# Mixing and perfect sampling in one-dimensional particle systems 

Ze Lei ${ }^{1,2}{ }^{(a)}$ and Werner $\mathrm{Krauth}^{1,2}$ (b)<br>${ }^{1}$ Laboratoire de Physique Statistique, Département de physique de l'ENS, Ecole Normale Supérieure, PSL Research University, Université Paris Diderot, Sorbonne Paris Cité, Sorbonne Universités, UPMC Univ. Paris 06, CNRS, 75005 Paris, France<br>${ }^{2}$ Max-Planck-Institut für Physik komplexer Systeme, Nöthnitzer Str. 38, 01187 Dresden, Germany

PACS 02.70.Tt - Justifications or modifications of Monte Carlo methods
PACS 02.50.Ng - Distribution theory and Monte Carlo studies


#### Abstract

We study the approach to equilibrium of the event-chain Monte Carlo (ECMC) algorithm for the one-dimensional hard-sphere model. Using the connection to the coupon-collector problem, we prove that a specific version of this local irreversible Markov chain realizes perfect sampling in $\mathcal{O}\left(N^{2} \log N\right)$ events, whereas the reversible local Metropolis algorithm requires $\mathcal{O}\left(N^{3} \log N\right)$ time steps for mixing. This confirms a special case of an earlier conjecture about $\mathcal{O}\left(N^{2} \log N\right)$ scaling of mixing times of ECMC and of the forward Metropolis algorithm, its discretized variant. We furthermore prove that sequential ECMC (with swaps) realizes perfect sampling in $\mathcal{O}\left(N^{2}\right)$ events. Numerical simulations indicate a cross-over towards $\mathcal{O}\left(N^{2} \log N\right)$ mixing for the sequential forward swap Metropolis algorithm, that we introduce here. We point out open mathematical questions and possible applications of our findings to higher-dimensional statisticalphysics models.


Sampling, mixing, perfect sampling, stopping rules. - Ever since the 1950s 11, Markov-chain Monte Carlo (MCMC) methods have ranked among the most versatile approaches in scientific computing. Monte Carlo algorithms strive to sample a probability distribution $\pi$. For an $N$-particle system in statistical mechanics, with particle $i \in\{1, \ldots, N\}$ described by coordinates $x_{i}$, sampling $\pi$ corresponds to generating configurations $\mathrm{x}=$ $\left\{x_{1}, \ldots, x_{N}\right\}$ distributed with the Boltzmann probability $\pi(\mathbf{x}) \propto \exp [-\beta E(\mathbf{x})]$, where $E$ is the system energy and $\beta$ the inverse temperature. For problems where $\mathbf{x}$ lies in a high-dimensional discrete or continuous space $\Omega$, this sampling problem can usually not be solved directly [2,3].

MCMC consists instead in sampling a probability distribution $\pi^{(t)}$ that evolves with a time $t$. The initial probability distribution, at time $t=0, \pi^{(t=0)}$, can be sampled directly. Often, it is simply composed of a single configuration, so that $\pi^{(0)}$ is a $\delta$-function on an explicitly given configuration $\left\{x_{1}, \ldots, x_{N}\right\}$. In the limit $t \rightarrow \infty$, the distribution $\pi^{(t)}$ evolves from the initial one towards the target probability distribution $\pi=\lim _{t \rightarrow \infty} \pi^{(t)}$. Besides the development of MCMC algorithms that approach the limit

[^0]distribution $\pi$ as quickly as possible for any initial distribution $\pi^{(0)}$, a key challenge in MCMC consists in estimating the time scale on which the time-dependent distribution $\pi^{(t)}$, which depends on $\pi^{(0)}$, is sufficiently close to $\pi$ that the two agree for all intents and purposes. This program has met with considerable success in some models of statistical physics, for example for the local Glauber dynamics in the two-dimensional Ising model 4,5 .

The difference between two (normalized) probability distributions $\pi$ and $\tilde{\pi}$ is often quantified by the total variation distance (TVD) 6.7.

$$
\begin{align*}
\|\tilde{\pi}-\pi\|_{\mathrm{TV}} & =\frac{1}{2} \int_{\Omega}|\tilde{\pi}(\mathbf{x})-\pi(\mathbf{x})| \mathrm{d} \mathbf{x}  \tag{1}\\
& =\max _{\mathcal{A} \subseteq \Omega}|\tilde{\pi}(\mathcal{A})-\pi(\mathcal{A})| \tag{2}
\end{align*}
$$

which satisfies $0 \leq\|\tilde{\pi}-\pi\|_{\mathrm{TV}} \leq 1$. The mixing time, the most relevant figure of merit for a Monte Carlo algorithm, is defined as the time $t$ after which the TVD (with $\tilde{\pi} \equiv \pi^{(t)}$ in eq. (2) ) is smaller than a given threshold $\epsilon$, for any initial distribution $\pi^{(0)}$. Although it is of great conceptual importance, the TVD cannot usually be computed. In statistical physics, this is already because the normalization of the Boltzmann weight, the partition function $Z=\int_{\Omega} \exp (-\beta E)$, is most often un-
known. Also, the distribution $\pi^{(t)}(\mathbf{x})$ is not known explicitly. It is because of this difficulty that practical simulations often carry systematic uncertainties that are difficult to quantify, and that heuristic convergence criteria for the approach towards equilibrium in MCMC abound 3, 8, 9]. They most often involve time-correlation functions of observables, rather than the probability distribution itself (as in eqs (1) and (2)).

In rare cases, MCMC algorithms allow for the definition of a stopping rule (based on the concept of a strong stationary time $[6]$ ), that yields a simulation-dependent value of $t$ at which the configuration is sampled exactly from the distribution $\pi$. The value of $t$ now often depends on the realization of the Markov chain (that is, on the individual sampled moves and, ultimately, on the drawn random numbers). The framework of stopping rules can be used to bound the mixing time [6]. Stopping rules exist for quite intricate models, as for example the Ising model, using the mixing-from-the-past framework 3,10 .

The great majority of Markov-chain Monte Carlo algorithms are reversible (they satisfy the detailed-balance condition). This is the case for example for all algorithms that are based on the Metropolis or the heat-bath algorithms [1,3], which allow reversible MCMC algorithms to be readily constructed for any distribution $\pi$, that is, for an arbitrary energy $E(\mathbf{x})$. In recent years, however, irreversible MCMC methods based on the global balance condition have shown considerable promise $11-15$. In these algorithms, $\pi^{(t)}$ approaches $\pi$ for long times, but the flows no longer vanish. One particular irreversible Markov chain, the event-chain Monte Carlo (ECMC) algorithm 13, 14, has proven useful for systems ranging from hardsphere models 16 to spin systems [17], polymers 18, 19 and to long-range interacting ensembles of molecules, such as water [20], where the Coulomb interaction plays a dominant role 21. Although there have been many indications of the algorithm's power, no exact results were known for the mixing behavior of ECMC, except for the case of a single particle, $N=122$.

In the present paper, we rigorously establish ECMC mixing times and stopping rules of the model of $N$ hard spheres on a one-dimensional line with periodic boundary conditions (a circle). Reversible MCMC algorithms for this model and its variants were analyzed rigorously 23, 24 and irreversible MCMC algorithm were discussed in detail [15]. The 1D hard-sphere model and reversible and irreversible MCMC algorithms are closely related to the symmetric exclusion process (SEP) on a periodic lattice 25 and to the totally asymmetric simple exclusion process (TASEP) 26 28. For ECMC, an algorithm that is closely related to the lifted Metropolis algorithm [15], we compute the TVD in a special case, and obtain the mixing times as a function of the parameter $\epsilon$. We confirm the $\mathcal{O}\left(N^{2} \log N\right)$ mixing time that had been conjectured on the basis of numerical simulations 15. Furthermore, we obtain a stopping rule for ECMC. We moreover present sequential variants of the forward Metropolis algorithm
and the ECMC algorithm. For the latter, we prove an $\mathcal{O}\left(N^{2}\right)$ exact-sampling result that seems however not to generalize the discretized version of the algorithm.

Hard spheres in 1D, reversible Monte Carlo. The mixing and convergence behavior of Markov chains for particle systems has been much studied. As for hard spheres in 2D and above, phase transitions have only been identified by numerical simulation $16,29,30$, it is natural that few rigorous results are available for the convergence and mixing behavior of MCMC algorithms in $D>1$ 31, 32 . We thus restrict our attention to the 1D hard-sphere model with periodic boundary conditions and treat both the discrete and the continuous cases.


Fig. 1: 1D hard-sphere model with periodic boundary conditions. (a): $N$ spheres of diameter $d$ on a ring of length $L$. (b): $N$ point particles on a ring of length $L_{\text {free }}=L-N d$. Configurations and local MCMC algorithms are equivalent for both representations.

The 1D hard-sphere model can be represented as $N$ spheres of diameter $d$ on a line of length $L$ with periodic boundary conditions (that is, on a ring, see Fig. 17). A valid configuration $a$ of $N$ spheres has unit statistical weight $\pi(a)=1$. Spheres do not overlap, so that the distance between sphere centers, and in particular between neighboring spheres, is larger than $d$. Each valid configuration of $N$ hard spheres is equivalent to a configuration of $N$ point particles on a ring of length $L_{\text {free }}=L-N d>0$ (see Fig. 1b), and the partition function of the model equals $S=L_{\text {free }}^{N}$, which proves the absence of a phase transition.

We only consider local Markov chains, where the move of sphere $i$ is accepted or rejected based solely on the position of its neighbors. One might think that this requires the distribution $p(\epsilon)$ to vanish for $|\epsilon|>2 d$. We rather implement locality by rejecting a move of sphere $i$ not only if the displacement leads to an overlap, but also if sphere $i$ would hop over one of its neighbors. In this way, any local Monte Carlo move of spheres on a circle corresponds to an equivalent move in the point-particle representation (for which there are no overlaps and moves are rejected only because they represent a hop over a neighbor). The dynamics of both models is the same. This implies that the MCMC dynamics of the 1D hard-sphere model has only a trivial density dependence.

Although we will study Markov chains that relabel spheres, we are interested only in the relaxation of quantities that can be expressed through the unlabeled distances
between neigboring spheres. This excludes the mixing in permutation space of labels or the self-correlation of a given sphere with itself (or another labeled sphere) at different times. Trivial uniform rotations are thus also neglected.

Detailed balance consists in requiring:

$$
\begin{equation*}
\pi(a) p(a \rightarrow b)=\pi(b) p(b \rightarrow a) \tag{3}
\end{equation*}
$$

where $p(a \rightarrow b)$ is the conditional probability to move from configuration $a$ to configuration $b$. The heat-bath algorithm is a local reversible MCMC algorithm. At each time step, it replaces a sampled sphere $i$ randomly in between its neighbors. The heat-bath algorithm mixes in at least $\mathcal{O}\left(N^{3}\right)$ and at most $\mathcal{O}\left(N^{3} \log N\right)$ time steps 24, although numerical simulations clearly favor the latter possibility $\left(\mathcal{O}\left(N^{3} \log N\right)\right) 15{ }^{1}$ For the one-dimensional hard-sphere model on a line without periodic boundary condition, the mixing time $\mathcal{O}\left(N^{3} \log N\right)$ is rigorously proven 23].

Analogous to the heat-bath algorithm, the reversible Metropolis algorithm also satisfies the detailed-balance condition: At each time step, a randomly chosen sphere $i$ attempts a move by $\epsilon$ taken from some probability distribution. The move is rejected if the proposed displacement $\epsilon$ is larger than the free space in the direction of the proposed move $\left(x_{i_{+}}-x_{i}-d\right.$ for $\left.\epsilon>0\right)$ or behind it ( $x_{i}-x_{i_{-}}-d$ for $\epsilon<0$ ) (where we suppose that $i_{+}$ is the right-hand-side neighbor of $i$, etc, and imply periodic boundary conditions). In the point-particle model, the equivalent move is rejected if the particle would hop over one or more of its neighbors and is accepted otherwise. Rigorous results for mixing times are unknown for the Metropolis algorithm, but numerical simulations clearly identify $\mathcal{O}\left(N^{3} \log N\right)$ mixing as for the heat-bath algorithm [15. In the discrete 1D hard-sphere model on the circle with $L$ sites and $N$ particles, the Metropolis algorithm is implemented in the socalled simple exclusion process (SEP), where at each time step, a randomly chosen particle attempts to move with equal probability to each of its two adjacent sites. The move is rejected if that site is already occupied. The mixing time of the SEP is $\sim\left(4 \pi^{2}\right)^{-1} N L^{2} \log N($ for $L \geq 2 N) 25$.

From the forward Metropolis to the event-chain algorithm. - Irreversible Monte Carlo algorithms violate the detailed-balance condition of eq. (3) but instead satisfy the weaker global-balance condition

$$
\begin{equation*}
\sum_{b} \pi(b) p(b \rightarrow a)=\pi(a) \tag{4}
\end{equation*}
$$

Together with the easily satisfiable aperiodicity and irreducibility conditions [6], the global-balance condition ensures that the steady-state solution corresponds to the probability $\pi$, but without necessarily cancelling the net flow $\pi(a) p(a \rightarrow b)-\pi(b) p(n \rightarrow a)$ between configurations $a$

[^1]

Fig. 2: Flow of the forward swap Metropolis algorithm into a configuration ( $a, i$ ) (the active sphere $i$ is shown in green). A rejected sphere move, by a displacement $\epsilon_{\mathcal{R}}$ (upper case), entails a swap and contributes $\mathcal{R}_{\epsilon_{\mathcal{R}}}$. An accepted sphere move, by a displacement $\epsilon_{\mathcal{A}}$ (lower case), contributes $\mathcal{A}_{\epsilon_{\mathcal{A}}}$. For any $\epsilon$, one of the flows equals one, and the other zero $\left[\left(\mathcal{R}_{\epsilon}, \mathcal{A}_{\epsilon}\right) \in\right.$ $\{(0,1),(1,0)\}]$.
and $b$ (cf eq. (3)). Here, we take up the forward Metropolis algorithm studied earlier, in a new variant that involves swaps. This allows us to arrive at an exact mixing result.

In the forward swap Metropolis algorithm ${ }^{2}$, at each time step, a randomly chosen sphere $i$ attempts to move by a random displacement $\epsilon$ with a predefined sign (that for clarity, we take to be positive). If the move is rejected (that is, the displacement $\epsilon$ does not yield a valid hardsphere configuration), the sphere swaps its label with the sphere responsible for the rejection (see the upper move in Fig. 22). Else, if the displacement $\epsilon$ is accepted, the sphere $i$ simply moves forward (see the lower move in Fig. 2). The total flow into a configuration $(a, i)$ (that is, the $N$-sphere configuration $a$ with the active sphere $i$ ) is:

$$
\begin{equation*}
\mathcal{F}(a, i)=\int_{0}^{\infty} \mathrm{d} \epsilon p(\epsilon) \underbrace{\left[\mathcal{A}_{\epsilon}(a, i)+\mathcal{R}_{\epsilon}(a, i)\right]}_{=1 \text { (see Fig. 2] }}=1=\pi(a), \tag{5}
\end{equation*}
$$

so that the algorithm satisfies global balance. The swap allows both the rejected and the accepted moves into the configuration $(a, i)$ to involve the sphere $i$ only. The forward swap Metropolis algorithm is equivalent (up to relabeling) to the forward Metropolis algorithm treated earlier if at each time step the active sphere $i$ is sampled randomly. The mixing time of this algorithm was conjectured to be $\mathcal{O}\left(N^{5 / 2}\right)$, based on numerical simulations [15]. This agrees with the proven mixing time scale of the totally asymmetric simple exclusion process (TASEP) 28].
The forward swap Metropolis algorithm satisfies global balance for any choice of the sphere $i$ and any step-size distribution $p(\epsilon)$. This implies that the active-sphere index $i$ need not be sampled randomly for the algorithm to remain valid. This distinguishes it from the forward Metropolis algorithm (without the swaps) treated in previous work 15]. In particular, the sphere $i$ can be active for several

[^2]

Fig. 3: Forward swap Metropolis algorithm, with configurations $\mathbf{x}_{t}, \ldots, \mathbf{x}_{t+5}$. The active sphere is sampled randomly at each time step so that the swaps have no other action than to relabel the spheres.
chains in a row. The algorithm, run with the following sequence of active spheres:

$$
\begin{equation*}
\ldots, \underbrace{i, i, \ldots, i, i}_{\text {chain } n}, \underbrace{j, j, \ldots, j, j}_{\text {chain } n+1}, \underbrace{k, k, \ldots, k, k}_{\text {chain } n+2}, \ldots \tag{6}
\end{equation*}
$$

is equivalent to the lifted forward Metropolis algorithm studied earlier [15], if the active spheres $i, j, k, \ldots$ in eq. (6) are sampled randomly. The algorithm naturally satisfies the global balance condition, and again, each individual move attempts a displacement by a distance $\epsilon$ sampled from a given distribution $p(\epsilon)$ that vanishes for negative $\epsilon$, and the chain lengths (number of repetitions of $i, j, k, \ldots$ ) $n, n+1, \ldots$ are sampled from a distribution. Numerical simulations have lead to the conjecture that this algorithm mixes in $\mathcal{O}\left(N^{2} \log N\right)$ time steps 15.

ECMC is the continuous-time limit of the lifted forward Metropolis algorithm, with the simultaneous limits $\epsilon \rightarrow 0$ and $l \rightarrow \infty$, but $(\langle\epsilon\rangle l) \rightarrow \ell$, where the chain length $\ell$ on the scale $L_{\mathrm{free}}$, is again sampled from a given probability distribution. In the point-particle representation of Fig. 1(b), one "event" chain of the ECMC algorithm simply moves the active particle $i$ from its initial position $x_{i}$ to $x_{i}+\ell$. It advances the time as $t \rightarrow t+\ell$, and increments the number of chains as $n \rightarrow n+1$. The number of eponymous "events" of ECMC (the number of changes of the active sphere) then grows approximately as $\sim\left(\ell / L_{\text {free }}\right) N$. When $\ell \sim \operatorname{unif}\left(0, L_{\text {free }}\right)$, this places particle $i$ at a random position on a ring. For this special uniform distribution of chain lengths, a perfect sample is clearly obtained once all particles were at least once picked as the active particle. This situation will now be analyzed in terms of the coupon-collector problem (see $[33,34]$ ).

For the ECMC with $\ell \sim \operatorname{unif}\left(0, L_{\text {free }}\right)$, the TVD can be expressed by the probability that at least one particle has never been picked as an active particle of a chain. Without
restriction, we suppose that the initial configuration is the compact state $\mathbf{x}=\{0,0, \ldots, 0\}$. We also measure time in the number of chains $n$ ( $n$ translates into an MCMC time as $t(n)=\langle\ell\rangle n$ and is easily converted into the number of events). In eq. (2), the set $\mathcal{A}$ is

$$
\begin{equation*}
\mathcal{A}=\left\{\mathbf{x} \mid \exists i \text { with } x_{i}=0\right\} . \tag{7}
\end{equation*}
$$

Also, clearly, $\pi^{(n)}(\mathcal{A})$ equals the probability that at least one particle has never been picked as an active particle of a chain, whereas $\pi(\mathcal{A})=0$, as it is a lower-dimensional subset of $\Omega$. From eqs (1) and (2), therefore (for $N \rightarrow \infty$ ):

$$
\begin{equation*}
\left\|\pi^{(n)}-\pi\right\|_{\mathrm{TV}} \sim 1-\exp \left[-\exp \left(-\frac{n-N \log N}{N}\right)\right] \tag{8}
\end{equation*}
$$

where we used the analytically known asymptotic tail probability for the coupon-collector problem $\sqrt{33}$ (see Fig. 4.

Rather than computing the difference between $\pi^{(n)}$ and $\pi$ at a fixed number $n$ of chains, one can simply run ECMC (with $\left.\ell \sim \operatorname{unif}\left(0, L_{\text {free }}\right)\right)$ until the time at which chains with any of the $N$ particles as active ones have completed. The expected number of chains or, in the language of the coupon-collector problem, the expected value of $n_{1}$ to "collect the last coupon" is given by

$$
\begin{equation*}
\left\langle n_{1}\right\rangle=N H_{N}=N \log N+\gamma N+\frac{1}{2}+\mathcal{O}(1 / N) \tag{9}
\end{equation*}
$$

where $H_{N}=\frac{1}{1}+\frac{1}{2}+\cdots+\frac{1}{N}$ is the $N$ th harmonic number and $\gamma=0.5772 \ldots$ is the Euler-Mascheroni constant. The distribution of this number of chains can be obtained from the tail distribution contained in eq. (8) (see Fig. 4). In both ways, we see that mixing takes place after $\mathcal{O}(N \log N)$ chains (corresponding to $\mathcal{O}\left(N^{2} \log N\right)$ events), confirming, for a special distribution of $\ell$, an earlier conjecture 15. The discussed mixing behavior of ECMC can more generally be obtained for distributions $\ell \sim \operatorname{unif}\left(c, c+L_{\text {free }}\right)$ with arbitrary (and even with negative) $c$. In our special case, choosing $c=-L_{\text {free }} / 2$ would lead to the smallest number of individual events. In view of the practical applications of ECMC, it appears important to understand whether this dependence on the distribution of $\ell$ rather than on its mean value has some relevance for the simulation of discrete 1D models, and whether it survives in higher dimensions, and for continuous (non-hard-sphere) potentials.

We next consider more general distributions, namely the uniform distribution $\ell \sim \operatorname{unif}\left(0, \lambda L_{\text {free }}\right)$, as well as the Gaussian distribution $\mathcal{N}\left(\mu, \sigma^{2}\right)$, where $\mu$ is the mean value and $\sigma$ the standard deviation. Again, particles are effectively independent and we conjecture the mixing time (which can now never lead to perfect sampling) to be governed by the particle which has moved the least number, $m$, of times. This is equivalent to the $m$-coupon generalization of the coupon-collector problem [33, whose tail probability is given by:

$$
\begin{equation*}
P\left(n_{m}\right)=\exp (-\Upsilon /(m-1)!) \tag{10}
\end{equation*}
$$

with

$$
\begin{equation*}
\Upsilon=\exp \left[-\frac{n_{m}-N \log N-(m-1) N \log \log N}{N}\right] \tag{11}
\end{equation*}
$$

(see Fig. 4). This means that the number of chains to collect each of the $N$ coupons at least $m$ times only add an $N \log \log N$ correction to the general $N \log N$ scale of chains.


Fig. 4: Cumulative probability of the coupon collector problem ( $m=1$ ), and of the $m$-coupon collector problem for $m=2$ and $m=3$. Numerical simulations for $N=2^{20}$ particles are compared to the asymptotic tail probability of eq. 11 .

To gain intuition, we now compute the TVD for the single-particle problem (for which $\ell \equiv \epsilon$ ). For simplicity, we set $L_{\text {free }}=1$ (measure the standard deviations in units of $\left.L_{\text {free }}\right)$. The TVD for chain lengths $\ell_{i} \sim \operatorname{unif}(0, \lambda)$, as discussed, equals the one for $\ell_{i} \sim \operatorname{unif}(-\lambda / 2, \lambda / 2)$. The sum of $m$ chains then follows the distribution:

$$
\begin{equation*}
p_{m}^{\mathrm{unif}}(x)=\int_{-\infty}^{\infty} \mathrm{d} t \mathrm{e}^{-2 \pi i t x}\left[\frac{\sin (\pi \lambda t)}{\pi \lambda t}\right]^{m} \tag{12}
\end{equation*}
$$

Using the Poisson summation formula and subtracting the equilibrium distribution $\pi=1$, we find:

$$
\sum_{k=-\infty}^{\infty} p_{m}^{\mathrm{unif}}(x+k)-1=\sum_{k \in \mathbb{N}^{+}} 2\left[\frac{\sin (\pi k \lambda)}{\pi k \lambda}\right]^{m} \cos (2 \pi k x)
$$

The total variation distance for chain lengths $\ell_{i} \sim$ unif $(0, \lambda)$ thus satisfies:

$$
\begin{align*}
&\left\|\pi^{(m)}-\pi\right\|_{\mathrm{TV}} \\
&=\int_{0}^{1} d x \mid \left.\sum_{k \in \mathbb{N}^{+}}\left[\frac{\sin (\pi k \lambda)}{\pi k \lambda}\right]^{m} \cos (2 \pi k x) \right\rvert\, \\
& \sim \frac{2}{\pi}\left|\frac{\sin (\pi \lambda)}{\pi \lambda}\right|^{m}(\text { for } \quad m \rightarrow \infty) \tag{13}
\end{align*}
$$

The TVD trivially vanishes for integer $\lambda$ (see Fig. 5a). Its peaks decay as $\frac{2}{\pi}(\pi \lambda)^{-m}$.

For Gaussian-distributed chain lengths $\ell_{i} \sim \mathcal{N}\left(\mu, \sigma^{2}\right)$, the sum of $m$ chains is distributed as:

$$
\begin{equation*}
\sum_{i=1}^{m} \ell_{i} \sim \mathcal{N}\left(m \mu, m \sigma^{2}\right) \tag{14}
\end{equation*}
$$

With $\vartheta_{3}$ the Jacobi theta function, we now have

$$
\begin{align*}
& \sum_{k=-\infty}^{\infty} p_{m}^{\text {Gauss }}(x+k)-1=\vartheta_{3}\left[\pi(x+\mu), \exp \left(-2 \pi^{2} m \sigma^{2}\right)\right] \\
& \quad=2 \sum_{k=1}^{\infty} \exp \left(-2 k^{2} \pi^{2} m \sigma^{2}\right) \cos [2 k \pi(x+m \mu)] \tag{15}
\end{align*}
$$

The total variation distance for the distribution of eq. 14 satisfies:

$$
\begin{align*}
\| \pi^{(m)}- & \pi \|_{\mathrm{TV}} \\
= & \int_{0}^{1} d x\left|\sum_{k=1}^{\infty} \exp \left(-2 k^{2} \pi^{2} m \sigma^{2}\right) \cos (2 k \pi x)\right| \\
& \sim \frac{2}{\pi} \exp \left(-2 \pi^{2} m \sigma^{2}\right)\left(\text { for } \quad m \sigma^{2} \rightarrow \infty\right) \tag{16}
\end{align*}
$$

(see Fig. 5b).


Fig. 5: TVD for a single sphere on a ring with uniform and Gaussian distributions of $\ell \equiv \epsilon$. (a): TVD after $m$ displacements $\epsilon \sim \operatorname{unif}\left(0, \lambda L_{\text {free }}\right)$. The TVD trivially vanishes for integer $\lambda$. Peaks decay as $\frac{2}{\pi}(\pi \lambda)^{-m}$ (for $m \rightarrow \infty$ ). (b): TVD for $m$ Gaussian displacements with standard deviation $\sigma$, compared with its first order approximation from the Jacobi $\vartheta$ function (see eq. 16). The inset illustrates the good agreement of the approximation on a logarithmic scale.

Both for the uniform and the Gaussian distribution, the single-sphere TVD decreases exponentially with the number $m$ of displacements (which are equivalent to singleparticle chains). We expect the same behavior for the $N$-sphere problem, where $m$ is now the number of chains for the $m$-coupon problem.

## Sequential forward Metropolis, sequential

 ECMC. - ECMC, with randomly sampled initial spheres and a standard deviation of the chain-length distribution $\sigma \sim L_{\text {free }}$, mixes in $\mathcal{O}\left(N^{2} \log N\right)$ events (corresponding to $\mathcal{O}(N \log N)$ chains). In the label-switching framework of ECMC, each chain consists in advancing the particle $i$ by a distance $\ell$ times, and both the ECMC and the forward-Metropolis versions are correct. Instead of sampling the active sphere for each chain, so that the coupon-collector aspect necessarily brings in the $\log N$ factor in the scaling of mixing times, we may also sequentially increment the active-sphere index for each| Time | Configuration | Chain length |
| :---: | :---: | :---: |
| $t$ | 3) | $l=3$ |
| $t+3$ |  | $l=4$ |
| $t+7$ |  | $l=3$ |
| $t+10$ |  | $l=1$ |
| $t+11$ |  | $l=2$ |
| $t+13$ |  | Final |

Fig. 6: Sequential lifted forward Metropolis algorithm (with swaps). Configurations $\mathbf{x}_{t}, \ldots, \mathbf{x}_{t+13}$ sampled through five chains with active sphere $1,2, \ldots, 5$ are shown. Chain lengths are $l_{1}=3, \ldots, l_{5}=2$. Each sphere displacement $\epsilon_{t}>0$ is either accepted or, if rejected, it induces a swap, so that the same sphere remains active throughout a chain.
chain (see Fig. 6):

$$
\begin{equation*}
\ldots, \underbrace{i, \ldots, i}_{\text {chain } i}, \underbrace{i+1, \ldots, i+1}_{\text {chain } i+1}, \underbrace{i+2, \ldots, i+2}_{\text {chain } i+2}, \ldots \tag{17}
\end{equation*}
$$

(where particle numbers are implied modulo $N$ ). Sequential ECMC, with a distribution $\ell_{i} \sim \operatorname{unif}\left(0, L_{\text {free }}\right)$ produces an exact sample in $\mathcal{O}\left(N^{2}\right)$ events (corresponding to exactly $N$ chains).

Evidently, the analysis of eqs (13) and can be applied to the sequential ECMC with distributions such as unif ( $0, \lambda L_{\text {free }}$ ) and, more generally, distributions with $\sigma \sim L_{\text {free }}$. After each "sweep" of chains, the TVD factorizes, and we expect mixing to take place after $\mathcal{O}(N)$ chains (corresponding to $\mathcal{O}\left(N^{2}\right)$ events).

ECMC is the limit of the lifted forward Metropolis algorithm, and the sequential ECMC the limit of the sequential lifted forward Metropolis algorithm for step sizes much smaller than the mean free space between spheres $\left(\langle\epsilon\rangle=L_{\text {free }} /(2 N \alpha)\right.$ with $\left.\alpha \gg 1\right)$. For a given discretization $2 / \alpha$, and for small $N$, the sequential lifted forward algorithm mimics the $\mathcal{O}\left(N^{2}\right)$ mixing of the sequential ECMC, but for large $N$, it seems to cross over into $\mathcal{O}\left(N^{2} \log N\right)$ mixing (see Fig. 77). $\mathcal{O}\left(N^{2}\right)$ mixing also emerges at fixed $N$ for large $\alpha$ (see Fig. 7 b ). (This is obtained using the heuristic mid-system variance $x_{i+N / 2}-x_{i}$ for ordered $x_{i}$, see [15].) In contrast, the random lifted forward Metropolis algorithm shows $\mathcal{O}\left(N^{2} \log N\right)$ mixing (see Fig. 7re), as discussed earlier [15]. This scaling is little influenced by the discretization (see Fig. 7 d ). It thus appears that the $N \rightarrow \infty$ limit of the sequential lifted forward Metropolis algorithm does not commute with the small discretization limit $\alpha \rightarrow \infty$.


Fig. 7: Crossover from the discrete lifted algorithm to ECMC, via the variance of the mid-system distance $x_{i+N / 2}-x_{i}$ for ordered $x_{i}$, started from compact initial condition (see [15]). Discrete step size with $\epsilon \sim \operatorname{unif}\left(0, L_{\text {free }} / N / \alpha\right)$, and chain length $l \sim \operatorname{unif}\{\alpha, \alpha N\}$ (a): Sequential lifted Metropolis with constant $\alpha=10$ for different $N$ : The cross-over from perfect sampling for small $N$ at a time scale $\mathcal{O}\left(N^{2}\right)$ towards $\mathcal{O}\left(N^{2} \log N\right)$ appears evident. (b): Sequential algorithm for $N=8192$, with increasing $\alpha$ : $\mathcal{O}\left(N^{2}\right)$ mixing scale emerges for large $\alpha$. (c): Random lifted Metropolis with $\alpha=10$ for different $N$ (legend as in (a)): $\mathcal{O}\left(N^{2} \log N\right)$ mixing time scale (conjectured earlier (15]). (d): Random lifted forward Metropolis algorithm: Limited role of $\alpha$ (color code for $\alpha$ as in (b)).

Conclusions. - In this paper we have proven that for 1D hard spheres, ECMC with a uniform distribution of chain length $\ell \sim\left[0, L_{\text {free }}\right]$, with $L_{\text {free }}=L-N d$ realizes a perfect sample in $\mathcal{O}\left(N^{2} \log N\right)$ events that correspond to $\mathcal{O}(N \log N)$ chains. This confirms, in a special case, an earlier conjecture 15 for the mixing time of ECMC. For this case, we can compute the TVD but also indicate a stopping rule for a time (which depends on the particular realization of the Markov chain), after which the Markov chain is in equilibrium. We have also provided numerical evidence that the $N^{2} \log N$ mixing prevails for other distributions of $\ell$, namely for the uniform distribution unif ( $0, \lambda L_{\text {free }}$ ) and the Gaussian, and used the coupon-collector approximation to justify this approximation.

We have furthermore discussed a sequential ECMC which mixes in a time $\mathcal{O}\left(N^{2}\right)$. For this algorithm, "particle swaps" are essential. We have checked that the discrete version of this algorithm, namely the sequential lifted forward Metropolis algorithm crosses over, as the number $N$ of spheres is increased, to an $\mathcal{O}\left(N^{2} \log N\right)$ mixing behavior. In this formula, the origin of the logarithm is unclear, as it can no longer stem from the coupon collector. It would be of great interest for the fundamental under-
standing of irreversible MCMC algorithms to extend the results from ECMC to discrete versions, that is the lifted forward Metropolis algorithm and its sequential variant, as well as to the corresponding lattice models that may be easier to treat.

The lessons from our analysis of 1D hard-sphere systems are threefold. First, irreversible Markov chains can be proven to mix on shorter time scales than reversible algorithms. Second, the speed of ECMC depends on the whole distribution of the chain lengths $\ell$, but not on its mean value. Third, sequential-update algorithms (that remain valid in higher dimensions) can mix on faster time scales than random-update versions. It remains to be seen how these lessons carry over to more intricate models and to higher dimensions.

$$
* * *
$$

We thank Florent Krzakala for a helpful discussion. W.K. acknowledges support from the Alexander von Humboldt Foundation.

## REFERENCES

[1] Metropolis N., Rosenbluth A. W., Rosenbluth M. N., Teller A. H. and Teller E., J. Chem. Phys., 21 (1953) 1087.
[2] Devroye L., Non-Uniform Random Variate Generation (Springer New York) 1986.
[3] Krauth W., Statistical Mechanics: Algorithms and Computations (Oxford University Press) 2006.
[4] Martinelli F., Lectures on Glauber Dynamics for Discrete Spin Models in Lectures on Probability Theory and Statistics: Ecole d'Eté de Probabilités de Saint-Flour XXVII - 1997, edited by Bernard P., (Springer, Berlin, Heidelberg) 1999 pp. 93-191.
[5] Lubetzky E. and Sly A., Comm. Math. Phys., 313 (2012) 815.
[6] Levin D. A., Peres Y. and Wilmer E. L., Markov Chains and Mixing Times (American Mathematical Society) 2008.
[7] Diaconis P., J. Stat. Phys., 144 (2011) 445.
[8] Berg B. A., Markov Chain Monte Carlo simulations and their statistical analysis: with web-based Fortran code (World Scientific) 2004.
[9] Landau D. and Binder K., A Guide to Monte Carlo Simulations in Statistical Physics (Cambridge University Press) 2013.
[10] Propp J. G. and Wilson D. B., Random Structures \& Algorithms, 9 (1996) 223.
[11] Turitsyn K. S., Chertkov M. and Vucelja M., Physica $D, 240$ (2011) 410.
[12] Fernandes H. C. and Weigel M., Comput. Phys. Commun., 182 (2011) 1856.
[13] Bernard E. P., Krauth W. and Wilson D. B., Phys. Rev. E, 80 (2009) 056704.
[14] Michel M., Kapfer S. C. and Krauth W., J. Chem. Phys., 140 (2014) 054116.
[15] Kapfer S. C. and Krauth W., Phys. Rev. Lett., 119 (2017) 240603.
[16] Bernard E. P. and Krauth W., Phys. Rev. Lett., 107 (2011) 155704.
[17] Lei Z. and Krauth W., EPL (Europhysics Letters), 121 (2018) 10008.
[18] Kampmann T. A., Boltz H.-H. and Kierfeld J., J. Chem. Phys., 143 (2015) 044105.
[19] Harland J., Michel M., Kampmann T. A. and Kierfeld J., EPL (Europhysics Letters), 117 (2017) 30001.
[20] Faulkner M. F., Qin L., Maggs A. C. and Krauth W., arXiv 1804.05795, (2018) .
[21] Kapfer S. C. and Krauth W., Phys. Rev. E, 94 (2016) 031302.
[22] Diaconis P., Holmes S. and Neal R. M., Annals of Applied Probability, 10 (2000) 726.
[23] Randall D. and Winkler P., Mixing Points on an Interval in proc. of Proceedings of the Seventh Workshop on Algorithm Engineering and Experiments and the Second Workshop on Analytic Algorithmics and Combinatorics, ALENEX /ANALCO 2005, Vancouver, BC, Canada, 22 January 20052005 pp. 218-221.
[24] Randall D. and Winkler P., Mixing Points on a Circle Vol. 3624 of Lecture Notes in Computer Science (Springer, Berlin, Heidelberg) 2005 pp. 426-435.
[25] Lacoin H., Ann Inst H Poincaré Probab Statist, 53 (2017) 1402.
[26] Gwa L.-H. and Spohn H., Phys. Rev. Lett., 68 (1992) 725.
[27] Chou T., Mallick K. and Zia R. K. P., Rep. Prog. Phys., 74 (2011) 116601.
[28] Baik J. and Liu Z., J. Stat. Phys., 165 (2016) 1051.
[29] Hoover W. G. and Ree F. H., J. Chem. Phys., 49 (1968) 3609.
[30] Alder B. J. and Wainwright T. E., Phys. Rev., 127 (1962) 359.
[31] Wilson D. B., Random Structures \& Algorithms, 16 (2000) 85.
[32] Kannan R., Mahoney M. W. and Montenegro R., Rapid mixing of several Markov chains for a hard-core model in proc. of 14 th annual ISAAC Lecture Notes in Computer Science (Springer, Berlin, Heidelberg) 2003 pp. 663-675.
[33] Erdős P. and Rényi A., Magyar Tud. Akad. Mat. Kutató Int. Közl., 6 (1961) 215-220.
[34] Blom G., Holst L. and Sandell D., Problems and Snapshots from the World of Probability (Springer New York) 1994.


[^0]:    ${ }^{(a)}$ ze.lei@ens.fr
    (b) werner.krauth@ens.fr

[^1]:    ${ }^{1}$ As mentioned, we do not consider uniform rotations of the total system, which would mix only on a time scale $\mathcal{O}\left(N^{4}\right)$.

[^2]:    ${ }^{2}$ The forward Metropolis algorithm introduced earlier 15 did not feature swaps.

