N = 32$ dimers of monomer radius σ in a periodic box of size L . Building on the Monte Carlo simulations of [24] and molecular dynamics simulations of [25], simulations are performed at densities $\rho = 2N\pi\sigma^2/L^2 = 0.5$ and $\rho = 0.7$, the latter being slightly under the observed phase transition. As commonly studied in isotropic-nematic transition [28], we now consider the relaxation of S_2 , the scalar parameter for the two-dimensional nematic order,

$$S_2 = \frac{1}{N} \sum_{j=1}^N (2 \cos^2(\theta_j - \bar{\theta}) - 1) \in [0, 1], \quad (57)$$

where $\bar{\theta}$ is the angle of the director, defined modulo π . This observable has an intrinsic dimer-flip symmetry $\theta_j \mapsto \theta_j + \pi$. Other observables, with or without the dimer-flip symmetry, can be constructed with the θ_j 's. They are presented in Appendix A and show a consistent behavior. In [25], the Ψ_6 of the equivalent sphere configuration is investigated, but we found that the global order of the θ_i 's shows longer correlation times at the considered densities.

We consider two different algorithms: The Rbisector scheme with a bisector rotational flow, as introduced in Sec. VB, and the Hybrid one with a flow switching at event between translational and bisector rota-


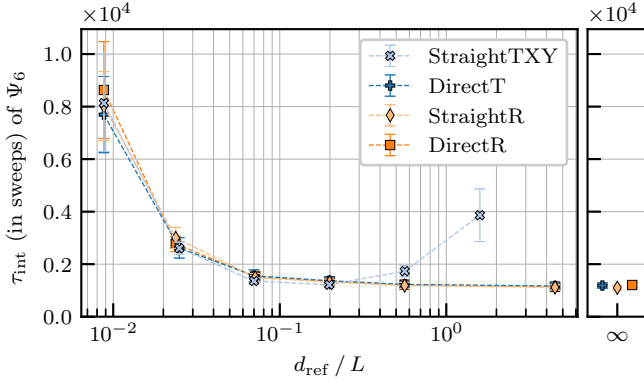


FIG. 2. Ψ_6 integrated autocorrelation time as a function of the refreshment distance d_{ref} , for $N = 64$ (left) and $N = 256$ (right) hard spheres at density $\rho = 0.708$. Rotations have a radius $\ell = 4\sigma$. One sweep corresponds to N event computations and L is the size of the system.




FIG. 3. Ψ_6 integrated autocorrelation time as a function of the number of particles N , for hard spheres at density $\rho = 0.708$. The StraightTXY has optimal refreshment, the other schemes have no refreshment. Rotations have a radius $\ell = 4\sigma$. One sweep corresponds to N event computations.

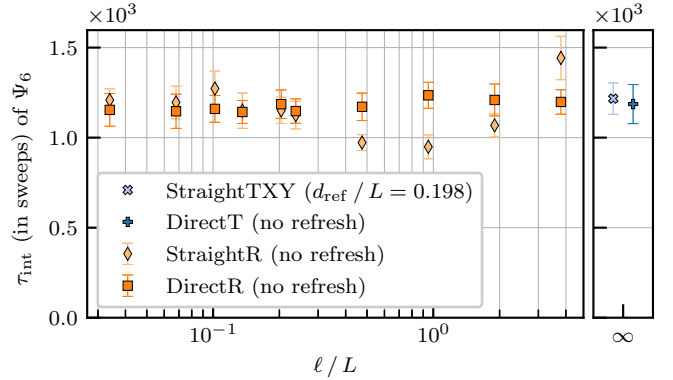


FIG. 4. Ψ_6 integrated autocorrelation time as a function of the radius ℓ , for $N = 64$ spheres at density $\rho = 0.708$. The fine-tuned versions of the translational schemes are shown for reference. One sweep corresponds to N event computations and L is the size of the system.

tional flows. Both are studied combined either with the straight and direct kernels. We compare them to the ECMC+Met algorithm introduced in [19] consisting in a ECMC scheme combined with Metropolis rotation proposals. In more details, the ECMC flow consists in translational moves, with all directions allowed. At each refreshment, a random number n of dimer rotations are also proposed. Their number n follows a geometric distribution of parameter p_{chain} and the proposed rotation angle increments are uniformly picked over $[-\Delta, \Delta]$, Δ being tuned at $\Delta = \pi$ for $\rho = 0.5$ and $\Delta = 0.4$ for $\rho = 0.7$. Following [19], p_{chain} is taken so that on average there is as many ECMC events as there are Metropolis proposals and both kind counts as one event computation. In addition to the straight kernel used in [19], we also simulate the ECMC+Met scheme with a direct kernel. At fixed-time refreshments, all lifting variables are fully drawn again according to their invariant measure. As all schemes have constant velocity, it corresponds to a fixed refreshment distance d_{ref} .

As shown in Fig. 5, the autocorrelation function for




FIG. 5. Normalized autocorrelation of the parameter S_2 , for 20 runs and their averaged autocorrelation (black). The system consists of $N = 32$ dimers at density $\rho = 0.7$. The Hybrid and Rbisector schemes with the direct kernel and no refreshment are used here. One sweep corresponds to N event computations.




FIG. 6. Nematic scalar order parameter $S_2 = \langle 2 \cos^2(\theta_j - \bar{\theta}) - 1 \rangle$ integrated autocorrelation time as a function of the refreshment distance d_{ref} , for $N = 32$ hard dimers at density $\rho = 0.5$ (**left**) and $\rho = 0.7$ (**right**). One sweep corresponds to N event computations and L is the size of the system.

S_2 appears close to an exponential decay, and in particular does not exhibit different time scales. Therefore, the integrated autocorrelation time is a well-suited measure of the decorrelation and is displayed in Fig. 6. For the dilute case $\rho = 0.5$, the ECMC+Met scheme shows an optimal value at finite d_{ref} whereas the Rbisector and Hybrid schemes reach an optimum for vanishing refreshment $d_{\text{ref}} \rightarrow \infty$. The direct schemes outperform their straight counterparts and the fastest decorrelation is achieved with the Rbisector algorithm (speed-up of ~ 3.0 compared to ECMC+Met, the Hybrid achieves a speed-up of ~ 2.1). Similarly to the rotations in sphere systems, the Rbisector and Hybrid schemes with straight kernel turn into purely deterministic processes as $d_{\text{ref}} \rightarrow \infty$. At the denser density of $\rho = 0.7$, similar observations can be drawn, regarding the direct versions of the algorithms better performing and the d_{ref} -tuning requirement only for the ECMC+Met algorithms. However, the Rbisector scheme now performs worse than the ECMC+Met scheme (speed-up of ~ 0.74) and the Hybrid algorithm shows the fastest decay (speed-up of ~ 2.6). Such behavior could be explained from the observed collision loops in this denser regime, where a rotating dimer i will collide with another dimer j , which will collide in turn with i and so on. There, introducing a flow switch between rotations and translations seems to control the appearance of such loops.

VII. CONCLUSION

Since the introduction of the ECMC method, important algorithmic and analytical efforts have been made towards the generalization of such schemes to any systems. Drawing on that line, the PDMP characterization of ECMC schemes allows us to derive the necessary fundamental symmetries of the flow conservation and minimum event rate, as required for invariance of the correct distribution π . It then makes clearly appear the more restrictive but sufficient symmetries most commonly im-

posed in order to derive explicit schemes. From there, the introduction of generalized flow is carried on in the same manner, by a careful characterization of necessary and sufficient symmetries. In particular, we define two classes of flows of interest: the ideal one, where events are suppressed, and the uniform-ideal ones, for which one can use the valid choices of rates and kernels, as developed and known for the usual translational flow.

In particular, the uniform-ideal class includes rotational flows and we devise ECMC schemes generating completely non-reversible rotations in bidimensional hard-sphere and hard-dimer systems. While the rotations can be simulated in many different ways in hard spheres, the contact constraint of the two monomers of each single dimer imposed strict requirements on the possible rotations, which reduce to the single bisector rotational flow in the case of normed monomer velocities. Importantly, the numerical simulations indicate, at least for meaningful observables and high densities, that for rotational flows the refreshment mechanism does not require some crucial fine tuning and may actually not be necessary for ergodicity, even for schemes then completely deterministic. This is reminiscent of numerical observations of similar behavior for translational ECMC with a reflection kernel in hard-sphere systems [8]. Proving such property constitutes a major mathematics challenge, as the question of proving ergodicity of ECMC schemes with refreshment for hard-particle systems is already not trivial [20] and an ergodicity proof with no refreshment, but with some conditions on the potential, is only known for the so-called Zig-Zag process [29]. In lack of such proof, a refreshment distance much larger than L could however be chosen.

Such possibility of achieving ergodicity without any refreshment would indicate how the chaotic nature of the considered system itself can be harnessed into producing efficient and robust non-reversible schemes. Thus, comparing the performances observed for the hard-sphere system and hard-dimer one, the latter presenting an additional nematic order, we already can observe how the

intrinsic stochastic nature of the system impacts the optimization of the algorithm scheme, as it appears that hard-sphere system, apart from the robustness to the refreshment fine tuning, does not particularly benefit from more involved schemes here close to the transition point. However, one can note that the StraightR scheme at $\ell \simeq L$ is better performing, while actually generating some very gradual sequential update of the direction. This behavior is reminiscent of sequential translational straight schemes where the direction is incremented by a fixed value at refreshments, for example in tethered dimer systems [18, 30]. On the other hand, in the hard-dimer case, speed up are observed, even for the small considered system size and at the dilute and denser densities. Interestingly, at the denser density, the Hybrid scheme allows to control the collision loops, which arises from the strong orientational correlation and may slow down the process when considering only bisector rotations. Thus, it appears that developing efficient non-reversible algorithms require a fine understanding of the correlations present in the system. An interesting prospect is to consider the implementations of rotational flows to other anisotropic-particle systems and to study the constraint on such flows as imposed by the particle shape itself and the impact of such anisotropy and the consequential correlations on the algorithm efficiency.

Finally, further work should also deal with the efficient numerical implementation of rotational flows, with in addition soft interactions. The challenge here lies in the simulations of the event times, which should computationally benefit from a thinning procedure. More generally, the question of the parallelization of such rotational flows in an efficient manner is an important one as to the

applications of such methods to large-scale simulations of molecular systems, all the more as recent parallelizing efforts have been made concerning the standard translational ECMC [31]. The parallelization of rotational flows could also allow to study the scaling and evolution of the observed speed-up in hard-dimer for large system sizes, as well as characterize more precisely their critical behavior.

ACKNOWLEDGMENTS

All the authors are grateful for the support of the French ANR under the grant ANR-20-CE46-0007 (*SuSa* project). A.G. is supported by the Institut Universitaire de France. Computations have been performed on the supercomputer facilities of the Mésocentre Clermont Auvergne.

CONFLICT OF INTEREST STATEMENT

The authors have no conflicts to disclose.

DATA AVAILABILITY STATEMENT

The data that support the findings of this study are available from the corresponding author upon reasonable request. The codes used to generate the data are available at <https://bitbucket.org/guyont/ecmc-rotational-flow/src/main/>.

-
- [1] N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller, Equation of state calculations by fast computing machines, *The journal of chemical physics* **21**, 1087 (1953).
 - [2] D. Frenkel and B. Smit, *Understanding molecular simulation: from algorithms to applications*, Vol. 1 (Elsevier, 2001).
 - [3] W. Janke, Statistical analysis of simulations: Data correlations and error estimation, *Quantum Simulations of Complex Many-Body Systems: From Theory to Algorithms* **10**, 423 (2002).
 - [4] D. A. Levin and Y. Peres, *Markov chains and mixing times*, Vol. 107 (American Mathematical Soc., 2017).
 - [5] P. C. Hohenberg and B. I. Halperin, Theory of dynamic critical phenomena, *Rev. Mod. Phys.* **49**, 435 (1977).
 - [6] R. H. Swendsen and J.-S. Wang, Nonuniversal critical dynamics in monte carlo simulations, *Physical review letters* **58**, 86 (1987).
 - [7] U. Wolff, Collective monte carlo updating for spin systems, *Physical Review Letters* **62**, 361 (1989).
 - [8] E. P. Bernard, W. Krauth, and D. B. Wilson, Event-chain monte carlo algorithms for hard-sphere systems, *Phys. Rev. E* **80**, 056704 (2009).
 - [9] M. Michel, S. C. Kapfer, and W. Krauth, Generalized event-chain monte carlo: Constructing rejection-free global-balance algorithms from infinitesimal steps, *The Journal of Chemical Physics* **140**, 054116 (2014), <https://doi.org/10.1063/1.4863991>.
 - [10] T. A. Kampmann, H.-H. Boltz, and J. Kierfeld, Monte carlo simulation of dense polymer melts using event chain algorithms, *The Journal of chemical physics* **143**, 044105 (2015).
 - [11] M. Michel, J. Mayer, and W. Krauth, Event-chain monte carlo for classical continuous spin models, *EPL (Europhysics Letters)* **112**, 20003 (2015).
 - [12] Y. Nishikawa, M. Michel, W. Krauth, and K. Hukushima, Event-chain algorithm for the heisenberg model: Evidence for $z \simeq 1$ dynamic scaling, *Physical Review E* **92**, 063306 (2015).
 - [13] J. Harland, M. Michel, T. A. Kampmann, and J. Kierfeld, Event-chain monte carlo algorithms for three- and many-particle interactions, *EPL (Europhysics Letters)* **117**, 30001 (2017).
 - [14] M. Michel, A. Durmus, and S. Sénécal, Forward event-chain monte carlo: Fast sampling by randomness control in irreversible markov chains, *Journal of Com-*

- putational and Graphical Statistics **29**, 689 (2020), <https://doi.org/10.1080/10618600.2020.1750417>.
- [15] M. Klement and M. Engel, Efficient equilibration of hard spheres with newtonian event chains, *The Journal of chemical physics* **150**, 174108 (2019).
- [16] P. Vanetti, A. Bouchard-Côté, G. Deligiannidis, and A. Doucet, Piecewise-deterministic markov chain monte carlo, *arXiv preprint arXiv:1707.05296* (2017).
- [17] J. Bierkens, S. Grazi, K. Kamatani, and G. Roberts, The boomerang sampler, in *Proceedings of the 37th International Conference on Machine Learning*, Proceedings of Machine Learning Research, Vol. 119, edited by H. D. III and A. Singh (PMLR, 2020) pp. 908–918.
- [18] P. Höllmer, A. Maggs, and W. Krauth, Hard-disk dipoles and non-reversible markov chains, *The Journal of Chemical Physics* **156**, 084108 (2022).
- [19] M. Klement, S. Lee, J. A. Anderson, and M. Engel, Newtonian event-chain monte carlo and collision prediction with polyhedral particles, *Journal of Chemical Theory and Computation* **17**, 4686 (2021).
- [20] A. Monemvassitis, A. Guillin, and M. Michel, PDMP characterisation of event-chain monte carlo algorithms for particle systems, *Journal of Statistical Physics* **190**, 66 (2023).
- [21] M. H. Davis, Piecewise-deterministic markov processes: a general class of non-diffusion stochastic models, *Journal of the Royal Statistical Society: Series B (Methodological)* **46**, 353 (1984).
- [22] M. H. A. Davis, *Markov models and optimization*, Monographs on Statistics and Applied Probability, Vol. 49 (Chapman & Hall, London, 1993) pp. xiv+295.
- [23] S. C. Kapfer and W. Krauth, Two-dimensional melting: From liquid-hexatic coexistence to continuous transitions, *Phys. Rev. Lett.* **114**, 035702 (2015).
- [24] K. W. Wojciechowski, A. C. Brańka, and D. Frenkel, Monte Carlo simulations of a two-dimensional hard dimer system, *Physica A: Statistical Mechanics and its Applications* **196**, 519 (1993).
- [25] L. F. Cugliandolo, P. Digregorio, G. Gonnella, and A. Suma, Phase coexistence in two-dimensional passive and active dumbbell systems, *Phys. Rev. Lett.* **119**, 268002 (2017).
- [26] E. P. Bernard and W. Krauth, Two-step melting in two dimensions: first-order liquid-hexatic transition, *Physical review letters* **107**, 155704 (2011).
- [27] K. J. Strandburg, Two-dimensional melting, *Rev. Mod. Phys.* **60**, 161 (1988).
- [28] R. Eppenga and D. Frenkel, Monte Carlo study of the isotropic and nematic phases of infinitely thin hard platelets, *Molecular Physics* **52**, 1303 (1984).
- [29] J. Bierkens, G. O. Roberts, and P.-A. Zitt, Ergodicity of the zigzag process, *Ann. Appl. Probab.* **29**, 2266 (2019).
- [30] L. Qin, P. Höllmer, and W. Krauth, Direction-sweep markov chains, *Journal of Physics A: Mathematical and Theoretical* **55**, 105003 (2022).
- [31] B. Li, S. Todo, A. C. Maggs, and W. Krauth, Multi-threaded event-chain monte carlo with local times, *Computer Physics Communications* **261**, 107702 (2021).

Supplementary Materials

Necessary and sufficient symmetries in Event-Chain Monte Carlo with generalized flows and Application to hard dimers

Appendix A: Additional dimer observables

The dimer angles θ_j can be seen twofold, either with π - or 2π -periodicity. If opposite directions can be discriminated, the global orientational order is encoded in the average of the directions $\delta_{\parallel}(\theta_j) = (\cos \theta_j, \sin \theta_j)$ of all dimers, namely the polarization,

$$\mathbf{p} = \frac{1}{N} \sum_{j=1}^N \delta_{\parallel}(\theta_j). \quad (\text{A1})$$

If the system has an intrinsic dimer-flip symmetry $\theta_j \mapsto \theta_j + \pi$, the relevant observable stems from the study of the isotropic-nematic phase transition [28] in three dimensions. The general matrix order parameter of such a transition is reduced in dimension 2 to a complex number,

$$z_2 = \frac{1}{N} \sum_{j=1}^N \exp(2i\theta_j) = S_2 \exp(2i\bar{\theta}), \quad (\text{A2})$$

where $\bar{\theta}$ is the angle of the director, defined modulo π , and S_2 is the scalar parameter for the two-dimensional nematic order,

$$S_2 = \frac{1}{N} \sum_{j=1}^N (2 \cos^2(\theta_j - \bar{\theta}) - 1) \in [0, 1]. \quad (\text{A3})$$

In the following, z_2 is called the nematic vector. Its expression is reminiscent of the Ψ_6 for hard spheres and plays the role of averaging in a π -periodic setting. The square box imposes the ensemble averages $\langle z_2 \rangle = 0$ and $\langle \mathbf{p} \rangle = 0$, which give a rudimentary check for ergodicity, respectively with and without dimer-flip symmetry.

Experiments (Fig. 7,8) show the polarization is slower than the nematic vector, but the need to flip all dimers is an artificial constrain in the case where opposite directions are equivalent. The scalar S_2 decorrelates faster than z_2 but the behavior of the algorithms is the same for all observables.

At low density, the ECMC+Met scheme has an artificial speed-up on the decorrelation of \mathbf{p} because the proposals are uniform on all angles and the $\theta_j \rightarrow \theta_j + \pi$ transform is always accepted.




FIG. 7. Nematic vector z_2 (left) and polarization \mathbf{p} (right) integrated autocorrelation time as a function of the refreshment distance d_{ref} , for $N = 32$ hard dimers at density $\rho = 0.5$. One sweep corresponds to N event computations and L is the size of the system.

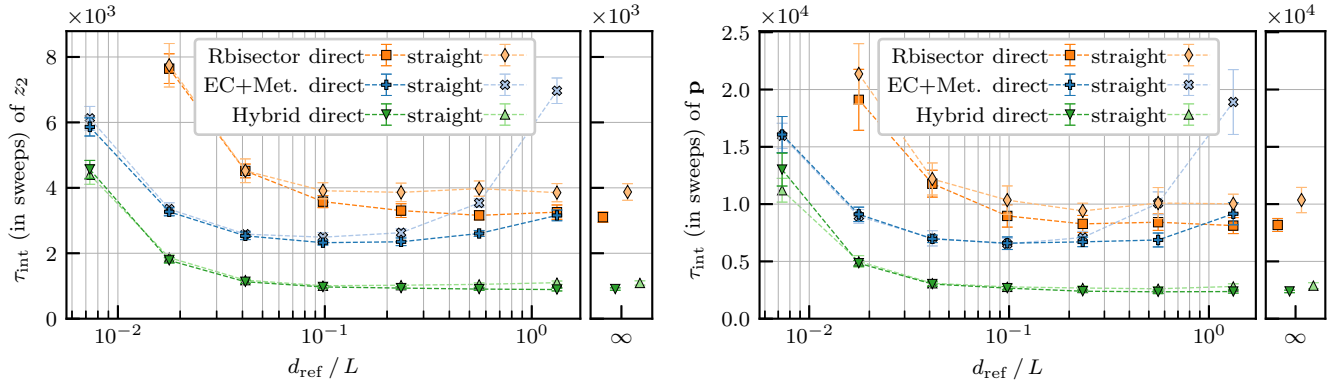


FIG. 8. Nematic vector z_2 (left) and polarization \mathbf{p} (right) integrated autocorrelation time as a function of the refreshment distance d_{ref} , for $N = 32$ hard dimers at density $\rho = 0.7$. One sweep corresponds to N event computations and L is the size of the system.