# THE INTERCHANGE PROCESS ON HIGH-DIMENSIONAL PRODUCTS 

By Jonathan Hermon ${ }^{1}$ and Justin Salez ${ }^{2}$<br>${ }^{1}$ Mathematics Department, University of British Columbia, jhermon@math.ubc.ca<br>${ }^{2}$ CEREMADE, Université Paris-Dauphine \& PSL, justin.salez@dauphine.psl.eu

We resolve a long-standing conjecture of Wilson (Ann. Appl. Probab. 14 (2004) 274-325), reiterated by Oliveira (2016), asserting that the mixing time of the interchange process with unit edge rates on the $n$-dimensional hypercube is of order $n$. This follows from a sharp inequality established at the level of Dirichlet forms, from which we also deduce that macroscopic cycles emerge in constant time, and that the log-Sobolev constant of the exclusion process is of order 1. Beyond the hypercube, our results apply to cartesian products of arbitrary graphs of fixed size, shedding light on a broad conjecture of Oliveira (Ann. Probab. 41 (2013) 871-913).

## 1. Introduction.

1.1. Interchange process. Let $G=\left(V_{G}, E_{G}\right)$ be a finite undirected connected graph. The interchange process (IP) on $G$ is the continuous-time random walk $\left(\xi_{t}\right)_{t \geq 0}$ on the symmetric group $\mathfrak{S}\left(V_{G}\right)$ with initial condition $\xi_{0}=\mathrm{id}$ and the following Markov generator: for all observables $f: \mathfrak{S}\left(V_{G}\right) \rightarrow \mathbb{R}$,

$$
\begin{equation*}
\mathcal{L}_{G}^{\mathrm{IP}} f(\sigma):=\sum_{e \in E_{G}}\left(f\left(\sigma \tau_{e}\right)-f(\sigma)\right) \tag{1}
\end{equation*}
$$

where $\tau_{e}$ denotes the transposition of the endpoints of the edge $e$. One may think of each vertex as carrying a labelled particle, and of the edges as being equipped with independent unit-rate Poisson clocks: whenever a clock rings, the particles sitting at the endpoints of the corresponding edge simply exchange their positions.

Since $\mathcal{L}_{G}^{\mathrm{IP}}$ is symmetric and irreducible, the law of $\xi_{t}$ converges to that of a uniform permutation $\xi_{\star}$ as $t \rightarrow \infty$. We shall here be interested in the time-scale on which this convergence occurs, as traditionally measured by the total-variation mixing time:

$$
\begin{equation*}
t_{\mathrm{MIX}}^{\mathrm{P}}(G):=\min \left\{t \geq 0: \max _{A \subseteq \subseteq\left(V_{G}\right)}\left|\mathbb{P}\left(\xi_{t} \in A\right)-\mathbb{P}\left(\xi_{\star} \in A\right)\right| \leq \frac{1}{e}\right\} \tag{2}
\end{equation*}
$$

Understanding the relation between this fundamental quantity and the geometry of $G$ is a challenging problem, to which a remarkable variety of tools have been applied: representation theory $[9,13]$, couplings [ $2,23,27,30$ ], eigenvectors [18, 30], functional inequalities [8, 22, $31]$, comparison methods [4,11], etcetera. The question is of course particularly meaningful when the number of states becomes large, and one is thus naturally led to study asymptotics along various growing sequences of graphs $\left(G_{n}\right)_{n \geq 1}$.

The case of the $n$-clique $G_{n}=\mathcal{K}_{n}$ has been extensively studied under the name random transposition shuffle. In particular, Diaconis and Shahshahani [13] proved that

$$
\begin{equation*}
t_{\mathrm{MIX}}^{\mathrm{IP}}\left(\mathcal{K}_{n}\right)=\frac{\log n}{n}(1+o(1)) \tag{3}
\end{equation*}
$$

In fact, this was shown for any precision $\varepsilon \in(0,1)$ instead of $\frac{1}{e}$ in $(2)$, thereby extablishing the very first instance of what is now called a cutoff phenomenon [10]. Another well-understood case is the $n$-path $\mathcal{P}_{n}$, for which Lacoin [21] recently proved cutoff at time

$$
\begin{equation*}
t_{\mathrm{MIX}}^{\mathrm{IP}}\left(\mathcal{P}_{n}\right)=\frac{n^{2} \log n}{2 \pi^{2}}(1+o(1)) \tag{4}
\end{equation*}
$$

There are, however, many simple graph sequences along which even the order of magnitude of $t_{\text {MIX }}^{\mathrm{IP}}\left(G_{n}\right)$ is unknown. An emblematic example (which was the initial motivation for our work) is the boolean hypercube $\mathbb{Z}_{2}^{n}$, for which Wilson [30] conjectured in 2004 that

$$
\begin{equation*}
t_{\mathrm{MIX}}^{\mathrm{IP}}\left(\mathbb{Z}_{2}^{n}\right) \asymp n . \tag{5}
\end{equation*}
$$

This was reiterated as Problem 4.2 of the AIM workshop Markov chains mixing times [26]. Here and throughout the paper, $\asymp$ and $\lesssim$ denote equality and inequality up to universal positive multiplicative constants. The current best estimates are

$$
\begin{equation*}
n \lesssim t_{\mathrm{MIX}}^{\mathrm{IP}}\left(\mathbb{Z}_{2}^{n}\right) \lesssim n \log n \tag{6}
\end{equation*}
$$

The lower bound is due to Wilson ([30], Section 9.1) and the upper bound was recently obtained by Alon and Kozma ([4], Corollary 10), as a special case of a much more general estimate which we will now discuss.
1.2. The big picture. An important observation about the IP is that the motion of a single particle is itself a Markov process. The generator is the usual graph Laplacian, which acts on functions $f: V_{G} \rightarrow \mathbb{R}$ by

$$
\begin{equation*}
\mathcal{L}_{G}^{\mathrm{RW}} f(x):=\sum_{y:\{x, y\} \in E_{G}}(f(y)-f(x)) . \tag{7}
\end{equation*}
$$

It is natural to expect the mixing properties of $\mathcal{L}_{G}^{\mathrm{IP}}$ and $\mathcal{L}_{G}^{\mathrm{RW}}$ to be intimately related. Indeed, a celebrated conjecture of Aldous, now resolved by Caputo, Liggett and Richthammer [8], asserts that the relaxation times (inverse spectral gaps) of these two operators coincide:

$$
\begin{equation*}
t_{\mathrm{REL}}^{\mathrm{P}}(G)=t_{\mathrm{REL}}^{\mathrm{RW}}(G) . \tag{8}
\end{equation*}
$$

Recall that $t_{\mathrm{REL}}^{\mathrm{RW}}(G)$ classically controls the mixing time $t_{\mathrm{MIX}}^{\mathrm{RW}}(G)$ of the single-particle dynamics (7), up to a correction which is only logarithmic in the number of vertices:

$$
\begin{equation*}
t_{\mathrm{MIX}}^{\mathrm{RW}}(G) \lesssim t_{\mathrm{REL}}^{\mathrm{RW}}(G) \log \left|V_{G}\right| . \tag{9}
\end{equation*}
$$

Inspired by the identity (8), Oliveira [27] conjectured that the same control applies to $t_{\mathrm{MIX}}^{\mathrm{IP}}(G)$. More precisely, he proposed the following simple-looking but far-reaching estimate, which is sharp in the three very different graph examples mentioned above (see Table 1).

Conjecture 1 (Oliveira [27]). For any connected graph $G$,

$$
\begin{equation*}
t_{\mathrm{MIX}}^{\mathrm{IP}}(G) \lesssim t_{\mathrm{REL}}^{\mathrm{RW}}(G) \log \left|V_{G}\right| \tag{10}
\end{equation*}
$$

TABLE 1
Some classical orders of magnitude

| $G$ | $\left\|V_{G}\right\|$ | $t_{\mathrm{REL}}^{\mathrm{RW}}(G)$ | $t_{\mathrm{MIX}}^{\mathrm{RW}}(G)$ | $t_{\mathrm{MIX}}^{\mathrm{IP}}(G)$ |
| :--- | :---: | :---: | :---: | :---: |
| $\mathcal{K}_{n}$ | $n$ | $1 / n$ | $1 / n$ | $(\log n) / n$ |
| $\mathcal{P}_{n}$ | $n$ | $n^{2}$ | $n^{2}$ | $n^{2} \log n$ |
| $\mathbb{Z}_{2}^{n}$ | $2^{n}$ | 1 | $\log n$ | $n$ |

Some partial progress on Conjecture 1 can be found in [16]. It is easy to prove that $t_{\text {rel }}^{\mathrm{RW}} \log |V|$ is comparable up to some universal constants to the mixing time of $|V|$ independent particles ([16], §2.1), and thus Oliveira's conjecture has the following probabilistic interpretation. It is saying that the mixing time of the interchange process is at most some universal constant multiple of the mixing time of $|V|$ independent particles (in fact, this is how it is phrased in [27]).

One of the most powerful techniques to bound the mixing time of a complicated Markov chain consists in comparing its Dirichlet form with that of a better understood chain having the same state space and stationary law, see the seminal paper by Diaconis and Saloff-Coste [11]. In the case of IP, the Dirichlet form is given by

$$
\begin{equation*}
\mathcal{E}_{G}^{\mathrm{IP}}(f):=\frac{1}{2\left|V_{G}\right|!} \sum_{\sigma \in \mathfrak{S}\left(V_{G}\right)} \sum_{e \in E_{G}}\left(f\left(\sigma \tau_{e}\right)-f(\sigma)\right)^{2} \tag{11}
\end{equation*}
$$

and a natural candidate for the comparison is the mean-field version $\mathcal{E}_{\mathcal{K}}^{\mathrm{IP}}$, where $\mathcal{K}$ denotes the complete graph on the same vertex set as $G$. Let us therefore define the comparison constant of the IP on $G$ as the smallest number $\chi_{G}^{\mathrm{IP}}$ such that the inequality

$$
\begin{equation*}
\mathcal{E}_{\mathcal{K}}^{\mathrm{IP}}(f) \leq \chi_{G}^{\mathrm{IP}} \mathcal{E}_{G}^{\mathrm{IP}}(f) \tag{12}
\end{equation*}
$$

holds for all $f: \mathfrak{S}\left(V_{G}\right) \rightarrow \mathbb{R}$. This constant is the optimal price to pay in order to systematically transfer quantitative estimates from IP on $\mathcal{K}$ to IP on $G$. In a recent breakthrough, Alon and Kozma [4], Theorem 1, established the following remarkably general estimate.

THEOREM 1.1 (Alon and Kozma [4]). For any regular connected graph $G$,

$$
\begin{equation*}
\chi_{G}^{\mathrm{IP}} \lesssim\left|V_{G}\right| t_{\mathrm{MIX}}^{\mathrm{RW}}(G) \tag{13}
\end{equation*}
$$

In particular, they deduced the following bound on the mixing times.
Corollary 1.2 (Alon and Kozma [4]). For any regular connected graph $G$,

$$
\begin{equation*}
t_{\mathrm{MIX}}^{\mathrm{IP}}(G) \lesssim t_{\mathrm{MIX}}^{\mathrm{RW}}(G) \log \left|V_{G}\right| \tag{14}
\end{equation*}
$$

Note that this proves Conjecture 1 along sequences $\left(G_{n}\right)_{n \geq 1}$ satisfying $t_{\text {MIX }}^{\mathrm{RW}}\left(G_{n}\right) \asymp$ $t_{\mathrm{REL}}^{\mathrm{RW}}\left(G_{n}\right)$. Examples include $\mathcal{K}_{n}, \mathcal{P}_{n}$, or the discrete tori $\mathbb{Z}_{n}, \mathbb{Z}_{n}^{2}, \overline{\mathbb{Z}}_{n}^{3}$, etc. On the other hand, for various other graphs such as the hypercube $\mathbb{Z}_{2}^{n}$ or bounded-degree expanders, one has

$$
\begin{equation*}
\frac{t_{\mathrm{REL}}^{\mathrm{RW}}\left(G_{n}\right)}{t_{\mathrm{MIX}}^{\mathrm{RW}}\left(G_{n}\right)} \xrightarrow[n \rightarrow \infty]{ } 0 \tag{15}
\end{equation*}
$$

and Theorem 1.1 fails at capturing the conjectured asymptotics. In light of this, the next step towards Conjecture 1 should naturally consist in understanding the mixing properties of the IP on graphs satisfying (15). This is precisely the program to which the present paper is intended to contribute.

## 2. Results.

2.1. Comparison constant and mixing time. A natural and important class of graphs satisfying (15) are the "high-dimensional" graphs obtained by taking cartesian products of a large number of small graphs. Recall that the cartesian product $G=G_{1} \times \cdots \times G_{n}$ of $n$ graphs $G_{1}, \ldots, G_{n}$ is the graph with vertex set $V_{G_{1}} \times \cdots \times V_{G_{n}}$, and where the neighbors of a vertex $x=\left(x_{1}, \ldots, x_{n}\right)$ are obtained by replacing an arbitrary coordinate $x_{i}(1 \leq i \leq n)$ with an arbitrary neighbor of $x_{i}$ in the graph $G_{i}$. Note that $G$ is connected as soon as $G_{1}, \ldots, G_{n}$
are. We will allow the dimension $n$ to grow arbitrarily but will keep the side-length fixed, meaning that

$$
\begin{equation*}
V_{G_{1}}=\cdots=V_{G_{n}}=\{1, \ldots, \ell\} \tag{16}
\end{equation*}
$$

for some fixed integer $\ell \geq 2$. Two simple examples are the $n$-dimensional torus $\mathbb{Z}_{\ell}^{n}=\mathbb{Z}_{\ell} \times$ $\cdots \times \mathbb{Z}_{\ell}$, and the $n$-dimensional Hamming graph $\mathcal{K}_{\ell}^{n}=\mathcal{K}_{\ell} \times \cdots \times \mathcal{K}_{\ell}$. In particular, when $\ell=2$, we recover the hypercube. Our main result is the determination of the exact order of magnitude of $\chi_{G}^{\mathrm{IP}}$ on all product graphs of fixed side-length.

THEOREM 2.1 (Comparison). All connected product graphs of side-length $\ell \geq 2$ satisfy

$$
\begin{equation*}
\chi_{G}^{\mathrm{IP}} \asymp \ell\left|V_{G}\right| \tag{17}
\end{equation*}
$$

where $\asymp_{\ell}$ means equality up to multiplicative constants that depend only on $\ell$.
Our estimate on $\chi_{G}^{\mathrm{IP}}$ classically yields an upper bound on the mixing time, even in the strong $L^{2}$ sense (see [4], Lemma 6). Moreover, a standard application of Wilson's method (see [18], Proposition 1.2) yields a matching lower bound. We thus obtain the following result, which confirms in particular Wilson's long-standing prediction (5).

Corollary 2.2 (Mixing time). All connected product graphs of side-length $\ell \geq 2$ satisfy

$$
\begin{equation*}
t_{\mathrm{MIX}}^{\mathrm{IP}}(G) \asymp \ell \log \left|V_{G}\right| \tag{18}
\end{equation*}
$$

Note that on product graphs, the single-particle dynamics (7) updates each coordinate independently. Consequently, any connected product graph of side-length $\ell \geq 2$ satisfies

$$
\begin{align*}
& t_{\mathrm{REL}}^{\mathrm{RW}}(G) \asymp \ell 1,  \tag{19}\\
& t_{\mathrm{MIX}}^{\mathrm{RW}}(G) \asymp \ell \log \log \left|V_{G}\right| . \tag{20}
\end{align*}
$$

(The double logarithm comes from the fact that there are $\log _{\ell}\left|V_{G}\right|$ coordinates, and that the time it takes to update all of them a constant number of times is logarithmic in the number of coordinates.) Thus, our Corollary 2.2 resolves Conjecture 1 for all product graphs of fixed side-length, in a regime where Theorem 1.1 always fails at doing so.

REMARK (Pre-cutoff). Let us comment on the constants hidden in our results, at least for the Hamming graph $\mathcal{K}_{\ell}^{n}$. Wilson's eigenvector method [30] produces the precise lower bound

$$
\begin{equation*}
\liminf _{n \rightarrow \infty}\left\{\frac{t_{\mathrm{MX}}^{\mathrm{P}}\left(\mathcal{K}_{\ell}^{n}\right)}{\ell^{n}}\right\} \geq c_{\ell} \tag{21}
\end{equation*}
$$

with the constant $c_{\ell}>0$ being completely explicit (e.g., we have $c_{2}=\frac{\ln 4}{2}$ ). On the other hand, our Corollary 2.2 guarantees that

$$
\begin{equation*}
\limsup _{n \rightarrow \infty}\left\{\frac{t_{\mathrm{MIX}}^{\mathrm{P}}\left(\mathcal{K}_{\ell}^{n}\right)}{\ell^{n}}\right\} \leq C_{\ell} \tag{22}
\end{equation*}
$$

for some constant $C_{\ell}<\infty$ that can certainly be made explicit as well, by a careful examination of our proof. However, we did not try to optimize the value of $C_{\ell}$, nor even to extract its rough dependency in $\ell$, because we believe that our comparison-based approach is inherently too rough to produce sharp constants anyway. Nevertheless, we note that neither Wilson's lower bound $c_{\ell}$ nor our upper bound $C_{\ell}$ change if we replace $\frac{1}{e}$ by any other precision $\varepsilon \in(0,1)$ in the definition (2), thereby establishing what is known as a pre-cutoff. Improving this to a true cutoff (i.e., $C_{\ell}=c_{\ell}$ ) remains a fascinating open problem.

We would like to close this section with a plausible extension of Theorem 2.1, inspired by an analogous result that we recently obtained for the Zero-Range Process [17], Corollary 3.

Conjecture 2 (General comparison). All finite connected graphs satisfy

$$
\begin{equation*}
\chi_{G}^{\mathrm{IP}} \lesssim\left|V_{G}\right| t_{\mathrm{REL}}^{\mathrm{RW}}(G) \tag{23}
\end{equation*}
$$

Note that a proof of this would immediately imply Conjecture 1.
2.2. Emergence of macroscopic cycles. One statistics of particular interest is the cycle structure of the random permutation $\xi_{t}$, as a function of the time $t$. On the infinite $d$-dimensional lattice $\mathbb{Z}^{d}$ with $d \geq 3$, a long-standing conjecture of Tóth [29] predicts a phase transition, indicated by the sudden emergence of infinite cycles at some critical time $t=t_{c} \in(0, \infty)$. This is related to a major open problem about the so-called quantum Heisenberg ferromagnet in statistical mechanics. To the best of our knowledge, the phase transition has only been proved on infinite regular trees [5, 15].

In the case of a large finite graph $G$, the relative lengths of cycles in a uniform random permutation asymptotically follow the Poisson-Dirichlet distribution (see, e.g., [28]). In particular, $\xi_{t}$ is likely to contain a macroscopic cycle at time $t \geq t_{\mathrm{MIX}}^{\mathrm{P}}(G)$. By analogy with Tóth's conjecture, one should however expect macroscopic cycles to emerge much before the mixing time. This was established in a precise sense by Schramm [28] in the mean-field case where $G=\mathcal{K}_{n}$, see also [6, 7]. Alas, results on other finite graphs are quite limited. In [3], Alon and Kozma obtained intriguing identities-involving the irreducible representations of the symmetric group-for the expected number of cycles of a given size in $\xi_{t}$ on any finite graph. Using these identities, they obtained a comparison-based estimate on the quantity

$$
\begin{equation*}
t_{\mathrm{CYC}}^{\mathrm{PP}}(G):=\min \left\{t \geq 0: \mathbb{P}\left(\xi_{t} \text { contains a cycle of length } \geq \frac{\left|V_{G}\right|}{2}\right) \geq \frac{1}{4}\right\} \tag{24}
\end{equation*}
$$

THEOREM 2.3 (Alon and Kozma [4]). All finite graphs G satisfy

$$
\begin{equation*}
t_{\mathrm{CYC}}^{\mathrm{IP}}(G) \lesssim \frac{\chi_{G}^{\mathrm{IP}}}{\left|V_{G}\right|} \tag{25}
\end{equation*}
$$

Thus, our main result implies that on high-dimensional graphs, macroscopic cycles do indeed emerge much before a single particle even mixes (recall (20)).

Corollary 2.4 (Giant cycles). All connected product graphs $G$ of side-length $\ell$ satisfy

$$
\begin{equation*}
t_{\mathrm{CYC}}^{\mathrm{IP}}(G) \lesssim \ell \tag{26}
\end{equation*}
$$

We note that Corollary 2.4 may not be sharp: it is actually quite possible that the macroscopic cycles already emerge at time $\Theta(1 / n)$ (where $n$ is the number of terms in the product), although proving this would require new ideas beyond the Alon-Kozma estimate (25). When specialized to the hypercube $\mathbb{Z}_{2}^{n}$, Corollary 2.4 complements a result of Koteckỳ, Miłoś and Ueltschi [20] regarding the appearance of mesoscopic cycles. It also complements a recent result by Adamczak, Kotowski and Miłoś [1], who established a phase transition for the emergence of macroscopic cycles on the 2-dimensional Hamming graph $\mathcal{K}_{n}^{2}$. Finally, we note that, by virtue of [4], Theorem 13, our main result also has direct implications on the magnetisation of the quantum Heisenberg ferromagnet.
2.3. Exclusion process. Another widely-studied interacting particle system is the exclusion process [14, 19, 24, 25]. For a finite graph $G$ and an integer $0<k<\left|V_{G}\right|$, the $k$-particle exclusion process (EX-k) on $G$ is a Markov chain on the set $\binom{V_{G}}{k}$ of $k$-element subsets of $V_{G}$, with generator given by

$$
\begin{equation*}
\mathcal{L}_{G}^{\mathrm{EX}-k} f(S):=\sum_{e \in \partial S}(f(S \oplus e)-f(S)), \tag{27}
\end{equation*}
$$

where $\oplus$ denotes the symmetric difference and $\partial S$ the edge-boundary of $S$ in $G$. This process describes the set occupied by $k$ fixed particles under the IP. More precisely, the EX-k $\left(\zeta_{t}\right)_{t \geq 0}$ with initial condition $S \in\binom{V_{G}}{k}$ can be constructed from the IP $\left(\xi_{t}\right)_{t \geq 0}$ by setting $\zeta_{t}:=\xi_{t}^{-1}(S)$. This observation, together with (8), easily implies that

$$
\begin{equation*}
\frac{t_{\mathrm{REL}}^{\mathrm{RW}}(G)}{t_{\mathrm{REL}}^{\mathrm{RW}}(\mathcal{K})} \leq \chi_{G}^{\mathrm{EX}-k} \leq \chi_{G}^{\mathrm{IP}} \tag{28}
\end{equation*}
$$

where $\mathcal{K}$ denotes the complete graph on $V_{G}$, and $\chi_{G}^{\mathrm{EX}-k}$ the optimal constant in the functional inequality $\mathcal{E}_{\mathcal{K}}^{\mathrm{EX}-k}(\cdot) \leq \chi_{G}^{\mathrm{EX}-k} \mathcal{E}_{G}^{\mathrm{EX}-k}(\cdot)$. Recalling (19) and the fact that $t_{\mathrm{REL}}^{\mathrm{RW}}(\mathcal{K})=1 /\left|V_{G}\right|$, we obtain the following corollary.

COROLLARY 2.5 (Comparison constant for EX-k). For all connected product graphs $G$ of side-length $\ell \geq 2$, and all $0<k<\left|V_{G}\right|$, we have $\chi_{G}^{\mathrm{EX}-k} \asymp_{\ell}\left|V_{G}\right|$.

As a consequence, one can transfer many quantitative estimates from $\mathcal{K}$ to $G$. This includes the inverse log-Sobolev constant $\rho_{G}^{\mathrm{EX}-k}$, defined as the smallest number such that

$$
\begin{equation*}
\mathbb{E}[f \log f]-\mathbb{E}[f] \log \mathbb{E}[f] \leq \rho_{G}^{\mathrm{EX}-k} \mathcal{E}_{G}^{\mathrm{EX}-k}(\sqrt{f}) \tag{29}
\end{equation*}
$$

for all $f:\binom{V_{G}}{k} \rightarrow(0, \infty)$, where $\mathbb{E}[\cdot]$ is expectation under the uniform law. This constant provides powerful controls on the underlying Markov semi-group [12]. It is easy to see that

$$
\begin{equation*}
\frac{1}{2} t_{\mathrm{REL}}^{\mathrm{RW}}(G) \leq \rho_{G}^{\mathrm{EX}-k} \leq \chi_{G}^{\mathrm{EX}-k} \rho_{\mathcal{K}}^{\mathrm{EX}-k} \tag{30}
\end{equation*}
$$

On the other hand, the log-Sobolev constant of the exclusion process on the complete graph (Bernoulli-Laplace model) was determined by Lee and Yau [22], Theorem 5:

$$
\begin{equation*}
\rho_{\mathcal{K}_{n}}^{\mathrm{EX}-k} \asymp \frac{1}{n} \log \frac{n^{2}}{k(n-k)} \tag{31}
\end{equation*}
$$

In particular, this allows us to pinpoint the exact order of $\rho_{G}^{\mathrm{EX}-k}$ in the dense-particle regime.
Corollary 2.6 (Log-Sobolev for EX-k). Fix $\varepsilon \in(0,1), \ell \geq 2$. Then, for all connected product graphs $G$ of side-length $\ell$ and all $k \in\left[\varepsilon\left|V_{G}\right|,(1-\varepsilon)\left|V_{G}\right|\right]$, we have $\rho_{G}^{\mathrm{EX}-k} \asymp_{\ell, \varepsilon} 1$.

Finally, we note that our main result also implies an upper bound of order $n$ (uniformly in $1 \leq k \leq 2^{n}$ ) on the $L^{2}$ mixing time of EX-k on the hypercube $\mathbb{Z}_{2}^{n}$, complementing a totalvariation estimate recently obtained by Hermon and Pymar [16] (as part of a much more general result).

## 3. Proof of the main result.

3.1. Proof outline. The lower bound in Theorem 2.1 is easy. Indeed, if $G$ is any finite graph and if $\mathcal{K}$ denotes the complete graph on $V_{G}$, then the very definition of $\chi_{G}^{\mathrm{IP}}$
implies

$$
\begin{equation*}
\chi_{G}^{\mathrm{IP}} \geq \frac{t_{\mathrm{REL}}^{\mathrm{IP}}(G)}{t_{\mathrm{REL}}^{\mathrm{P}}(\mathcal{K})}=\frac{t_{\mathrm{REL}}^{\mathrm{RW}}(G)}{t_{\mathrm{REL}}^{\mathrm{RW}}(\mathcal{K})}=\left|V_{G}\right| t_{\mathrm{REL}}^{\mathrm{RW}}(G) \tag{32}
\end{equation*}
$$

where the first equality uses (8). For a graph product of side-length $\ell$, we deduce

$$
\begin{equation*}
\chi_{G}^{\mathrm{IP}} \gtrsim \ell\left|V_{G}\right| \tag{33}
\end{equation*}
$$

The remainder of the paper is devoted to proving a matching upper bound. To do so, we combine four simple ideas, each one corresponding to a step of the proof.

1. Our first step consists in reducing the analysis of IP on a general $n$-dimensional graphproduct $G$ of side-length $\ell$ to the special case of the Hamming graph $\mathcal{K}_{\ell}^{n}$. This reduction relies on the classical method of canonical paths. An important simplification is that, by a standard path-lifting procedure, it is actually enough to just compare the single-particle on $G$ to that on $\mathcal{K}_{\ell}^{n}$. See Section 3.2 for details.
2. Our second step consists in re-interpreting the single-particle dynamics on $\mathcal{K}_{\ell}^{n}$ as a random walk on the additive group $\mathbb{Z}_{n}^{\ell}$, with the increment law $\mu$ being uniform over vectors with a single nonzero coordinate. This algebraic reformulation is performed in Section 3.3. It will allow one to use group-theoretical methods.
3. The third step consists in exploiting the celebrated octopus inequality [8], Theorem 2.3, to compare the IP with increment law $\mu$ to the IP with increment law $\mu^{\star t}=\mu \star \cdots \star \mu$ ( $t$-fold convolution), at a cost of order $t$. This is directly inspired by what Alon and Kozma did in [4]. However, instead of taking $t=\Theta(n \log n)$ so as to ensure that $\mu^{\star t}$ is close to uniform (all coordinates being refreshed with high probability), we crucially take $t=\Theta(n)$ only, with the prefactor being carefully adjusted so that only roughly half of the coordinates get refreshed under $\mu^{\star t}$. This important point is made rigorous by an application of the de Moivre-Laplace local limit theorem, see Section 3.4.
4. Finally, the last step consists in showing that, although the increment law $\mu^{\star t}$ is still very far from uniform (because of our choice of $t$ ), the associated Dirichlet form is actually comparable to the one with uniform increments. This is achieved by constructing canonical paths of length 2 , the underlying intuition being that randomizing all coordinates of a vector can be achieved by randomizing half the coordinates in one step, and the other half in a second step. This is described in Section 3.5.
3.2. Canonical paths. Our starting point is a powerful tool for comparing Dirichlet forms known as canonical paths; see, for example, [11]. As a warm-up, consider the single-particle dynamics (7) with Dirichlet form

$$
\begin{equation*}
\mathcal{E}_{G}^{\mathrm{RW}}(f):=\frac{1}{2\left|V_{G}\right|} \sum_{\{x, y\} \in E_{G}}(f(x)-f(y))^{2} \tag{34}
\end{equation*}
$$

As usual, a path in $G$ will be a finite sequence of vertices $\gamma=\left(\gamma_{0}, \ldots, \gamma_{k}\right)(k \geq 0)$ such that $e_{i}:=\left\{\gamma_{i-1}, \gamma_{i}\right\} \in E_{G}$ for each $1 \leq i \leq k$. We call $k$ the length of the path and denote it by $|\gamma|$. Also, we refer to $\gamma_{0}, \gamma_{k}$ as the endpoints of the path $\gamma$, and to $e_{1}, \ldots, e_{k}$ as the traversed edges. By a random path in $G$, we simply mean a random variable taking value in the (countable) set of all paths in $G$. We write $\mathbb{E}[\cdot]$ for the corresponding expectation.

ThEOREM 3.1 (Canonical paths, see e.g., [11]). Let $G, H$ be connected graphs on the same vertex set. For each edge $f \in E_{H}$, let $\gamma_{f}$ be any random path in $G$ with the same endpoints as $f$. Then, $\mathcal{E}_{H}^{\mathrm{RW}} \leq \kappa \mathcal{E}_{G}^{\mathrm{RW}}$ where $\kappa$ is the congestion, defined as follows:

$$
\begin{equation*}
\kappa=\max _{e \in E_{G}}\left\{\sum_{f \in E_{H}} \mathbb{E}\left[\left|\gamma_{f}\right| \mathbf{1}_{\left(\gamma_{f} \text { traverses } e\right)}\right]\right\} . \tag{35}
\end{equation*}
$$

We now make three elementary but important remarks.
REMARK (Trivial choice). Even in the worst-case situation where $H$ is the complete graph on $V_{G}$, we can always achieve the poor bound

$$
\begin{equation*}
\kappa \leq \frac{\left|V_{G}\right|^{3}}{4} \tag{36}
\end{equation*}
$$

by considering a spanning tree $T$ of $G$ and letting $\gamma_{f}$ be the unique simple path in $T$ connecting the endpoints of $f$. Note that this path is actually nonrandom. Exploiting randomness and the particular structure of $H$ to design paths with a low congestion is a matter of art.

REMARK (Congestion behaves well under products). If for $1 \leq i \leq n$, we can compare $G_{i}$ to $H_{i}$ with congestion $\kappa_{i}$, then we can compare $G_{1} \times \cdots \times G_{n}$ to $H_{1} \times \cdots \times H_{n}$ with congestion

$$
\begin{equation*}
\kappa=\max \left(\kappa_{1}, \ldots, \kappa_{n}\right), \tag{37}
\end{equation*}
$$

by considering paths that only evolve along a single coordinate, in the obvious way.
REMARK (Cayley graphs). Theorem 3.1 simplifies when $G=\operatorname{Cay}(\mathbb{G}, A)$ and $H=$ $\operatorname{Cay}(\mathbb{G}, B)$ are Cayley graphs generated by subsets $A, B$ of a finite group $\mathbb{G}$. Indeed, any word $\omega=\left(\omega_{1}, \ldots, \omega_{k}\right) \in A^{k}$ can be used to define a path

$$
\begin{equation*}
\gamma_{x, x b}:=\left(x, x \omega_{1}, \ldots, x \omega_{1} \cdots \omega_{k}\right) \tag{38}
\end{equation*}
$$

in $G$ from $x \in \mathbb{G}$ to $x b$, where $b=\omega_{1} \cdots \omega_{k}$ is the evaluation of $\omega$ in $\mathbb{G}$. Consequently, we only have to specify, for each $b \in B$, a random word $\omega_{b}$ over $A$ whose evaluation is $b$. Moreover, a straightforward computation shows that the resulting congestion is simply

$$
\begin{equation*}
\kappa=\max _{a \in A}\left\{\sum_{b \in B} \mathbb{E}\left[\left|\omega_{b}\right| N\left(a, \omega_{b}\right)\right]\right\} \tag{39}
\end{equation*}
$$

where $|\omega|$ denotes the length of a word $\omega$, and $N(a, \omega)$ the number of occurrences of $a$ in it.
Remark 3.2 applies in particular to the IP on any graph $G$. Indeed, one has

$$
\begin{equation*}
\mathcal{E}_{G}^{\mathrm{IP}}=\mathcal{E}_{\operatorname{Cay}(\mathbb{G}, A)}^{\mathrm{RW}}, \tag{40}
\end{equation*}
$$

with $\mathbb{G}=\mathfrak{S}\left(V_{G}\right)$ and $A=\left\{\tau_{e}: e \in E_{G}\right\}$. Moreover, any path in $G$ with endpoints $f=\{x, y\}$ and traversed edges $e_{1}, \ldots, e_{k}$ can be lifted to a word over $A$ that evaluates to $\tau_{f}$, namely:

$$
\begin{equation*}
\omega:=\left(\tau_{e_{1}}, \ldots, \tau_{e_{k-1}}, \tau_{e_{k}}, \tau_{e_{k-1}}, \ldots, \tau_{e_{1}}\right) \tag{41}
\end{equation*}
$$

Since the congestion is multiplied by at most 4 ( 2 for the length of the word, and 2 for the number of occurrences of a letter in it), we obtain the following classical result.

Corollary 3.2 (From canonical paths for RW to canonical paths for IP). Under the exact same assumptions (and notation) as in Theorem 3.1, we also have

$$
\begin{equation*}
\mathcal{E}_{H}^{\mathrm{IP}} \leq 4 \kappa \mathcal{E}_{G}^{\mathrm{IP}} \tag{42}
\end{equation*}
$$

Combining this with Remarks 3.2 and 3.2, we obtain the following inequality, which reduces the upper bound of Theorem 2.1 to the extremal case where $G$ is the ( $n, \ell$ )-Hamming graph:

$$
\begin{equation*}
\mathcal{K}_{\ell}^{n}:=\underbrace{\mathcal{K}_{\ell} \times \cdots \times \mathcal{K}_{\ell}}_{n \text { times }} . \tag{43}
\end{equation*}
$$

Corollary 3.3. For any n-dimensional connected product graph $G$ of side-length $\ell$,

$$
\begin{equation*}
\mathcal{E}_{\mathcal{K}_{\ell}^{n}}^{\mathrm{IP}} \leq \ell^{3} \mathcal{E}_{G}^{\mathrm{IP}} . \tag{44}
\end{equation*}
$$

In light of this result, the upper bound in Theorem 2.1 now boils down to the claim

$$
\begin{equation*}
\sup _{n \geq 1}\left\{\frac{\chi_{\mathcal{K}_{\ell}^{n}}^{\mathrm{IP}}}{\ell^{n}}\right\}<\infty \tag{45}
\end{equation*}
$$

for each $\ell \geq 2$, to which the remainder of the paper is devoted.
3.3. The octopus inequality. From now on, we fix the side-length $\ell \geq 2$ and the dimension $n \geq 1$. Writing $\mathcal{K}$ for the complete graph on $\{1, \ldots, \ell\}^{n}$, our goal is to establish the comparison

$$
\begin{equation*}
\mathcal{E}_{\mathcal{K}}^{\mathrm{IP}} \leq c \ell^{n} \mathcal{E}_{\mathcal{K}_{\ell}^{n}}^{\mathrm{IP}} \tag{46}
\end{equation*}
$$

where $c$ does not depend on $n$. We start by observing that the random walks on $\mathcal{K}_{\ell}^{n}$ and on $\mathcal{K}$ can both be conveniently viewed as random walks on the group

$$
\begin{equation*}
\mathbb{G}:=\mathbb{Z}_{\ell}^{n} \tag{47}
\end{equation*}
$$

equipped with coordinate-wise addition $\bmod \ell$ (which we will simply denote by + ). Given a probability measure $\mu$ on $\mathbb{G}$, we recall that the random walk with increment law $\mu$ has Dirichlet form

$$
\begin{equation*}
\mathcal{E}_{\mu}^{\mathrm{RW}}(f):=\frac{1}{2|\mathbb{G}|} \sum_{x, z \in \mathbb{G}} \mu(z)(f(x+z)-f(x))^{2}, \tag{48}
\end{equation*}
$$

for all $f: \mathbb{G} \rightarrow \mathbb{R}$. In particular, we have the representation

$$
\begin{equation*}
\mathcal{E}_{\mathcal{K}}^{\mathrm{RW}}=\left(\ell^{n}-1\right) \mathcal{E}_{\pi}^{\mathrm{RW}}, \quad \mathcal{E}_{\mathcal{K}_{\ell}^{n}}^{\mathrm{RW}}=n(\ell-1) \mathcal{E}_{\rho_{1}}^{\mathrm{RW}} \tag{49}
\end{equation*}
$$

where $\pi$ and $\rho_{k}(0 \leq k \leq n)$ respectively denote the uniform distributions on $\mathbb{G}$ and on

$$
\begin{equation*}
\mathbb{G}_{k}:=\{x \in \mathbb{G}:|\operatorname{supp}(x)|=k\} . \tag{50}
\end{equation*}
$$

Here $\operatorname{supp}(x)=\left\{i: x_{i} \neq 0\right\}$ naturally denotes the support of $x=\left(x_{1}, \ldots, x_{n}\right) \in \mathbb{G}$. Similarly, the IP on $\mathbb{G}$ with increment law $\mu$ has Dirichlet form

$$
\begin{equation*}
\mathcal{E}_{\mu}^{\mathrm{IP}}(f):=\frac{1}{2|\mathbb{G}|!} \sum_{\sigma \in \mathfrak{S}(\mathbb{G})} \sum_{x, z \in \mathbb{G}} \mu(z)\left(f\left(\sigma \tau_{\{x, x+z\}}\right)-f(\sigma)\right)^{2}, \tag{51}
\end{equation*}
$$

for $f: \mathfrak{S}(\mathbb{G}) \rightarrow \mathbb{R}$, with the interpretation $\tau_{\{x, x+z\}}=$ id when $z=0$. In view of (49) (with IP instead of RW), our claim (46) rewrites as

$$
\begin{equation*}
\mathcal{E}_{\pi}^{\mathrm{IP}} \leq c n \mathcal{E}_{\rho_{1}}^{\mathrm{IP}} \tag{52}
\end{equation*}
$$

for some (possibly different) constant $c<\infty$ that is only allowed to depend on $\ell$. The proof will crucially rely on the following elegant application of the octopus inequality [8], Theorem 2.3, which we borrow from Alon and Kozma [4]. We include a short proof as our setting here is slightly different. The convolution of two probability measures $\mu, v$ on $\mathbb{G}$ is defined by

$$
\begin{equation*}
(\mu \star v)(x):=\sum_{z \in \mathbb{G}} \mu(z) v(x-z) . \tag{53}
\end{equation*}
$$

Also, we say that a measure $\mu$ on $\mathbb{G}$ is symmetric if $\mu(z)=\mu(-z)$ for all $z \in \mathbb{G}$.

Lemma 3.4 (Comparison for convolutions). For any symmetric probability measure $\mu$ on $\mathbb{G}$,

$$
\begin{equation*}
\mathcal{E}_{\mu \star \mu}^{\mathrm{IP}} \leq 2 \mathcal{E}_{\mu}^{\mathrm{IP}} \tag{54}
\end{equation*}
$$

Proof. If $\mu(0)=0$, the octopus inequality [8], Theorem 2.3, asserts that

$$
\begin{aligned}
& \sum_{\sigma \in \mathfrak{S}(\mathbb{G}), z \in \mathbb{G}} \mu(z)\left(f\left(\sigma \tau_{\{x, x+z\}}\right)-f(\sigma)\right)^{2} \\
& \geq \frac{1}{2} \sum_{\sigma \in \mathfrak{S}(\mathbb{G}),(u, v) \in \mathbb{G}^{2}} \mu(u) \mu(v)\left(f\left(\sigma \tau_{\{x+u, x+v\}}\right)-f(\sigma)\right)^{2},
\end{aligned}
$$

for all $f: \mathfrak{S}(\mathbb{G}) \rightarrow \mathbb{R}$ and $x \in \mathbb{G}$ (the factor $\frac{1}{2}$ on the right-hand side compensates for the fact that we are here summing over all ordered pairs $(u, v) \in \mathbb{G}^{2}$ ). Averaging over all $x \in \mathbb{G}$, and applying the (bijective) change of variables $(x, u, v) \mapsto(x+u,-u, v-u)$ on the right-hand side, we obtain

$$
\begin{equation*}
2 \mathcal{E}_{\mu}^{\mathrm{IP}}(f) \geq \frac{1}{2|\mathbb{G}|!} \sum_{\sigma \in \mathfrak{S}(\mathbb{G}),(x, u, v) \in \mathbb{G}^{3}} \mu(-u) \mu(v-u)\left(f\left(\sigma \tau_{\{x, x+v\}}\right)-f(\sigma)\right)^{2} \tag{55}
\end{equation*}
$$

which is precisely $2 \mathcal{E}_{\mu}^{\mathrm{IP}}(f) \geq \mathcal{E}_{\mu \star \mu}^{\mathrm{IP}}(f)$ by symmetry of $\mu$. This proves the claim when $\mu(0)=$ 0 . In the general case, we write $\mu=(1-\theta) \rho_{0}+\theta v$ with $v(0)=0$, and we observe that

$$
\mathcal{E}_{\mu}^{\mathrm{IP}}=\theta \mathcal{E}_{v}^{\mathrm{IP}}, \quad \mathcal{E}_{\mu \star \mu}^{\mathrm{IP}}=\theta^{2} \mathcal{E}_{v \star v}^{\mathrm{IP}}+2 \theta(1-\theta) \mathcal{E}_{v}^{\mathrm{IP}}
$$

Thus, the claim $\mathcal{E}_{\mu \star \mu}^{\mathrm{IP}} \leq 2 \mathcal{E}_{\mu}^{\mathrm{IP}}$ is equivalent to $\mathcal{E}_{\nu \star \nu}^{\mathrm{IP}} \leq 2 \mathcal{E}_{v}^{\mathrm{IP}}$.
For reasons that will become clear later, we henceforth set

$$
\begin{align*}
t & :=2^{\left[\log _{2} n\right\rceil} \in[n, 2 n],  \tag{56}\\
\theta & :=n\left(1-e^{-\frac{\ln 2}{t}}\right) \in(0,1),  \tag{57}\\
p & :=\frac{\ell-1}{\ell} \tag{58}
\end{align*}
$$

Let us introduce the measure $\mu$ defined by

$$
\begin{equation*}
\mu:=(1-\theta p) \rho_{0}+\theta p \rho_{1} \tag{59}
\end{equation*}
$$

Since $\mu$ is symmetric and $t$ is a power of 2 , we may iterate Lemma 3.4 to get

$$
\begin{equation*}
\mathcal{E}_{\mu^{\star t}}^{\mathrm{IP}} \leq t \mathcal{E}_{\mu}^{\mathrm{IP}}=\theta p t \mathcal{E}_{\rho_{1}}^{\mathrm{IP}} \leq 2 n \mathcal{E}_{\rho_{1}}^{\mathrm{IP}} \tag{60}
\end{equation*}
$$

where $\mu^{\star t}=\mu \star \cdots \star \mu$ denotes the $t$-fold convolution of $\mu$. Thus, our goal (52) now boils down to showing that

$$
\begin{equation*}
\mathcal{E}_{\pi}^{\mathrm{IP}} \leq c \mathcal{E}_{\mu^{\star t}}^{\mathrm{IP}} \tag{61}
\end{equation*}
$$

for some constant $c<\infty$ that only depends on $\ell$. To this end, we analyze the convolution $\mu^{\star t}$ accurately using the de Moivre-Laplace local limit theorem.
3.4. Local limit theorem. As a warm-up, consider the binomial expansion of the uniform law: $\pi=\sum_{k=0}^{n} b_{k} \rho_{k}$, where

$$
\begin{equation*}
b_{k}:=\binom{n}{k}(1-p)^{k} p^{n-k} \tag{62}
\end{equation*}
$$

The classical de Moivre-Laplace local limit theorem provides uniform estimates on the coefficients $b_{0}, \ldots, b_{n}$. Although a specific value of $p$ was chosen at (58), the statement is of course valid for any $p \in(0,1)$.

THEOREM 3.5 (de Moivre-Laplace). There is $C<\infty$ depending only on $p$ such that

$$
\begin{equation*}
\left|b_{k}-\frac{e^{-\frac{x^{2}}{2}}}{\sqrt{2 \pi n p(1-p)}}\right| \leq \frac{C}{n^{3 / 2}} \tag{63}
\end{equation*}
$$

for all $0 \leq k \leq n$, with $x=(k-n p) / \sqrt{n p(1-p)}$.
We can use this to approximate $\mathcal{E}_{\pi}$ with $\mathcal{E}_{\rho_{I}}$, where $\rho_{I}$ is defined as follows:

$$
\begin{align*}
\rho_{I} & :=\frac{1}{|I|} \sum_{k \in I} \rho_{k} ;  \tag{64}\\
I & :=(n p-2 \sqrt{n p(1-p)}, n p+2 \sqrt{n p(1-p)}) \cap\{0, \ldots, n\} . \tag{65}
\end{align*}
$$

LEMMA 3.6 (Plateau proxy for $\mathcal{E}_{\pi}^{\mathrm{IP}}$ ). There is $c<\infty$ depending on $\ell$ only, such that

$$
\begin{equation*}
\mathcal{E}_{\pi}^{\mathrm{IP}} \leq c \mathcal{E}_{\rho_{I}}^{\mathrm{IP}} \tag{66}
\end{equation*}
$$

Proof. If $v$ is a symmetric distribution on $\mathbb{G}$, the Cauchy-Schwarz inequality yields

$$
\begin{aligned}
\sqrt{\frac{(v \star v)(x)}{\pi(x)}} & \geq \sum_{z \in \mathbb{G}} \sqrt{v(z) v(z-x)} \\
& \geq \sum_{z \in \mathbb{G}} v(z) \wedge v(z-x) \\
& \geq \sum_{z \in \mathbb{G}}\left[v(z)-\left(v(z)-\frac{1}{|\mathbb{G}|}\right)_{+}-\left(\frac{1}{|\mathbb{G}|}-v(z-x)\right)_{+}\right] \\
& =1-2 d_{\mathrm{TV}}(v, \pi)
\end{aligned}
$$

for all $x \in \mathbb{G}$, where $d_{\mathrm{TV}}(\cdot, \cdot)$ denotes the total-variation distance. In particular, when $d_{\mathrm{TV}}(\nu, \pi) \leq \frac{1}{4}$, we obtain $(\nu \star v)(x) \geq \pi(x) / 4$ for all $x \in \mathbb{G}$ and hence

$$
\begin{equation*}
\mathcal{E}_{\pi}^{\mathrm{IP}} \leq 4 \mathcal{E}_{v \star v}^{\mathrm{IP}} . \tag{67}
\end{equation*}
$$

Let us now apply this general observation to the restriction of $\pi$ to $\bigcup_{k \in I} \mathbb{G}_{k}$ :

$$
\begin{equation*}
v:=\frac{1}{q} \sum_{k \in I} b_{k} \rho_{k}, \quad q:=\sum_{k \in I} b_{k} \tag{68}
\end{equation*}
$$

Note that $d_{\mathrm{TV}}(v, \pi)=1-q$, and that $q \geq 3 / 4$ thanks to our definition of $I$ and Chebychev's inequality for the $\operatorname{Binomial}(n, p)$. Thus, (67) applies and yields

$$
\begin{aligned}
\mathcal{E}_{\pi}^{\mathrm{IP}} & \leq 4 \mathcal{E}_{\nu \star \nu}^{\mathrm{IP}} \\
& \leq 8 \mathcal{E}_{v}^{\mathrm{IP}}
\end{aligned}
$$

$$
\begin{aligned}
& =\frac{8}{q} \sum_{k \in I} b_{k} \mathcal{E}_{\rho_{k}}^{\mathrm{IP}} \\
& \leq \frac{32|I|}{3}\left(\max _{k \in I} b_{k}\right) \mathcal{E}_{\rho_{I}}^{\mathrm{IP}},
\end{aligned}
$$

where the second inequality uses Lemma 3.4, and the third $q \geq 3 / 4$. Finally, Theorem 3.5 ensures that $|I| \max _{k \in I} b_{k}$ is bounded by a quantity which only depends on $p$.

In order to establish (61), we will now approximate $\mu^{\star t}$ by the distribution

$$
\begin{align*}
\rho_{J} & :=\frac{1}{|J|} \sum_{k \in J} \rho_{k}  \tag{69}\\
J & :=\left(\frac{n p}{2}-2 \sqrt{n p(1-p)}, \frac{n p}{2}+2 \sqrt{n p(1-p)}\right) \cap\{0, \ldots, n\} . \tag{70}
\end{align*}
$$

Note that the center of $J$ is twice smaller than that of $I$.
Lemma 3.7 (Plateau proxy for $\mu^{\star t}$ ). There is $c>0$ depending on $\ell$ only, such that

$$
\begin{equation*}
\mu^{\star t} \geq c \rho_{J} \tag{71}
\end{equation*}
$$

Proof. The convolution with $\mu$ describes the following transformation on $\mathbb{G}$-valued random variables: pick one of the $n$ coordinates uniformly at random and, with probability $\theta$, replace it with a fresh uniform sample from $\mathbb{Z}_{\ell}$. Consequently, we may construct a random variable $X=\left(X_{1}, \ldots, X_{n}\right)$ with law $\mu^{\star t}$ by setting

$$
X_{i}:= \begin{cases}Z_{i} & \text { if } i \in\left\{U_{1}, \ldots, U_{N}\right\}  \tag{72}\\ 0 & \text { otherwise }\end{cases}
$$

where $N, U_{1}, \ldots, U_{t}, Z_{1}, \ldots, Z_{n}$ are independent random variables with the following laws:

- $N$ is binomial with parameters $t$ and $\theta$;
- $U_{1}, \ldots, U_{t}$ are uniform on $\{1, \ldots, n\}$;
- $Z_{1}, \ldots, Z_{n}$ are uniform on $\mathbb{Z}_{\ell}$.

In particular, setting $S:=|\operatorname{supp}(X)|$, we have

$$
\begin{equation*}
\mu^{\star t}=\sum_{k=0}^{n} \mathbb{P}(S=k) \rho_{k}, \tag{73}
\end{equation*}
$$

and our proof boils down to establishing that

$$
\begin{equation*}
\min _{k \in J} \mathbb{P}(S=k) \geq \frac{c}{\sqrt{n}} \tag{74}
\end{equation*}
$$

for some constant $c>0$ that only depends on $\ell$. Now, conditionally on $N$, the variable

$$
\begin{equation*}
R:=\left|\left\{U_{1}, \ldots, U_{N}\right\}\right| \tag{75}
\end{equation*}
$$

counts the number of distinct coupons collected by time $N$ in the standard coupon-collector problem of size $n$. Thus,

$$
\begin{equation*}
\mathbb{E}[R \mid N]=n\left(1-\left(1-\frac{1}{n}\right)^{N}\right), \quad \operatorname{Var}(R \mid N) \leq \frac{n}{4} \tag{76}
\end{equation*}
$$

Recalling our definitions (57), and since $N$ is a $\operatorname{Binomial}(t, \theta)$ variable, we easily deduce

$$
\begin{equation*}
\mathbb{E}[R]=n\left(1-\left(1-\frac{\theta}{n}\right)^{t}\right)=\frac{n}{2}, \quad \operatorname{Var}(R) \leq n \tag{77}
\end{equation*}
$$

Consequently, Chebychev's inequality yields

$$
\begin{equation*}
\mathbb{P}\left(R \in\left[\frac{n}{2}-2 \sqrt{n}, \frac{n}{2}+2 \sqrt{n}\right]\right) \geq \frac{3}{4} . \tag{78}
\end{equation*}
$$

Now, conditionally on $R$, the random variable $S$ is just a Binomial with parameters $R, p$. In particular, Theorem 2.1 with $R$ instead of $n$ ensures that

$$
\begin{equation*}
\min _{k \in J} \mathbb{P}\left(S=k \left\lvert\, R \in\left[\frac{n}{2}-2 \sqrt{n}, \frac{n}{2}+2 \sqrt{n}\right]\right.\right) \geq \frac{c}{\sqrt{n}} \tag{79}
\end{equation*}
$$

where $c>0$ only depends on $p$. Combining (78) and (79) establishes the claim.
3.5. Final comparison. In view of Lemmas 3.6 and 3.7, our objective (61) now reduces to establishing the following.

Proposition 3.8 (Final comparison). There exists $c<\infty$ depending only on $\ell$, such that

$$
\begin{equation*}
\mathcal{E}_{\rho_{I}}^{\mathrm{IP}} \leq c \mathcal{E}_{\rho_{J}}^{\mathrm{IP}} \tag{80}
\end{equation*}
$$

The crucial ingredient of the proof is the following lemma.
Lemma 3.9. For any $i, j \in\{0, \ldots, n\}$ with $i+j \in\{0, \ldots, n\}$, we have

$$
\begin{equation*}
\mathcal{E}_{\rho_{i+j}}^{\mathrm{IP}} \leq \frac{8\left|\mathbb{G}_{i}\right|}{\left|\mathbb{G}_{i}\right| \wedge\left|\mathbb{G}_{j}\right|} \mathcal{E}_{\rho_{i}}^{\mathrm{IP}}+\frac{8\left|\mathbb{G}_{j}\right|}{\left|\mathbb{G}_{i}\right| \wedge\left|\mathbb{G}_{j}\right|} \mathcal{E}_{\rho_{j}}^{\mathrm{IP}} . \tag{81}
\end{equation*}
$$

Proof. Let $(X, Y)$ denote a random element from the set

$$
\begin{equation*}
\left\{(x, y) \in \mathbb{G}_{i} \times \mathbb{G}_{j}: \operatorname{supp}(x) \cap \operatorname{supp}(y)=\varnothing\right\} \tag{82}
\end{equation*}
$$

Then $\omega:=(X, Y)$ is a random word of length 2 over $\mathbb{G}_{i} \cup \mathbb{G}_{j}$, whose evaluation $X+Y$ is uniform over $\mathbb{G}_{i+j}$. By Corollary 3.2 and Remark 3.2, we deduce that

$$
\begin{equation*}
\mathcal{E}_{\mathrm{Cay}\left(\mathbb{G}, \mathbb{G}_{i+j}\right)}^{\mathrm{IP}} \leq 4 \kappa \mathcal{E}_{\mathrm{Cay}\left(\mathbb{G}, \mathbb{G}_{i} \cup \mathbb{G}_{j}\right)}^{\mathrm{IP}}, \tag{83}
\end{equation*}
$$

where the congestion $\kappa$ is given by

$$
\begin{align*}
\kappa & =2\left|\mathbb{G}_{i+j}\right| \max _{b \in \mathbb{G}_{i} \cup \mathbb{G}_{j}}\{\mathbb{P}(X=b)+\mathbb{P}(Y=b)\}  \tag{84}\\
& =\frac{2\left|\mathbb{G}_{i+j}\right|}{\left|\mathbb{G}_{i}\right| \wedge\left|\mathbb{G}_{j}\right|}\left(1+\mathbf{1}_{(i=j)}\right) . \tag{85}
\end{align*}
$$

The second line follows from the observation that $X$ and $Y$ are uniformly distributed on $\mathbb{G}_{i}$ and $\mathbb{G}_{j}$, respectively. On the other hand, the definitions of $\rho_{i}, \rho_{j}, \rho_{i+j}$ imply

$$
\begin{align*}
\mathcal{E}_{\mathrm{Cay}\left(\mathbb{G}, \mathbb{G}_{i+j}\right)}^{\mathrm{IP}} & =\left|\mathbb{G}_{i+j}\right| \mathcal{E}_{\rho_{i+j}}^{\mathrm{IP}}  \tag{86}\\
\left(1+\mathbf{1}_{(i=j)}\right) \mathcal{E}_{\mathrm{Cay}\left(\mathbb{G}, \mathbb{G}_{i} \cup \mathbb{G}_{j}\right)}^{\mathrm{IP}} & =\left|\mathbb{G}_{i}\right| \mathcal{E}_{\rho_{i}}^{\mathrm{IP}}+\left|\mathbb{G}_{j}\right| \mathcal{E}_{\rho_{j}}^{\mathrm{IP}} . \tag{87}
\end{align*}
$$

The claim readily follows.

Proof of Proposition 3.8. Our definitions of $I$, $J$ ensure that $|I|=|J|$ and that

$$
\begin{equation*}
I \subseteq\{i+j:(i, j) \in \mathcal{P}\}, \text { where } \mathcal{P}:=\left\{(i, j) \in J^{2}: j \in\{i, i+1\}\right\} \tag{88}
\end{equation*}
$$

In particular, we have

$$
\begin{equation*}
\mathcal{E}_{\rho_{I}}^{\mathrm{IP}} \leq \frac{1}{|J|} \sum_{(i, j) \in \mathcal{P}} \mathcal{E}_{\rho_{i+j}}^{\mathrm{IP}} \tag{89}
\end{equation*}
$$

Now, since $\left|\mathbb{G}_{i}\right|=\binom{n}{i}(\ell-1)^{i}$ for all $i \in\{0, \ldots, n\}$, we have

$$
\begin{equation*}
\frac{\left|\mathbb{G}_{i+1}\right|}{\left|\mathbb{G}_{i}\right|}=\frac{(\ell-1)(n-i)}{(i+1)} \tag{90}
\end{equation*}
$$

As $i$ varies across $J$, this ratio remains bounded away from 0 and $\infty$ uniformly in $n$. Consequently, Lemma 3.9 ensures that for all $(i, j) \in \mathcal{P}$,

$$
\begin{equation*}
\mathcal{E}_{\rho_{i+j}}^{\mathrm{IP}} \leq c \mathcal{E}_{\rho_{i}}^{\mathrm{IP}}+c \mathcal{E}_{\rho_{j}}^{\mathrm{IP}} \tag{91}
\end{equation*}
$$

where $c<\infty$ depends only on $\ell$. Inserting this above, we obtain

$$
\begin{align*}
\mathcal{E}_{\rho_{I}}^{\mathrm{IP}} & \leq \frac{c}{|J|} \sum_{(i, j) \in \mathcal{P}}\left(\mathcal{E}_{\rho_{i}}^{\mathrm{IP}}+\mathcal{E}_{\rho_{j}}^{\mathrm{IP}}\right)  \tag{92}\\
& \leq \frac{4 c}{|J|} \sum_{j \in J} \mathcal{E}_{\rho_{j}}^{\mathrm{IP}}=4 c \mathcal{E}_{\rho_{J}}^{\mathrm{IP}} \tag{93}
\end{align*}
$$

This concludes the proof.

## REFERENCES

[1] AdAMcZak, R., Kotowski, M. and MiŁoś, P. (2018). Phase transition for the interchange and quantum Heisenberg models on the Hamming graph. Preprint. Available at arXiv:1808.08902.
[2] Aldous, D. and Fill, J. (2002). Reversible Markov chains and random walks on graphs. Unfinished manuscript. Available at http://www.stat.berkeley.edu/~aldous/RWG/book.html.
[3] Alon, G. and Kozma, G. (2013). The probability of long cycles in interchange processes. Duke Math. J. 162 1567-1585. MR3079255 https://doi.org/10.1215/00127094-2266018
[4] Alon, G. and Kozma, G. (2018). Comparing with octopi. Preprint. Available at arXiv:1811.10537.
[5] Angel, O. (2003). Random infinite permutations and the cyclic time random walk. In Discrete Random Walks (Paris, 2003). Discrete Math. Theor. Comput. Sci. Proc., AC 9-16. Assoc. Discrete Math. Theor. Comput. Sci., Nancy. MR2042369
[6] Berestycki, N. (2011). Emergence of giant cycles and slowdown transition in random transpositions and $k$-cycles. Electron. J. Probab. 16 152-173. MR2754801 https://doi.org/10.1214/EJP.v16-850
[7] Berestycki, N. and Kozma, G. (2015). Cycle structure of the interchange process and representation theory. Bull. Soc. Math. France $\mathbf{1 4 3}$ 265-280. MR3351179 https://doi.org/10.24033/bsmf. 2686
[8] Caputo, P., Liggett, T. M. and Richthammer, T. (2010). Proof of Aldous' spectral gap conjecture. J. Amer. Math. Soc. 23 831-851. MR2629990 https://doi.org/10.1090/S0894-0347-10-00659-4
[9] Diaconis, P. (1988). Group Representations in Probability and Statistics. Institute of Mathematical Statistics Lecture Notes-Monograph Series 11. IMS, Hayward, CA. MR0964069
[10] Diaconis, P. (1996). The cutoff phenomenon in finite Markov chains. Proc. Natl. Acad. Sci. USA 93 16591664. MR1374011 https://doi.org/10.1073/pnas.93.4.1659
[11] Diaconis, P. and Saloff-Coste, L. (1993). Comparison theorems for reversible Markov chains. Ann. Appl. Probab. 3 696-730. MR 1233621
[12] Diaconis, P. and Saloff-Coste, L. (1996). Logarithmic Sobolev inequalities for finite Markov chains. Ann. Appl. Probab. 6 695-750. MR1410112 https://doi.org/10.1214/aoap/1034968224
[13] Diaconis, P. and Shahshahani, M. (1981). Generating a random permutation with random transpositions. Z. Wahrsch. Verw. Gebiete 57 159-179. MR0626813 https://doi.org/10.1007/BF00535487
[14] Durrett, R., Liggett, T. M., Spitzer, F. and Sznitman, A.-S. (2012). Interacting Particle Systems at Saint-Flour. Probability at Saint-Flour. Springer, Heidelberg. MR3075635
[15] Hammond, A. (2015). Sharp phase transition in the random stirring model on trees. Probab. Theory Related Fields 161 429-448. MR3334273 https://doi.org/10.1007/s00440-013-0543-7
[16] Hermon, J. and Pymar, R. (2018). The exclusion process mixes (almost) faster than independent particles. Ann. Probab. To appear. Preprint available at arXiv:1808.10846.
[17] Hermon, J. and SAlez, J. (2019). A version of Aldous' spectral-gap conjecture for the zero range process. Ann. Appl. Probab. 29 2217-2229. MR3984254 https://doi.org/10.1214/18-AAP1449
[18] Jonasson, J. (2012). Mixing times for the interchange process. ALEA Lat. Am. J. Probab. Math. Stat. 9 667-683. MR3069380
[19] Kipnis, C. and Landim, C. (1999). Scaling Limits of Interacting Particle Systems. Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences] 320. Springer, Berlin. MR1707314 https://doi.org/10.1007/978-3-662-03752-2
[20] Kotecký, R., MiŁoś, P. and Ueltschi, D. (2016). The random interchange process on the hypercube. Electron. Commun. Probab. 21 Paper No. 4. MR3485373 https://doi.org/10.1214/16-ECP4540
[21] Lacoin, H. (2016). Mixing time and cutoff for the adjacent transposition shuffle and the simple exclusion. Ann. Probab. 44 1426-1487. MR3474475 https://doi.org/10.1214/15-AOP1004
[22] LEE, T.-Y. and YAU, H.-T. (1998). Logarithmic Sobolev inequality for some models of random walks. Ann. Probab. 26 1855-1873. MR1675008 https://doi.org/10.1214/aop/1022855885
[23] Levin, D. A. and Peres, Y. (2017). Markov Chains and Mixing Times. Amer. Math. Soc., Providence, RI. Second edition of [MR2466937]. MR3726904
[24] Liggett, T. M. (1999). Stochastic Interacting Systems: Contact, Voter and Exclusion Processes. Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences] 324. Springer, Berlin. MR1717346 https://doi.org/10.1007/978-3-662-03990-8
[25] Liggett, T. M. (2005). Interacting Particle Systems. Classics in Mathematics. Springer, Berlin. Reprint of the 1985 original. MR2108619 https://doi.org/10.1007/b138374
[26] Oliveira, R. (2016). AimPL: Markov chain mixing times. Available at http://aimpl.org/markovmixing.
[27] Oliveira, R. I. (2013). Mixing of the symmetric exclusion processes in terms of the corresponding singleparticle random walk. Ann. Probab. 41 871-913. MR3077529 https://doi.org/10.1214/11-AOP714
[28] Schramm, O. (2005). Compositions of random transpositions. Israel J. Math. 147 221-243. MR2166362 https://doi.org/10.1007/BF02785366
[29] TóTh, B. (1993). Improved lower bound on the thermodynamic pressure of the spin $1 / 2$ Heisenberg ferromagnet. Lett. Math. Phys. 28 75-84. MR1224836 https://doi.org/10.1007/BF00739568
[30] Wilson, D. B. (2004). Mixing times of Lozenge tiling and card shuffling Markov chains. Ann. Appl. Probab. 14 274-325. MR2023023 https://doi.org/10.1214/aoap/1075828054
[31] YAU, H.-T. (1997). Logarithmic Sobolev inequality for generalized simple exclusion processes. Probab. Theory Related Fields 109 507-538. MR1483598 https://doi.org/10.1007/s004400050140

