



HAL
open science

A sharp log-Sobolev inequality for the multislice

Justin Salez

► **To cite this version:**

Justin Salez. A sharp log-Sobolev inequality for the multislice. *Annales Henri Lebesgue*, 2021, 4, pp.1143-1161. <10.5802/ahl.99>. <hal-03353026>

HAL Id: hal-03353026

<https://hal.science/hal-03353026v1>

Submitted on 6 Jan 2025

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.



HAL Authorization

A sharp log-Sobolev inequality for the multislice

Justin Salez*

Abstract

We determine the log-Sobolev constant of the multi-urn Bernoulli-Laplace diffusion model with arbitrary parameters, up to a small universal multiplicative constant. Our result extends a classical estimate of Lee and Yau (1998) and confirms a conjecture of Filmus, O’Donnell and Wu (2018). Among other applications, we completely quantify the “small-set expansion” phenomenon on the multislice, and obtain sharp mixing-time estimates for the colored exclusion process on various graphs.

Contents

1	Introduction	2
1.1	The multislice	2
1.2	Functional inequalities	3
1.3	Related works	5
2	Results	7
2.1	Main estimate	7
2.2	Small-set expansion	8
2.3	Colored exclusion process	9
3	Proofs	12
3.1	General strategy	12
3.2	Main recursion	14
3.3	Putting things together	15
3.4	Coarsening argument	18

*CEREMADE, CNRS, UMR 7534, Université Paris-Dauphine, PSL University, 75016 Paris, France

1 Introduction

1.1 The multislice

Consider a sequence of positive integers $\kappa = (\kappa_1, \dots, \kappa_L)$ of arbitrary length $L \geq 2$, and set

$$n = \kappa_1 + \dots + \kappa_L. \quad (1)$$

We will refer to the elements of $[L] = \{1, \dots, L\}$ as *colors*, and write Ω_κ for the set of $[L]$ -valued sequences in which each color $\ell \in [L]$ appears exactly κ_ℓ times:

$$\Omega_\kappa := \left\{ \omega = (\omega_1, \dots, \omega_n) \in [L]^n : \sum_{i=1}^n \mathbf{1}_{(\omega_i=\ell)} = \kappa_\ell \text{ for each } \ell \in [L] \right\}.$$

This natural combinatorial set is sometimes called a *multislice*. It provides a canonical interpretation for the classical multinomial coefficient:

$$|\Omega_\kappa| = \binom{n}{\kappa_1, \dots, \kappa_L}.$$

The symmetric group \mathfrak{S}_n acts transitively on the multislice in the obvious way, by permuting coordinates. In particular, transpositions induce a natural local random walk on Ω_κ , which consists in repeatedly picking two positions $1 \leq i < j \leq n$ uniformly at random and replacing the current state $\omega \in \Omega_\kappa$ with the new state

$$\omega^{ij} := (\omega_1, \dots, \omega_{i-1}, \omega_j, \omega_{i+1}, \dots, \omega_{j-1}, \omega_i, \omega_{j+1}, \dots, \omega_n).$$

This Markov chain is known as the *transposition walk on the multislice*, or *multi-urn Bernoulli-Laplace diffusion model* with parameter κ . It can also be viewed as a random walk on the Schreier graph $\mathbb{G}_\kappa = (\Omega_\kappa, E_\kappa)$, whose edge-set is given by

$$E_\kappa := \left\{ \{\omega, \omega'\} \subseteq \Omega_\kappa : \sum_{i=1}^n \mathbf{1}_{(\omega_i \neq \omega'_i)} = 2 \right\}.$$

Thanks to the degree of freedom in the choice of the parameter κ , the model is rich enough to encompass several classical special cases, including:

- (i) the random walk on the complete graph of order n , corresponding to $\kappa = (1, n-1)$;
- (ii) the k -particle Bernoulli-Laplace diffusion on n sites, corresponding to $\kappa = (k, n-k)$;

(iii) the transposition walk on \mathfrak{S}_n , corresponding to $\kappa = (1, \dots, 1)$.

These fundamental examples have been studied in full detail, see in particular [12, 13, 2, 18, 33, 38, 30, 39]. In the general case, however, understanding the precise impact of the parameter κ on the mixing properties of the graph \mathbb{G}_κ was suggested as an open problem several times [13, 9, 15]. Beyond the traditional “mixing times of Markov chains” perspective, this question was recently shown in [16, 15, 14] to have remarkable applications to the theory of Boolean functions on the multislice, see Section 2.1 below for more details. In particular, the present paper was motivated by a conjecture from [16] regarding the so-called *log-Sobolev constant* of the multislice, whose definition will be recalled in the next section.

Remark 1 (Coarsening). *There is an obvious partial ordering on our parameter space: say that κ' is coarser than κ if it can be obtained from κ by repeatedly merging two entries into one. Note that this operation simply amounts to identifying certain colors, so that the transposition walk on $\Omega_{\kappa'}$ is a projection of the one on Ω_κ . In particular, the mixing behavior of the chain can only improve as κ becomes coarser, with the case $\kappa = (1, \dots, 1)$ of example (iii) being the worst. Our main result will precisely quantify this qualitative statement.*

1.2 Functional inequalities

One of the most powerful ways to quantify the mixing properties of a Markov chain consists in establishing appropriate functional inequalities for the underlying Dirichlet form. We shall here only recall the relevant definitions, and refer to the seminal papers [8, 3] or the excellent survey [34] for a detailed account. We start by turning the multislice Ω_κ into a probability space by equipping it with the uniform distribution. In particular, we regard functions $f: \Omega_\kappa \rightarrow \mathbb{R}$ as random variables, and write $\mathbb{E}_\kappa[f]$ for the corresponding expectation:

$$\mathbb{E}_\kappa[f] := \frac{1}{|\Omega_\kappa|} \sum_{\omega \in \Omega_\kappa} f(\omega).$$

The *Dirichlet form* of our chain is defined for every $f, g: \Omega_\kappa \rightarrow \mathbb{R}$ by

$$\mathcal{E}_\kappa(f, g) := \frac{1}{2n} \sum_{1 \leq i < j \leq n} \mathbb{E}_\kappa [(\nabla^{ij} f) (\nabla^{ij} g)], \quad (2)$$

where $(\nabla^{ij} f)(\omega) := f(\omega^{ij}) - f(\omega)$ is the discrete gradient.

Remark 2 (Scaling). *We have here chosen to work under the natural continuous-time scaling where each of the $\binom{n}{2}$ possible transpositions occurs at rate $1/n$, so that a coordinate gets refreshed at rate 1. We emphasize that this is a matter of convention only: switching to discrete time amounts to nothing more than multiplying the above Dirichlet form by $2/(n-1)$.*

Since $\mathcal{E}_\kappa(f, f)$ measures the *local* variation of the observable f along a typical transition of the chain, it is natural to compare it with the *variance* $\text{Var}_\kappa(f)$ or the *entropy* $\text{Ent}_\kappa(f)$, which quantify the *global* variation of f across the whole state space:

$$\begin{aligned}\text{Var}_\kappa(f) &:= \mathbb{E}_\kappa[f^2] - \mathbb{E}_\kappa[f]^2, \\ \text{Ent}_\kappa(f) &:= \mathbb{E}_\kappa[f \log f] - \mathbb{E}_\kappa[f] \log \mathbb{E}_\kappa[f].\end{aligned}$$

All logs appearing in this paper are natural logarithms, and the last definition is of course restricted to non-negative functions, with the standard convention $0 \log 0 = 0$. With this notation in hands, the three classical functional inequalities read as follows:

- The *Poincaré inequality* holds with constant τ if

$$\text{Var}_\kappa(f) \leq \tau \mathcal{E}_\kappa(f, f), \quad \text{for all } f: \Omega_\kappa \rightarrow \mathbb{R}. \quad (3)$$

- The *modified log-Sobolev inequality* holds with constant τ if

$$\text{Ent}_\kappa(f) \leq \tau \mathcal{E}_\kappa(f, \log f), \quad \text{for all } f: \Omega_\kappa \rightarrow \mathbb{R}_+. \quad (4)$$

- The *log-Sobolev inequality* holds with constant τ if

$$\text{Ent}_\kappa(f) \leq \tau \mathcal{E}_\kappa(\sqrt{f}, \sqrt{f}), \quad \text{for all } f: \Omega_\kappa \rightarrow \mathbb{R}_+. \quad (5)$$

The optimal values of τ in these functional inequalities are respectively known as the (inverse) *Poincaré*, *modified log-Sobolev*, and *log-Sobolev* constants of the chain. They will here be denoted by $\tau_{\text{REL}}(\kappa)$, $\tau_{\text{MLS}}(\kappa)$ and $\tau_{\text{LS}}(\kappa)$. These fundamental parameters provide powerful controls on the underlying Markov semi-group, and have tight connections to mixing times, concentration of measure, small-set expansion, and hypercontractivity. We again refer to [8, 3, 34] for a detailed account, and to [19] for new characterizations. Let us simply note that the statements (3), (4), (5) are essentially increasing in strength, in the sense that

$$2\tau_{\text{REL}}(\kappa) \leq 4\tau_{\text{MLS}}(\kappa) \leq \tau_{\text{LS}}(\kappa). \quad (6)$$

Perhaps surprisingly, the first two quantities turn out to be too rough to capture the precise impact of κ on the mixing properties of the multislice Ω_κ . Specifically, we note the following dramatic insensitivity result, see Section 3.4 for details.

Lemma 1 (Insensitivity of the Poincaré and modified log-Sobolev constants). *We have*

$$\tau_{\text{REL}}(\kappa) = 1 \quad \text{and} \quad \tau_{\text{MLS}}(\kappa) \in \left[\frac{1}{2}, 2 \right],$$

regardless of the choice of the parameter κ .

In contrast, the much finer *log-Sobolev constant* $\tau_{\text{LS}}(\kappa)$ happens to depend on κ in a non-trivial way, and understanding the exact nature of this dependency is precisely the aim of the present paper. Before we state our results, let us give a brief account on this general problem and its broad range of applications.

1.3 Related works

As already mentioned, the multi-urn Bernoulli-Laplace model encompasses various well-studied special cases. The simplest one is the random walk on the complete n -vertex graph, obtained with $\kappa = (1, n-1)$. This example belongs to the short list of chains whose log-Sobolev constant is known exactly, see the seminal paper [8] by Diaconis and Saloff-Coste.

Theorem 1 (Random walk on the complete graph, see Theorem A.1 in [8]).

$$\tau_{\text{LS}}(1, n-1) = \begin{cases} \frac{n \log(n-1)}{n-2} & \text{if } n \geq 3 \\ 2 & \text{if } n = 2. \end{cases}$$

A much richer example is the famous “Random Transposition” walk on the symmetric group \mathfrak{S}_n , which corresponds to the choice $\kappa = (1, \dots, 1)$. A sharp estimate on the log-Sobolev constant of this fundamental chain can be deduced from the detailed representation-theoretic analysis conducted by Diaconis and Shahshahani in their pioneering work [12].

Theorem 2 (Random transposition on the symmetric group, see [12]).

$$\log n \leq \tau_{\text{LS}}(\underbrace{1, \dots, 1}_{n \text{ times}}) \leq 4 \log n.$$

Several years later, Lee and Yau found a more direct proof, based on what is now known as the “martingale method” [30]. This approach also allowed them to determine the order of magnitude of the log-Sobolev constant of the k -particle Bernoulli-Laplace diffusion on n sites, thereby resolving an open problem raised by Diaconis and Saloff-Coste in [8].

Theorem 3 (Two-urn Bernoulli-Laplace diffusion model, see Theorem 5 in [30]). *There exists a universal constant $\varepsilon > 0$ such that for all $0 < k < n$,*

$$\varepsilon \log \left(\frac{n^2}{k(n-k)} \right) \leq \tau_{\text{LS}}(k, n-k) \leq \frac{2}{\log 2} \log \left(\frac{n^2}{k(n-k)} \right).$$

The implications of Theorems 2-3 are too numerous to be all cited. One particularly active direction consists in “transferring” these log-Sobolev estimates to models with less symmetry in order to obtain sharp mixing-time bounds, via the celebrated “comparison method” introduced by Diaconis and Saloff-Coste [10, 11]. Recent successful examples include the *interchange process* on arbitrary graphs [1], or the *exclusion process* on high-dimensional product graphs [24]. Beyond Markov chains, the well-known connection between log-Sobolev inequalities and *hypercontractivity* provides another extremely fertile ground for applications in discrete analysis and computer science. We refer to the book [36, Chapters 9 & 10] for details, and to the recent work [16] for an impressive list of references from combinatorics, computational learning, property testing or Boolean functions, where Theorems 2-3 played a crucial role. Motivated by these applications, Filmus, O’Donnell and Wu [16] initiated the study of the log-Sobolev constant $\tau_{\text{LS}}(\kappa)$ for general κ . Their main result is as follows.

Theorem 4 (General bound, see Theorem 1 in [16]). *For any choice of the parameter κ ,*

$$\tau_{\text{LS}}(\kappa) \leq \frac{2}{\log 2} \sum_{\ell=1}^L \log \left(\frac{4n}{\kappa_{\ell}} \right).$$

Several remarkable consequences of this estimate can be found in the recent works [16, 14]. A quick comparison with Theorems 1, 2 and 3 shows that the bound is of the right order of magnitude in the extreme case $L = 2$, but is off by a factor of order n at the other extreme, $L = n$. Regarding what the correct order of magnitude of $\tau_{\text{LS}}(\kappa)$ should be for *all* ranges of κ , Filmus, O’Donnell and Wu proposed the following beautifully simple dependency.

Conjecture 1 (See page 3 in [16]). *For any choice of the parameter κ ,*

$$\tau_{\text{LS}}(\kappa) \asymp \log \left(\frac{n}{\kappa_{\text{MIN}}} \right),$$

where $\kappa_{\text{MIN}} := \min\{\kappa_1, \dots, \kappa_L\}$ and where \asymp means equality up to universal pre-factors.

Note that the right-hand side decreases smoothly from $\log n$ down to 0 as κ becomes coarser and coarser, in agreement with Remark 1. To better appreciate this conjecture,

consider the *single-site dynamics* obtained by projecting the multislice onto a fixed coordinate $i \in [n]$: under our transposition walk, the variable ω_i simply gets refreshed at unit rate according to the marginal distribution

$$\mathbb{P}_\kappa(\omega_i = \ell) = \frac{\kappa_\ell}{n}, \quad \ell \in [L].$$

The log-Sobolev constant of this trivial chain is well-known to be

$$\tau_{\text{LS}}^{\text{triv}}(\kappa) = \frac{n}{n - 2\kappa_{\text{MIN}}} \log \left(\frac{n}{\kappa_{\text{MIN}}} - 1 \right) \asymp \log \left(\frac{n}{\kappa_{\text{MIN}}} \right),$$

see [8, Theorem A.1]. Although our probability space Ω_κ is far from being a product space, the above conjecture asserts that the transposition walk mixes essentially as well as if the coordinates $\omega_1, \dots, \omega_n$ were being refreshed independently. A brief look at Theorems 1, 2 and 3 will convince the reader that this intuition is correct in all known special cases.

2 Results

2.1 Main estimate

Our main result is the determination of the log-Sobolev constant $\tau_{\text{LS}}(\kappa)$ for all values of the parameter κ , up to a (small) universal multiplicative constant.

Theorem 5 (The log-Sobolev constant of the multislice). *For all values of κ ,*

$$\log \left(\frac{n}{\kappa_{\text{MIN}}} \right) \leq \tau_{\text{LS}}(\kappa) \leq \frac{4}{\log 2} \log \left(\frac{n}{\kappa_{\text{MIN}}} \right).$$

This confirms Conjecture 1. We note that the improvement upon Theorem 4 can be considerable if the dimension L is large. Specifically, the upper bound of Filmus, O’Donnell and Wu is *always* super-linear in L , since the convexity of $t \mapsto t \log t$ yields

$$\sum_{\ell=1}^L \log \left(\frac{n}{\kappa_\ell} \right) \geq L \log L,$$

for any choice of the parameter κ . In contrast, our result shows that

$$\tau_{\text{LS}}(\kappa) \asymp \log L, \tag{7}$$

as long as the vector $\kappa = (\kappa_1, \dots, \kappa_L)$ is reasonably *balanced*, in the (weak) sense that its lowest entry is of the same order as the mean entry. In particular, our estimate can be

readily used to sharpen the dependency in L in the various quantitative results that were derived from Theorem 4 in [16]. To avoid a lengthy detour through hypercontractivity, we choose to leave the details to the reader, and to instead describe two different applications: a sharp quantification of the “small-set expansion” phenomenon for the multislice, and a general log-Sobolev inequality for the colored exclusion processes.

Remark 3 (Sharpness of universal constants). *In our lower bound, the pre-factor in front of the logarithm can not be replaced by any larger universal constant, since we have*

$$\tau_{\text{LS}}(\kappa) = (1 + o(1)) \log \left(\frac{n}{\kappa_{\text{MIN}}} \right),$$

in the special case $\kappa = (1, n-1)$, as per Theorem 1. Regarding the upper bound, our pre-factor can not be improved by more than a $\log 2$ factor. Indeed, we will show that

$$\tau_{\text{LS}}(\kappa) \geq (4 - o(1)) \log \left(\frac{n}{\kappa_{\text{MIN}}} \right),$$

in the important special case $\kappa = (\lfloor n/2 \rfloor, \lceil n/2 \rceil)$, see (26). In fact, the possibly loose $\log 2$ term comes directly from the one appearing in Theorem 3, and any improvement of the latter will immediately imply the same improvement in our upper bound.

2.2 Small-set expansion

Recall that the multislice is naturally equipped with a graph structure by declaring two vertices $\omega, \omega' \in \Omega_\kappa$ to be adjacent if they differ at exactly two coordinates. Following standard graph-theoretical notation, we write $|\partial A|$ for the *edge boundary* of a subset $A \subseteq \Omega_\kappa$, i.e., the set of edges having one end-point in A and the other outside A . Let us consider the problem of finding a constant $\iota(\kappa)$, as large as possible, such that the isoperimetric inequality

$$\frac{|\partial A|}{|A|} \geq \iota(\kappa) \log \left(\frac{|\Omega_\kappa|}{|A|} \right), \tag{8}$$

holds for all non-empty subsets $A \subseteq \Omega_\kappa$. The left-hand side measures the *conductance* of A , i.e. the facility for the walk to escape from A , given that it currently lies in A . The presence of the logarithmic term on the other side constitutes a notable improvement upon the more standard *Cheeger inequality*: instead of being constant, the right-hand side of (8) gets larger as the set A gets smaller, thereby capturing the celebrated *small-set expansion* phenomenon [29, 31, 16]. Our log-Sobolev estimate allows us to determine the fundamental quantity $\iota(\kappa)$ for all values of κ , up to a small universal constant.

Corollary 1 (Small-set expansion for the multislice). *The optimal constant in (8) satisfies*

$$\frac{\log 2}{4} \frac{n}{\log\left(\frac{n}{\kappa_{\text{MIN}}}\right)} \leq \iota(\kappa) \leq \frac{n}{\log\left(\frac{n}{\kappa_{\text{MIN}}}\right)}. \quad (9)$$

The proof will be given in Section 3.3. As in Remark 3, the universal constants appearing in our estimate can not be improved, apart from perhaps removing the $\log 2$ term.

2.3 Colored exclusion process

A far-reaching generalization of the transposition walk on the multislice Ω_κ consists in allowing each of the $\binom{n}{2}$ possible transpositions to occur at a different (possibly zero) rate. More precisely, we fix a non-negative symmetric array $G = (G_{ij})_{1 \leq i, j \leq n}$ (which we interpret as a weighted graph) and consider the following weighted version of the Dirichlet form (2):

$$\mathcal{E}_\kappa^G(f, g) := \frac{1}{2} \sum_{1 \leq i < j \leq n} G_{ij} \mathbb{E}_\kappa [(\nabla^{ij} f)(\nabla^{ij} g)]. \quad (10)$$

The canonical setting – to which we shall here stick for simplicity – consists in taking G to be the transition matrix of the simple random walk on a regular graph, which we henceforth identify with G . The resulting process is known as the κ -colored exclusion process on G , see [4]. By varying the parameter κ , we obtain a rich family of diffusion models on G including:

- (i) the *simple random walk* on G , when $\kappa = (1, n - 1)$;
- (ii) the *k -particle exclusion process* on G , when $\kappa = (k, n - k)$;
- (iii) the *interchange process* on G , when $\kappa = (1, \dots, 1)$.

Comparing the mixing properties of these three processes constitutes a rich and active research problem, see [40, 35, 4, 28, 37, 6, 20, 1]. Perhaps the most celebrated result in this direction is the remarkable fact that their Poincaré constants coincide, as conjectured by Aldous and established by Caputo, Liggett and Richthammer [4].

Theorem 6 (Insensitivity of the Poincaré constant, see [4]). *The Poincaré constant $\tau_{\text{REL}}(\kappa, G)$ of the κ -colored exclusion process on G does not depend on κ . In particular, it equals the Poincaré constant $\tau_{\text{REL}}(G)$ of the simple random walk on G .*

In a sense, this result asserts that the Poincaré constant is too “rough” to capture the influence of the color profile κ on the mixing properties of the colored exclusion process. It is thus natural to turn one’s attention to the finer log-Sobolev constant.

Question 1. *How does the log-Sobolev constant $\tau_{\text{LS}}(\kappa, G)$ depend upon the parameter κ ?*

Our main result answers this question in the simple *mean-field* setting, where G is the complete graph. However, it implies an estimate of $\tau_{\text{LS}}(\kappa, G)$ for arbitrary G , by means of the celebrated “comparison method” introduced by Diaconis and Saloff-Coste [10, 11]. A particularly pleasant observation here is that we do not even need to build a comparison theory for the colored exclusion process: we can simply recycle the one that has already been developed for the interchange process. Specifically, let $c(G)$ be the smallest number such that the functional inequality

$$\mathcal{E}_{(1,\dots,1)}^G(f, f) \leq c(G) \mathcal{E}_{(1,\dots,1)}(f, f), \quad (11)$$

holds for all $f: \Omega_{(1,\dots,1)} \rightarrow \mathbb{R}$. This fundamental quantity is known as the *comparison constant* of the interchange process on G . It was shown in [1] that

$$c(G) \lesssim \tau_{\text{MIX}}(G),$$

where \lesssim means inequality up to a universal multiplicative constant, and where $\tau_{\text{MIX}}(G)$ denotes the mixing time of the simple random walk on G . It is in fact believed that

$$c(G) \asymp \tau_{\text{REL}}(G), \quad (12)$$

see Conjecture 2 in [22]. This refinement, inspired by an analogous relation for the *Zero-Range process* [25], is already known to hold for several natural families of graphs ranging from low-dimensional tori [1] to high-dimensional products [22]. Those estimates can be combined with our main result to yield a general log-Sobolev inequality for the colored exclusion process (see Section 3.4 for details):

Corollary 2 (Log-Sobolev inequality for the colored exclusion process). *We have*

$$\max \left\{ 2\tau_{\text{REL}}(G), \log \left(\frac{n}{\kappa_{\text{MIN}}} \right) \right\} \leq \tau_{\text{LS}}(\kappa, G) \leq \frac{4}{\log 2} c(G) \log \left(\frac{n}{\kappa_{\text{MIN}}} \right).$$

To appreciate the sharpness of this general inequality, note that the lower and upper bounds are of the same order in the following two generic situations:

- For families of graphs with $c(G) \asymp 1$ (i.e. “well-connected” graphs), we obtain

$$\tau_{\text{LS}}(\kappa, G) \asymp \log\left(\frac{n}{\kappa_{\text{MIN}}}\right),$$

exactly as in the mean-field case. Note that this potentially constitutes a considerable extension of our main result, since the class of graphs satisfying $c(G) \asymp 1$ is believed to contain all expanders, as per (12).

- For graphs satisfying the conjecture (12), in the regime $\kappa_{\text{MIN}} \geq \varepsilon n$ ($\varepsilon > 0$ fixed), we get

$$\tau_{\text{LS}}(\kappa, G) \asymp \tau_{\text{REL}}(G).$$

This constitutes a multi-colored generalization of several estimates obtained for two colors, including [30, Theorem 4] on the cycle and [22, Corollary 5] on the hypercube.

Remark 4 (Mixing times). *One of the many interests of those log-Sobolev estimates is that they provide powerful controls on the strong L^∞ -mixing time of the process, see e.g., [34]. Let us here just give one concrete example: on the d -dimensional hypercube, our work implies that the balanced colored exclusion process with an arbitrarily fixed number $L \geq 2$ of colors mixes in time $\Theta(d^2)$. The special case $L = 2$ of this statement had been conjectured several years ago by Wilson [40], and was settled only recently [20].*

We end this section with an intriguing possibility, which arises naturally in view of Theorem 6 and of what happens in the mean-field case (Lemma 1).

Question 2 (Sensitivity of the modified log-Sobolev constant). *Can the choice of the parameter κ affect $\tau_{\text{MLS}}(\kappa, G)$ by more than a universal multiplicative constant?*

A negative answer would, in particular, substantially improve our current knowledge on the mixing times of the interchange and exclusion processes on general graphs. We note that, unlike our main result, the estimate on $\tau_{\text{MLS}}(\kappa)$ provided by Lemma 1 can *not* be directly transferred to more general graphs, since the modified log-Sobolev constant is notoriously *not* amenable to comparison techniques. This severe drawback constitutes a strong point in favor of log-Sobolev inequalities (as opposed to their modified versions) for mean-field interacting particle models, and was one of the motivations for the present work.

3 Proofs

3.1 General strategy

Let us start with an elementary but crucial observation about the multislice.

Remark 5 (Recursive structure). *If $(\omega_1, \dots, \omega_n)$ is uniformly distributed on Ω_κ , then the conditional law of $(\omega_1, \dots, \omega_{i-1}, \omega_{i+1}, \dots, \omega_n)$ given $\{\omega_i = \ell\}$ is uniform on $\Omega_{\kappa'}$, where*

$$\kappa' = (\kappa_1, \dots, \kappa_{\ell-1}, \kappa_\ell - 1, \kappa_{\ell+1}, \dots, \kappa_L).$$

Such a simple recursive structure suggests the possibility of proving Theorem 5 by induction over the dimension n , using the “chain rule” for entropy (see formula (13) below). This is in fact a classical strategy for establishing functional inequalities, known as the “martingale method”. Introduced by Lu & Yau [32] in the context of Kawasaki and Glauber dynamics, it has been successfully applied to various interacting particle systems [41, 30, 17, 18, 5, 21], as well as other Markov chains enjoying an appropriate recursive structure [7, 26, 27, 16, 23]. In particular, this is how Theorem 3 was proved. However, as explained in detail in [16], moving from the special case $L = 2$ covered by Theorem 3 to the general case studied in Theorem 4 significantly complicates the inductive argument, resulting in the loose $L \log L$ dependency mentioned at (7). Here we introduce two simple ideas to bypass those complications and prove Conjecture 1:

- (i) instead of just a single site, we condition on a whole region being colored with $\ell \in [L]$;
- (ii) when averaging the contributions from the various colors, we assign more weight to rare colors, which are the one which really govern $\tau_{\text{LS}}(\kappa)$. More precisely, our decomposition (16) below gives weight $1 - \frac{\kappa_\ell}{n}$ to the ℓ -colored region, whereas the traditional uniform average over all sites would give it the weight $\frac{\kappa_\ell}{n}$.

Let us now implement those ideas. We fix an observable $f: \Omega_\kappa \rightarrow \mathbb{R}_+$ once and for all. To lighten notation, we drop the index κ from our expectations, and write simply

$$\text{Ent}(f) := \mathbb{E}[f \log f] - \mathbb{E}[f] \log \mathbb{E}[f],$$

for the entropy of f . If Z is a random variable on Ω_κ , we define the *conditional entropy* of f given Z by simply replacing all expectations with conditional expectations, i.e.

$$\text{Ent}(f|Z) := \mathbb{E}[f \log f|Z] - \mathbb{E}[f|Z] \log \mathbb{E}[f|Z].$$

We then have the following elementary “chain rule”:

$$\text{Ent}(f) = \mathbb{E} [\text{Ent}(f|Z)] + \text{Ent}(\mathbb{E}[f|Z]). \quad (13)$$

The choice $Z = \omega_i$ is of course natural in light of Remark 5, and this was the one adopted in the proofs of Theorems 3 and 4. However, as mentioned in (i) above, we choose here to condition instead on the whole ℓ -colored region, i.e., on the random set

$$\xi_\ell := \{i \in [n] : \omega_i = \ell\}. \quad (14)$$

With $Z = \xi_\ell$, the formula (13) becomes

$$\text{Ent}(f) = \mathbb{E} [\text{Ent}(f|\xi_\ell)] + \text{Ent}(\mathbb{E}[f|\xi_\ell]). \quad (15)$$

Following our second idea (ii), we multiply both sides of this identity by the “unusual” weight $1 - \frac{\kappa_\ell}{n}$ and then sum over all colors $\ell \in [L]$. Recalling (1), we obtain the following formula, which will constitute the basis of our induction:

$$(L-1)\text{Ent}(f) = \underbrace{\sum_{\ell=1}^L \left(1 - \frac{\kappa_\ell}{n}\right) \mathbb{E} [\text{Ent}(f|\xi_\ell)]}_{\Sigma_1} + \underbrace{\sum_{\ell=1}^L \left(1 - \frac{\kappa_\ell}{n}\right) \text{Ent}(\mathbb{E}[f|\xi_\ell])}_{\Sigma_2}. \quad (16)$$

Our main task will consist in estimating the two terms Σ_1 and Σ_2 on the right-hand side, in terms of the log-Sobolev constants of certain lower-dimensional multislices. More precisely, we let $\kappa^{\setminus \ell}$ denote the parameter obtained from κ by removing the ℓ -th entry, i.e.

$$\kappa^{\setminus \ell} := (\kappa_1, \dots, \kappa_{\ell-1}, \kappa_{\ell+1}, \dots, \kappa_L),$$

and we will prove in the next section that

$$\Sigma_1 \leq (L-2) \max_{\ell \in [L]} \{\tau_{\text{LS}}(\kappa^{\setminus \ell})\} \mathcal{E}_\kappa(\sqrt{f}, \sqrt{f}); \quad (17)$$

$$\Sigma_2 \leq \max_{\ell \in [L]} \left\{ 2 \left(1 - \frac{\kappa_\ell}{n}\right) \tau_{\text{LS}}(\kappa_\ell, n - \kappa_\ell) \right\} \mathcal{E}_\kappa(\sqrt{f}, \sqrt{f}). \quad (18)$$

Plugging those estimates into (16) yields a log-Sobolev inequality for Ω_κ , thereby establishing the following recursive estimate.

Proposition 1 (Recursive log-Sobolev estimate). *We have*

$$(L-1)\tau_{\text{LS}}(\kappa) \leq (L-2) \max_{\ell \in [L]} \{\tau_{\text{LS}}(\kappa^{\setminus \ell})\} + \max_{\ell \in [L]} \left\{ 2 \left(1 - \frac{\kappa_\ell}{n}\right) \tau_{\text{LS}}(\kappa_\ell, n - \kappa_\ell) \right\}.$$

From this, the upper bound in Theorem 5 follows by an easy induction over the number L of colors, using the known log-Sobolev estimate for $L = 2$ (Theorem 2). The details, as well as the proof of the lower bound, are provided in Section 3.3.

3.2 Main recursion

This section is devoted to proving the two technical estimates (17) and (18) which, in view of the decomposition (16), establish Proposition 1.

Proof of the first estimate (17). Conditionally on the ℓ -colored region ξ_ℓ , f may be regarded as a function of the remaining coordinates $(\omega_i : i \in [n] \setminus \xi_\ell)$, which form a uniformly distributed element of $\Omega_{\kappa \setminus \ell}$. Consequently, the log-Sobolev inequality for the multislice $\Omega_{\kappa \setminus \ell}$ gives

$$\text{Ent}(f|\xi_\ell) \leq \frac{\tau_{\text{LS}}(\kappa \setminus \ell)}{2(n - \kappa_\ell)} \sum_{1 \leq i < j \leq n} \mathbb{E} \left[\left(\nabla^{ij} \sqrt{f} \right)^2 \mathbf{1}_{(i \notin \xi_\ell, j \notin \xi_\ell)} \middle| \xi_\ell \right].$$

Note that the event in the indicator can be rewritten as $\{\ell \notin \{\omega_i, \omega_j\}\}$, and that we may impose the restriction $\{\omega_i \neq \omega_j\}$ at no cost, since $\nabla^{ij} \sqrt{f} = 0$ on the event $\{\omega_i = \omega_j\}$. Taking expectations and rearranging, we arrive at

$$\left(1 - \frac{\kappa_\ell}{n}\right) \mathbb{E}[\text{Ent}(f|\xi_\ell)] \leq \frac{\tau_{\text{LS}}(\kappa \setminus \ell)}{2n} \sum_{1 \leq i < j \leq n} \mathbb{E} \left[\left(\nabla^{ij} \sqrt{f} \right)^2 \mathbf{1}_{(\omega_i \neq \omega_j)} \mathbf{1}_{(\ell \notin \{\omega_i, \omega_j\})} \right].$$

Summing over all $\ell \in [L]$ yields

$$\sum_{\ell=1}^L \left(1 - \frac{\kappa_\ell}{n}\right) \mathbb{E}[\text{Ent}(f|\xi_\ell)] \leq (L - 2) \max_{\ell \in [L]} \{\tau_{\text{LS}}(\kappa \setminus \ell)\} \mathcal{E}_\kappa(\sqrt{f}, \sqrt{f}),$$

which is exactly the claim made at (17). \square

Proof of the second estimate (18). Fix $\ell \in [L]$, and let us write

$$\mathbb{E}[f|\xi_\ell] = F(\xi_\ell), \tag{19}$$

for some non-negative function $F = F_\ell$. The distribution of ξ_ℓ is uniform over all κ_ℓ -element subsets of $[n]$, and this is precisely the stationary distribution of the occupied set in the κ_ℓ -particle Bernoulli-Laplace diffusion model on n sites. When applied to the function F , the log-Sobolev inequality for this process reads as follows:

$$\text{Ent}(\mathbb{E}[f|\xi_\ell]) \leq \frac{\tau_{\text{LS}}(\kappa_\ell, n - \kappa_\ell)}{2n} \sum_{1 \leq i < j \leq n} \mathbb{E} \left[\left(\sqrt{F(\xi_\ell^{ij})} - \sqrt{F(\xi_\ell)} \right)^2 \right], \tag{20}$$

where A^{ij} denotes the set obtained from A by swapping the membership status of i and j :

$$A^{ij} := \begin{cases} A \cup \{j\} \setminus \{i\} & \text{if } i \in A, j \notin A \\ A \cup \{i\} \setminus \{j\} & \text{if } i \notin A, j \in A \\ A & \text{else.} \end{cases}$$

Now, fix $1 \leq i < j \leq n$ and a κ_ℓ -element set $A \subseteq [n]$. First, by definition of F , we have

$$\mathbb{E}[f|\xi_\ell = A] = F(A).$$

On the other hand, since the involution $\tau^{ij}: \omega \mapsto \omega^{ij}$ preserves the uniform law on Ω_κ and maps the event $\{\xi_\ell = A\}$ onto the event $\{\xi_\ell = A^{ij}\}$, we have

$$\mathbb{E}[f \circ \tau^{ij}|\xi_\ell = A] = \mathbb{E}[f|\xi_\ell = A^{ij}] = F(A^{ij}).$$

But the function $\Phi: (u, v) \mapsto (\sqrt{u} - \sqrt{v})^2$ is convex on \mathbb{R}_+^2 , so Jensen's inequality yields

$$\begin{aligned} \left(\sqrt{F(A^{ij})} - \sqrt{F(A)}\right)^2 &= \Phi(\mathbb{E}[f \circ \tau^{ij}|\xi_\ell = A], \mathbb{E}[f|\xi_\ell = A]) \\ &\leq \mathbb{E}[\Phi(f \circ \tau^{ij}, f)|\xi_\ell = A] = \mathbb{E}\left[\left(\nabla^{ij} \sqrt{f}\right)^2 \Big| \xi_\ell = A\right]. \end{aligned}$$

Moreover, we have $\left(\sqrt{F(A^{ij})} - \sqrt{F(A)}\right)^2 = 0$ when A contains neither i nor j , so we obtain

$$\left(\sqrt{F(A^{ij})} - \sqrt{F(A)}\right)^2 \leq \mathbb{E}\left[\left(\nabla^{ij} \sqrt{f}\right)^2 \Big| \xi_\ell = A\right] (\mathbf{1}_{(i \in A)} + \mathbf{1}_{(j \in A)}).$$

Averaging this inequality over all possible κ_ℓ -element set $A \subseteq [n]$ yields

$$\mathbb{E}\left[\left(\sqrt{F(\xi_\ell^{ij})} - \sqrt{F(\xi_\ell)}\right)^2\right] \leq \mathbb{E}\left[\left(\nabla^{ij} \sqrt{f}\right)^2 (\mathbf{1}_{(i \in \xi_\ell)} + \mathbf{1}_{(j \in \xi_\ell)})\right].$$

We may now plug this estimate back into (20) to arrive at

$$\text{Ent}(\mathbb{E}[f|\xi_\ell]) \leq \frac{\tau_{\text{LS}}(\kappa_\ell, n - \kappa_\ell)}{2n} \sum_{1 \leq i < j \leq n} \mathbb{E}\left[\left(\nabla^{ij} \sqrt{f}\right)^2 (\mathbf{1}_{(\omega_i = \ell)} + \mathbf{1}_{(\omega_j = \ell)})\right].$$

Finally, multiplying by $(1 - \frac{\kappa_\ell}{n})$ and summing over all $\ell \in [L]$ gives

$$\sum_{\ell=1}^L \left(1 - \frac{\kappa_\ell}{n}\right) \text{Ent}(\mathbb{E}[f|\xi_\ell]) \leq \max_{\ell \in [L]} \left\{2 \left(1 - \frac{\kappa_\ell}{n}\right) \tau_{\text{LS}}(\kappa_\ell, n - \kappa_\ell)\right\} \mathcal{E}_\kappa(\sqrt{f}, \sqrt{f}),$$

which is precisely the claim (18). □

3.3 Putting things together

To complete the proof of Theorem 5, we only need an estimate on the second term appearing on the right-hand side of our recursive log-Sobolev inequality. We of course use Theorem 3.

Lemma 2 (Two-color estimate). *For any $\ell \in [L]$, we have*

$$\left(1 - \frac{\kappa_\ell}{n}\right) \tau_{\text{LS}}(\kappa_\ell, n - \kappa_\ell) \leq \frac{2}{\log 2} \log \left(\frac{n}{\kappa_{\text{MIN}}}\right).$$

Proof. By Theorem 3, we have

$$\left(1 - \frac{\kappa_\ell}{n}\right) \tau_{\text{LS}}(\kappa_\ell, n - \kappa_\ell) \leq \frac{2}{\log 2} \left(1 - \frac{\kappa_\ell}{n}\right) \log \left(\frac{n^2}{\kappa_\ell(n - \kappa_\ell)}\right).$$

Since the right-hand side is maximized when $\kappa_\ell = \kappa_{\text{MIN}}$, our task boils down to establishing

$$\left(1 - \frac{\kappa_{\text{MIN}}}{n}\right) \log \left(\frac{n^2}{\kappa_{\text{MIN}}(n - \kappa_{\text{MIN}})}\right) \leq \log \left(\frac{n}{\kappa_{\text{MIN}}}\right).$$

But this is exactly the special case $t = \frac{\kappa_{\text{MIN}}}{n}$ of the inequality

$$t \log t - (1 - t) \log(1 - t) \leq 0,$$

which is valid for all $t \in [0, \frac{1}{2}]$. To see this, note that the left-hand side is a convex function of $t \in [0, \frac{1}{2}]$ (as can be easily checked by differentiating) and that it equals zero at the two boundary points $t = 0$ and $t = \frac{1}{2}$. \square

We are now in position to prove our main result.

Proof of the upper bound in Theorem 5. Our aim is to prove that

$$\tau_{\text{LS}}(\kappa) \leq \Phi(\kappa) := \frac{4}{\log 2} \log \left(\frac{n}{\kappa_{\text{MIN}}}\right). \quad (21)$$

We proceed by induction over the dimension L of the parameter $\kappa = (\kappa_1, \dots, \kappa_L)$. By combining Proposition 1 and Lemma 2, we have

$$(L - 1) \tau_{\text{LS}}(\kappa) \leq \Phi(\kappa) + (L - 2) \max_{\ell \in [L]} \tau_{\text{LS}}(\kappa^{\setminus \ell}), \quad (22)$$

which already establishes the claim in the base case $L = 2$. Now, assume that $L \geq 3$ and that the claim already holds for lower values of L . In particular, we know that

$$\tau_{\text{LS}}(\kappa^{\setminus \ell}) \leq \Phi(\kappa^{\setminus \ell}),$$

for all $\ell \in [L]$. But $\Phi(\kappa^{\setminus \ell}) \leq \Phi(\kappa)$, since removing an entry from the parameter κ can only decrease the value of the sum $n = \kappa_1 + \dots + \kappa_L$ and increase the value of the minimum $\kappa_{\text{MIN}} = \min\{\kappa_1, \dots, \kappa_L\}$. Consequently, (22) gives

$$(L - 1) \tau_{\text{LS}}(\kappa) \leq \Phi(\kappa) + (L - 2) \Phi(\kappa) = (L - 1) \Phi(\kappa),$$

and (21) is established. \square

Our upper bound on $\tau_{\text{LS}}(\kappa)$ implies the lower bound on $\iota(\kappa)$ given in Corollary 1, thanks to the well-known relation between log-Sobolev inequalities and small-set expansion:

Lemma 3 (Log-Sobolev inequality and small-set expansion). *We have*

$$\iota(\kappa)\tau_{\text{LS}}(\kappa) \geq n.$$

Proof. This follows from the definitions of $\iota(\kappa)$ and $\tau_{\text{LS}}(\kappa)$, once we have observed that

$$\begin{aligned} \mathcal{E}_\kappa(\mathbf{1}_A, \mathbf{1}_A) &= \frac{|\partial A|}{n|\Omega_\kappa|}, \\ \text{Ent}_\kappa(\mathbf{1}_A) &= \frac{|A|}{|\Omega_\kappa|} \log\left(\frac{|\Omega_\kappa|}{|A|}\right), \end{aligned}$$

for any event $A \subseteq \Omega_\kappa$. □

The inequality in Lemma 3 is obtained by restricting the definition of the log-Sobolev inequality to indicator functions, and could therefore be rather loose. However, it turns out to be sharp in the present case, as we will now see.

Proof of the remaining halves of Theorem 5 and Corollary 1. By definition, we have

$$\iota(\kappa) \leq \frac{|\partial A|}{|A| \log\left(\frac{|\Omega_\kappa|}{|A|}\right)}, \quad (23)$$

for any non-empty event $A \subseteq \Omega_\kappa$. We fix $\ell \in [L]$ such that $\kappa_\ell = \kappa_{\text{MIN}}$, and consider the choice

$$A := \{\xi_\ell = \{1, \dots, \kappa_{\text{MIN}}\}\}, \quad (24)$$

where we recall that ξ_ℓ is the ℓ -colored region. Since ξ_ℓ is uniformly distributed over all κ_{MIN} -element subsets of $[n]$, we have

$$|A| = \frac{|\Omega_\kappa|}{\binom{n}{\kappa_{\text{MIN}}}}.$$

On the other hand, from any state inside A , there are precisely $\kappa_{\text{MIN}}(n - \kappa_{\text{MIN}})$ transpositions that result in a state outside A , and hence

$$|\partial A| = \kappa_{\text{MIN}}(n - \kappa_{\text{MIN}})|A|.$$

Thus, the inequality (23) gives

$$\begin{aligned} \iota(\kappa) &\leq \frac{\kappa_{\text{MIN}}(n - \kappa_{\text{MIN}})}{\log \binom{n}{\kappa_{\text{MIN}}}} \\ &\leq \frac{n}{\log \left(\frac{n}{\kappa_{\text{MIN}}} \right)}, \end{aligned} \tag{25}$$

where the second line uses the classical binomial estimate $\binom{n}{k} \geq \left(\frac{n}{k}\right)^k$, valid for all $1 \leq k \leq n$. This establishes the upper bound in Corollary 1, as well as the lower bound in Theorem 5, by Lemma 3. Finally, note that in the case $\kappa = (\lfloor n/2 \rfloor, \lfloor n/2 \rfloor)$, the estimate (25) yields

$$\frac{n}{\iota(\lfloor \frac{n}{2} \rfloor, \lfloor \frac{n}{2} \rfloor)} \geq 4 \log \left(\frac{n}{\kappa_{\text{MIN}}} \right) - o(1). \tag{26}$$

Thus, our pre-factor can not be improved by more than $\log 2$, as claimed in Remark 3. \square

3.4 Coarsening argument

It now remains to prove Lemma 1 and Corollary 2. Both will rely on the elementary observation, already alluded to in Remark 1, that the multislice Ω_κ is a “coarsened” version of the “free” multislice $\Omega_{(1, \dots, 1)}$, where $(1, \dots, 1)$ denotes the all-one vector of length n . To formalize this, let us introduce the projection $\Psi: [n] \rightarrow [L]$ defined by the relation

$$\Psi(i) = \ell \iff i \in [\kappa_1 + \dots + \kappa_{\ell-1} + 1, \kappa_1 + \dots + \kappa_\ell],$$

and extend this definition to sequences by coordinate-wise application:

$$\Psi(\omega_1, \dots, \omega_n) := (\Psi(\omega_1), \dots, \Psi(\omega_n)).$$

The mapping Ψ “projects” the multislice $\Omega_{(1, \dots, 1)}$ onto Ω_κ in the following precise sense.

Lemma 4 (Coarsening). *For any observable $f: \Omega_\kappa \rightarrow \mathbb{R}$, we have*

$$\mathbb{E}_\kappa[f] = \mathbb{E}_{(1, \dots, 1)}[f \circ \Psi].$$

Moreover, for any $f, g: \Omega_\kappa \rightarrow \mathbb{R}$ and any weighted graph $G = (G_{ij})_{1 \leq i, j \leq n}$,

$$\mathcal{E}_\kappa^G(f, g) = \mathcal{E}_{(1, \dots, 1)}^G(f \circ \Psi, g \circ \Psi).$$

Proof. By construction, we have $|\Psi^{-1}(\{\ell\})| = \kappa_\ell$ for each color $\ell \in [L]$, and hence Ψ maps $\Omega_{(1,\dots,1)}$ to Ω_κ . The first claim asserts that the Ψ -image of the uniform measure on $\Omega_{(1,\dots,1)}$ is the uniform measure on Ω_κ , which is nothing more than the observation that each element of Ω_κ admits the same number of pre-images under Ψ (namely $\kappa_1! \cdots \kappa_L!$). The second claim follows from the first and the definition (10), once we note that the commutativity relation

$$\nabla^{ij}(f \circ \Psi) = (\nabla^{ij} f) \circ \Psi,$$

trivially holds for all $1 \leq i < j \leq n$ and all $f: \Omega_\kappa \rightarrow \mathbb{R}$. \square

We can now easily establish our log-Sobolev estimate for the colored exclusion process.

Proof of Corollary 2. We use Lemma 4 and the definitions of $\tau_{\text{LS}}(\kappa)$ and $c(G)$ to write

$$\begin{aligned} \text{Ent}_\kappa(f) &\leq \tau_{\text{LS}}(\kappa) \mathcal{E}_\kappa(\sqrt{f}, \sqrt{f}) \\ &= \tau_{\text{LS}}(\kappa) \mathcal{E}_{(1,\dots,1)}(\sqrt{f} \circ \Psi, \sqrt{f} \circ \Psi) \\ &\leq \tau_{\text{LS}}(\kappa) c(G) \mathcal{E}_{(1,\dots,1)}^G(\sqrt{f} \circ \Psi, \sqrt{f} \circ \Psi) \\ &= \tau_{\text{LS}}(\kappa) c(G) \mathcal{E}_\kappa^G(\sqrt{f}, \sqrt{f}). \end{aligned}$$

Since $f: \Omega_\kappa \rightarrow \mathbb{R}_+$ was arbitrary, we have just proved

$$\tau_{\text{LS}}(\kappa, G) \leq c(G) \tau_{\text{LS}}(\kappa).$$

The claimed upper bound now follows from our main estimate on $\tau_{\text{LS}}(\kappa)$. The lower bound

$$\tau_{\text{LS}}(\kappa, G) \geq 2\tau_{\text{REL}}(G)$$

is obtained by combining the general inequality $\tau_{\text{LS}}(\cdot) \geq 2\tau_{\text{REL}}(\cdot)$ with Theorem 6. To prove the other lower bound, we choose the test function $f = \mathbf{1}_A$ in the definition of the log-Sobolev inequality, with A as in (24). We have already seen that $|A| = |\Omega_\kappa| / \binom{n}{\kappa_{\text{MIN}}}$. Moreover, we now have $|\partial A| \leq |A|d\kappa_{\text{MIN}}$ where d denotes the degree in G , since moving from A to A^c requires transposing some site in $\{1, \dots, \kappa_{\text{MIN}}\}$ with one of its d neighbors. We thus obtain

$$\tau_{\text{LS}}(\kappa, G) \geq \frac{|A|d \log \frac{|\Omega_\kappa|}{|A|}}{|\partial A|} \geq \frac{1}{\kappa_{\text{MIN}}} \log \binom{n}{\kappa_{\text{MIN}}} \geq \log \left(\frac{n}{\kappa_{\text{MIN}}} \right),$$

and the proof is complete. \square

Proof of Lemma 1. The statement $\tau_{\text{REL}}(\kappa) = 1$ is a (simple) special case of Theorem 6. This immediately implies $\tau_{\text{MLS}}(\kappa) \geq \frac{1}{2}$, by the general relation (6). To prove the more interesting statement $\tau_{\text{MLS}}(\kappa) \leq 2$, we take an arbitrary function $f: \Omega_\kappa \rightarrow \mathbb{R}_+$ and use Lemma 4 to write

$$\begin{aligned} \text{Ent}_\kappa(f) &= \text{Ent}_{(1, \dots, 1)}(f \circ \Psi) \\ &\leq \tau_{\text{MLS}}(1, \dots, 1) \mathcal{E}_{(1, \dots, 1)}(f \circ \Psi, \log f \circ \Psi) \\ &= \tau_{\text{MLS}}(1, \dots, 1) \mathcal{E}_\kappa(f, \log f). \end{aligned}$$

This shows that $\tau_{\text{MLS}}(\kappa) \leq \tau_{\text{MLS}}(1, \dots, 1)$, and the desired conclusion now follows from the classical estimate $\tau_{\text{MLS}}(1, \dots, 1) \leq \frac{1}{2}$, due to Goel [18, Corollary 3.1]. \square

Acknowledgment

The author warmly thanks Jonathan Hermon for his valuable comments on a preliminary version of this work.

References

- [1] Gil Alon and Gady Kozma. Comparing with octopi. *arXiv e-prints*, page arXiv:1811.10537, November 2018.
- [2] Nathanaël Berestycki and Rick Durrett. A phase transition in the random transposition random walk. *Probab. Theory Related Fields*, 136(2):203–233, 2006.
- [3] Sergey Bobkov and Prasad Tetali. Modified log-Sobolev inequalities, mixing and hypercontractivity. In *Proceedings of the Thirty-Fifth Annual ACM Symposium on Theory of Computing*, pages 287–296. ACM, New York, 2003.
- [4] Pietro Caputo, Thomas M. Liggett, and Thomas Richthammer. Proof of Aldous’ spectral gap conjecture. *J. Amer. Math. Soc.*, 23(3):831–851, 2010.
- [5] Pietro Caputo and Gustavo Posta. Entropy dissipation estimates in a zero-range dynamics. *Probab. Theory Related Fields*, 139(1-2):65–87, 2007. MR2322692.
- [6] Stephen B. Connor and Richard J. Pymar. Mixing times for exclusion processes on hypergraphs. *Electron. J. Probab.*, 24:Paper No. 73, 48, 2019.

- [7] Paolo Dai Pra, Anna Maria Paganoni, and Gustavo Posta. Entropy inequalities for unbounded spin systems. *Ann. Probab.*, 30(4):1959–1976, 2002.
- [8] P. Diaconis and L. Saloff-Coste. Logarithmic Sobolev inequalities for finite Markov chains. *Ann. Appl. Probab.*, 6(3):695–750, 1996.
- [9] Persi Diaconis. *Group representations in probability and statistics*, volume 11 of *Institute of Mathematical Statistics Lecture Notes—Monograph Series*. Institute of Mathematical Statistics, Hayward, CA, 1988.
- [10] Persi Diaconis and Laurent Saloff-Coste. Comparison techniques for random walk on finite groups. *Ann. Probab.*, 21(4):2131–2156, 1993.
- [11] Persi Diaconis and Laurent Saloff-Coste. Comparison theorems for reversible Markov chains. *Ann. Appl. Probab.*, 3(3):696–730, 1993.
- [12] Persi Diaconis and Mehrdad Shahshahani. Generating a random permutation with random transpositions. *Z. Wahrsch. Verw. Gebiete*, 57(2):159–179, 1981.
- [13] Persi Diaconis and Mehrdad Shahshahani. Time to reach stationarity in the Bernoulli-Laplace diffusion model. *SIAM J. Math. Anal.*, 18(1):208–218, 1987.
- [14] Yuval Filmus. FKN theorem for the multislice, with applications. *Combin. Probab. Comput.*, 29(2):200–212, 2020.
- [15] Yuval Filmus and Ferdinand Ihringer. Boolean constant degree functions on the slice are juntas. *Discrete Math.*, 342(12):111614, 7, 2019.
- [16] Yuval Filmus, Ryan O’Donnell, and Xinyu Wu. A log-Sobolev inequality for the multi-slice, with applications. *arXiv e-prints*, page arXiv:1809.03546, September 2018.
- [17] Fuqing Gao and Jeremy Quastel. Exponential decay of entropy in the random transposition and Bernoulli-Laplace models. *Ann. Appl. Probab.*, 13(4):1591–1600, 2003.
- [18] Sharad Goel. Modified logarithmic Sobolev inequalities for some models of random walk. *Stochastic Process. Appl.*, 114(1):51–79, 2004.

- [19] Jonathan Hermon and Yuval Peres. A characterization of L_2 mixing and hypercontractivity via hitting times and maximal inequalities. *Probab. Theory Related Fields*, 170(3-4):769–800, 2018.
- [20] Jonathan Hermon and Richard Pymar. The exclusion process mixes (almost) faster than independent particles, 2018.
- [21] Jonathan Hermon and Justin Salez. Entropy dissipation estimates for inhomogeneous zero-range processes, 2019.
- [22] Jonathan Hermon and Justin Salez. The interchange process on high-dimensional products, 2019.
- [23] Jonathan Hermon and Justin Salez. Modified log-Sobolev inequalities for strong-Rayleigh measures. *arXiv e-prints*, page arXiv:1902.02775, Feb 2019.
- [24] Jonathan Hermon and Justin Salez. The interchange process on high-dimensional products. *arXiv e-prints*, page arXiv:1905.02146, May 2019.
- [25] Jonathan Hermon and Justin Salez. A version of Aldous’ spectral-gap conjecture for the zero range process. *Ann. Appl. Probab.*, 29(4):2217–2229, 2019.
- [26] Mark Jerrum and Jung-Bae Son. Spectral gap and log-sobolev constant for balanced matroids. In *The 43rd Annual IEEE Symposium on Foundations of Computer Science, 2002. Proceedings.*, pages 721–729, Nov 2002.
- [27] Mark Jerrum, Jung-Bae Son, Prasad Tetali, and Eric Vigoda. Elementary bounds on Poincaré and log-Sobolev constants for decomposable Markov chains. *Ann. Appl. Probab.*, 14(4):1741–1765, 2004.
- [28] Johan Jonasson. Mixing times for the interchange process. *ALEA Lat. Am. J. Probab. Math. Stat.*, 9(2):667–683, 2012.
- [29] J. Kahn, G. Kalai, and N. Linial. The influence of variables on boolean functions. In *[Proceedings 1988] 29th Annual Symposium on Foundations of Computer Science*, pages 68–80, 1988.
- [30] Tzong-Yow Lee and Horng-Tzer Yau. Logarithmic Sobolev inequality for some models of random walks. *Ann. Probab.*, 26(4):1855–1873, 1998.

- [31] László Lovász and Ravi Kannan. Faster mixing via average conductance. In *Annual ACM Symposium on Theory of Computing (Atlanta, GA, 1999)*, pages 282–287. ACM, New York, 1999.
- [32] Sheng Lin Lu and Horng-Tzer Yau. Spectral gap and logarithmic Sobolev inequality for Kawasaki and Glauber dynamics. *Comm. Math. Phys.*, 156(2):399–433, 1993.
- [33] Peter Matthews. A strong uniform time for random transpositions. *J. Theoret. Probab.*, 1(4):411–423, 1988.
- [34] Ravi Montenegro and Prasad Tetali. Mathematical aspects of mixing times in Markov chains. *Found. Trends Theor. Comput. Sci.*, 1(3):x+121, 2006.
- [35] Ben Morris. The mixing time for simple exclusion. *Ann. Appl. Probab.*, 16(2):615–635, 2006.
- [36] Ryan O’Donnell. *Analysis of Boolean functions*. Cambridge University Press, New York, 2014.
- [37] Roberto Imbuzeiro Oliveira. Mixing of the symmetric exclusion processes in terms of the corresponding single-particle random walk. *Ann. Probab.*, 41(2):871–913, 2013.
- [38] Oded Schramm. Compositions of random transpositions. *Israel J. Math.*, 147:221–243, 2005.
- [39] Lucas Teyssier. Limit profile for random transpositions. *arXiv e-prints*, page arXiv:1905.08514, May 2019.
- [40] David Bruce Wilson. Mixing times of Lozenge tiling and card shuffling Markov chains. *Ann. Appl. Probab.*, 14(1):274–325, 2004.
- [41] Horng-Tzer Yau. Logarithmic Sobolev inequality for generalized simple exclusion processes. *Probab. Theory Related Fields*, 109(4):507–538, 1997.