



Justin Salez

Cutoff for non-negatively curved Markov chains

Received March 22, 2021

Abstract. Discovered by Aldous, Diaconis and Shahshahani in the context of card shuffling, the cutoff phenomenon has since then been established for a variety of Markov chains. However, proving cutoff remains a delicate affair, which requires a very detailed knowledge of the chain. Identifying the general mechanisms underlying this phase transition, without having to pinpoint its precise location, remains one of the most fundamental open problems in the area of mixing times. In the present paper, we make a step in this direction by establishing cutoff for all Markov chains with non-negative curvature, under a suitably refined product condition. The result applies, in particular, to the random walk on abelian Cayley expanders satisfying a mild degree assumption, hence to the random walk on *almost all* abelian Cayley graphs. Our proof relies on a quantitative *entropic concentration principle*, which we believe to lie behind all cutoff phenomena.

Keywords. Markov chains, cutoff phenomenon, varentropy, curvature

1. Introduction

1.1. The cutoff phenomenon

Setup. Throughout the paper, we consider a stochastic matrix P on a finite state space \mathcal{X} , and we let $(\mathcal{P}_t)_{t \geq 0}$ denote the associated heat kernel, defined for $t \in [0, \infty)$ and $x, y \in \mathcal{X}$ by

$$\mathcal{P}_t(x, y) := e^{-t} \sum_{k=0}^{\infty} \frac{P^k(x, y)t^k}{k!}.$$

Any continuous-time Markov semigroup on a finite state space takes this form, after a trivial time scaling ensuring that jumps occur at rate at most 1. As soon as P is irreducible, we have

$$\mathcal{P}_t(x, y) \xrightarrow[t \rightarrow \infty]{} \pi(y),$$

where $\pi = \pi P$ denotes the unique invariant law. The rate at which this convergence occurs is captured by the so-called *mixing time*, defined for any precision $\varepsilon \in (0, 1)$ by

$$t_{\text{MIX}}(\varepsilon) := \min \left\{ t \geq 0 : \max_{x \in \mathcal{X}} \|\mathcal{P}_t(x, \cdot) - \pi\|_{\text{TV}} \leq \varepsilon \right\},$$

Justin Salez: CEREMADE, Université Paris-Dauphine and PSL, 75775 Paris, France;
justin.salez@dauphine.psl.eu

Mathematics Subject Classification (2020): 60J27

where $\|\mu - \pi\|_{\text{TV}} = \max_{A \subseteq \mathcal{X}} |\mu(A) - \pi(A)|$ denotes total-variation distance. Understanding how this fundamental parameter depends on the underlying transition matrix P constitutes a fascinating area of research, at the crossroad between probability, discrete geometry, spectral analysis and functional inequalities; see the books [43, 51] for an introduction.

The cutoff phenomenon. Quantifying the approach to equilibrium is of course particularly relevant when the number of states is large. One is thus naturally led to consider a sequence $(P_n)_{n \geq 1}$ of irreducible transition matrices and investigate the asymptotic behavior of their mixing times as $n \rightarrow \infty$. To lighten our exposition, we will completely drop the subscript n from all notations. For many natural chains, a remarkable phase transition known as a *cutoff* has been established: the distance to equilibrium remains close to its maximal value for a long time, and then suddenly drops to zero on a much shorter time-scale. More formally, for any fixed $\varepsilon \in (0, 1)$,

$$\frac{t_{\text{MIX}}(1 - \varepsilon)}{t_{\text{MIX}}(\varepsilon)} = 1 + o(1), \tag{1}$$

where $o(1)$ denotes a quantity that vanishes as $n \rightarrow \infty$. The first instances of this phenomenon were discovered in the 80's by Aldous, Diaconis and Shahshahani in the context of card shuffling [1, 2, 21]. Since then, cutoff has been established on a variety of examples; see in particular the surveys [20, 59]. The present paper is concerned with the more fundamental problem of singling out abstract conditions under which this phase transition occurs, without having to pinpoint its precise location. This question was raised by Aldous and Diaconis in their seminal works [1, 2], and constitutes one of the most important open problems in the modern study of Markov chains.

The product condition. In the reversible case (i.e., $\pi(x)P(x, y) = \pi(y)P(y, x)$ for all $x, y \in \mathcal{X}$), cutoff is easily seen to imply the so-called *product condition*: for each fixed $\varepsilon \in (0, 1)$,

$$t_{\text{REL}} \ll t_{\text{MIX}}(\varepsilon), \tag{2}$$

where the notation $a \ll b$ means $a/b = o(1)$, and where t_{REL} denotes the *relaxation time* (inverse spectral gap) of the matrix P . The interest of this criterion is that it only involves orders of magnitude: unlike the definition (1), it can be checked without having to determine the precise prefactor in front of mixing times. In the 2004 AIM workshop on mixing times, Peres [57] conjectured that (2) is also sufficient for cutoff. This has been verified for birth-and-death chains [22] and, more generally, random walks on trees [7]. Unfortunately, counter-examples have been constructed in the general case; see [18, Section 6]. In fact, this failure is *generic*, in the sense that (2) is stable under a general rank-one perturbation which completely destroys cutoff; see [43, Example 18.7]. Thus, the product condition will incorrectly predict cutoff for many natural Markov chains, including certain random walks on abelian groups.

Other criteria. A few other criteria for cutoff of reversible chains have been proposed. In particular, Basu, Hermon and Peres provided a formulation in terms of concentration of

hitting times [7]. Unfortunately, verifying this condition requires determining the precise prefactor in front of hitting times, which, to the best of our knowledge, has not been practically carried out beyond random walks on trees. In the specific context of random walks on regular graphs with fixed degree $d \geq 3$, the delicate *asymptotic Ramanujan property*

$$t_{\text{REL}} = \frac{d}{d - 2\sqrt{d-1}} + o(1),$$

is known to imply cutoff [30, 44, 55]; see [16] for a generalization. Note, however, that unlike the product condition (2), verifying the Ramanujan property *does* require determining the relaxation time up to an $o(1)$ term, a notoriously challenging task even on random instances [13, 28]. Finally, let us mention an impressive series of works by Lubetzky and Sly [46–49] (see also [29]), which develops a general framework for proving cutoff in certain spin systems at sufficiently high temperature, without having to determine the cutoff location.

Our contribution. For a broad class of chains known as *non-negatively curved* chains, we provide a sufficient “product-like” condition for cutoff, which only involves comparing orders of magnitude of t_{REL} and t_{MIX} . Moreover, we do not require reversibility, but only symmetry of the support:

$$P(x, y) > 0 \iff P(y, x) > 0. \quad (3)$$

Before stating the result, let us briefly recall the two notions of curvature on which our approach relies: the Bakry–Émery curvature, and the Ollivier–Ricci curvature.

1.2. Non-negative curvature

Ricci curvature is a central concept in Riemannian geometry; see e.g. the book [39]. Over the past decades, considerable efforts have been made to develop a satisfactory analogue of this notion in discrete settings, such as graphs and Markov chains. In particular, Bakry and Émery [4, 5] proposed a fruitful approach based on Γ -calculus; see also the book [6]. More recently, Ollivier [53] put forward a different definition of curvature based on optimal transport, which applies to arbitrary metric spaces. Our results will apply to both definitions indifferently.

Ollivier–Ricci curvature. Under the symmetry assumption (3), we may turn the state space \mathcal{X} into a metric space by equipping it with the distance

$$\text{dist}(x, y) := \min \{k \in \mathbb{Z}_+ : P^k(x, y) > 0\}.$$

We use the graph-theoretical notation $x \sim y$ to mean that $\text{dist}(x, y) = 1$. As any distance, the above metric can be “lifted” to the distributional level via optimal transport. Specifically, the L^1 -Wasserstein distance between two probability measures μ and ν on \mathcal{X} is defined as

$$\mathcal{W}_1(\mu, \nu) := \min_{\gamma \in \Pi(\mu, \nu)} \sum_{x, y \in \mathcal{X}} \gamma(x, y) \text{dist}(x, y),$$

where $\Pi(\mu, \nu)$ is the set of all *couplings* of μ and ν , i.e. probability measures on $\mathcal{X} \times \mathcal{X}$ whose marginals are μ and ν . The *Ollivier–Ricci curvature* is the largest number κ such that the inequality

$$\mathcal{W}_1(\mathcal{P}_t(x, \cdot), \mathcal{P}_t(y, \cdot)) \leq e^{-\kappa t} \text{dist}(x, y)$$

holds for all $t \in \mathbb{R}_+$ and all $x, y \in \mathcal{X}$. By Kantorovich duality, this is equivalent to the inequality

$$\|\mathcal{P}_t f\|_{\text{LIP}} \leq e^{-\kappa t} \|f\|_{\text{LIP}} \tag{4}$$

for all $t \in \mathbb{R}_+$ and all observables $f: \mathcal{X} \rightarrow \mathbb{R}$, where $\|f\|_{\text{LIP}} = \max_{x \sim y} |f(x) - f(y)|$. In particular, non-negative Ollivier–Ricci curvature simply means that the semigroup is a contraction for the Lipschitz norm. This natural property constitutes the essence of the powerful *path coupling method* [17], and its consequences in terms of geometry, mixing, and concentration of measure have been massively investigated. The literature is too vast for an exhaustive account, and we refer the reader to the survey [54] for details. Establishing non-negative curvature is easier than one may think. Indeed, by convexity of \mathcal{W}_1 , it is enough to prove the one-step estimate

$$\mathcal{W}_1(P(x, \cdot), P(y, \cdot)) \leq (1 - \kappa) \text{dist}(x, y).$$

Furthermore, by the triangle inequality, we may restrict our attention to the case where $x \sim y$.

Bakry–Émery curvature. We now turn to Bakry–Émery curvature. We only introduce the necessary notation, and refer to the beautiful lecture notes [61] for details. Write $\mathcal{L} = P - \text{Id}$ for the generator of the semigroup $(\mathcal{P}_t)_{t \geq 0}$. This operator acts on any observable $f: \mathcal{X} \rightarrow \mathbb{R}$ via

$$(\mathcal{L}f)(x) := \sum_{y \in \mathcal{X}} P(x, y)(f(y) - f(x)),$$

and the associated *carré du champ* Γ is given, for all $f, g: \mathcal{X} \rightarrow \mathbb{R}$, by the formula

$$\Gamma(f, g)(x) := \frac{1}{2} \sum_{y \in \mathcal{X}} P(x, y)(f(y) - f(x))(g(y) - g(x)).$$

The *Bakry–Émery curvature* is then defined as the largest number κ such that the inequality

$$\frac{1}{2} \mathcal{L}\Gamma(f, f) \geq \Gamma(f, \mathcal{L}f) + \kappa \Gamma(f, f)$$

holds pointwise for all $f: \mathcal{X} \rightarrow \mathbb{R}$. This functional inequality, often denoted $\text{CD}(\kappa, \infty)$, is equivalent to the following subcommutativity relation between the carré du champ and the semigroup:

$$\forall t \in \mathbb{R}_+, \quad \Gamma(\mathcal{P}_t f, \mathcal{P}_t f) \leq e^{-2\kappa t} \mathcal{P}_t \Gamma(f, f). \tag{5}$$

For consequences in terms of geometry, mixing, and concentration of measure, see [25].

Non-negatively curved chains. In the discrete setting, there is apparently no known relation between the Bakry–Émery curvature and the Ollivier–Ricci curvature, although these two notions share many similarities. We emphasize that our results will apply to both definitions indifferently. Thus, by a *non-negatively curved chain*, we will henceforth simply mean a chain that has non-negative curvature in the Bakry–Émery sense *or* the Ollivier–Ricci sense. Non-negatively curved chains are ubiquitous, and appear in a broad variety of contexts. Classical examples include:

- all random walks on abelian groups [42];
- all conjugacy-invariant random walks on groups [11, 26];
- simple random walks on all Bruhat graphs of Coxeter groups [60];
- all monotone birth-and-death chains [40];
- the zero-range process with non-decreasing rates, and many other particle systems [62];
- the Glauber dynamics for various spin systems at high temperature [26].

To complement this list of examples, we note that non-negative curvature is closed under several natural operations, such as composition, superposition and L^1 -tensorization [53].

1.3. Results and implications

Recall that P is *not* required to be reversible, i.e., to coincide with its adjoint

$$P^*(x, y) := \frac{\pi(y)P(y, x)}{\pi(x)}.$$

Consequently, we define the relaxation time t_{REL} of the chain to be the inverse spectral gap of the additive reversibilization $\frac{P+P^*}{2}$. Equivalently, t_{REL} is the optimal constant in the Poincaré inequality

$$\mathbf{Var}(f) \leq t_{\text{REL}} \mathbf{E}[\Gamma(f, f)], \quad (6)$$

valid for all observables $f: \mathcal{X} \rightarrow \mathbb{R}$, where \mathbf{E} and \mathbf{Var} denote expectation and variance on the finite probability space (\mathcal{X}, π) . We also introduce the *sparsity parameter*

$$\Delta := \max \left\{ \frac{1}{P(x, y)} : x, y \in \mathcal{X}, x \sim y \right\},$$

which is simply the maximum degree when P is the transition matrix of simple random walk on a graph. To avoid degeneracies, we will assume that $|\mathcal{X}| \geq 3$, so that (under our assumption (3)),

$$t_{\text{REL}} \geq 1/2 \quad \text{and} \quad \Delta \geq 2. \quad (7)$$

Finally, we recall that the notation $a \ll b$ means that the ratio a/b vanishes as our implicit parameter n tends to infinity. Similarly, $a \lesssim b$ means that a/b is bounded above uniformly in n . We are now ready to state our main result, in which *non-negative curvature* can be understood either in the Ollivier–Ricci or the Bakry–Émery sense.

Theorem 1 (Main result). *Consider a sequence of irreducible transition matrices with symmetric support and non-negative curvature. Suppose that for every fixed $\varepsilon \in (0, 1)$, we have*

$$t_{\text{MIX}}(\varepsilon) \gg (t_{\text{REL}} \log \Delta)^2. \tag{8}$$

Then the sequence exhibits cutoff. More precisely, for every $\varepsilon \in (0, 1/2)$, we have

$$t_{\text{MIX}}(\varepsilon) - t_{\text{MIX}}(1 - \varepsilon) \lesssim \sqrt{t_{\text{MIX}}(1/4)} t_{\text{REL}} \log \Delta.$$

Let us now comment on the “product-like” condition (8). First, it is *effective*, in the sense that its verification only requires comparing the orders of magnitude of t_{MIX} and t_{REL} , as promised. Second, it implies the original product condition, by (7). Third, the presence of an additional sparsity term is unavoidable, because of the generic counter-example in [43, Example 18.7]. More precisely, let P be the transition matrix of the random walk with increment law μ on an abelian group \mathcal{X} (such chains are non-negatively curved), where the pair (\mathcal{X}, μ) is chosen so that $t_{\text{REL}} \lesssim 1 \ll t_{\text{MIX}}(1/4)$. Then we can destroy cutoff without affecting this property by simply replacing μ with $(1 - \theta)\mu + \theta\pi$, where $\theta \in (0, 1)$ satisfies $\frac{1}{t_{\text{MIX}}(1/4)} \ll \theta \ll 1$. However, this perturbation will drastically increase the sparsity parameter Δ , and the role of the latter in our condition (8) is precisely to preclude this type of pathologies. Finally, we emphasize that there is a variety of methods for estimating the orders of magnitude of $t_{\text{MIX}}(\varepsilon)$ and t_{REL} ; see the books [43, 51]. For example, a simple diameter bound (see Lemma 11 below) ensures that our “product-like” condition (8) holds whenever

$$\text{diam}(\mathcal{X}) \gg (t_{\text{REL}} \log \Delta)^2.$$

This condition may be further simplified by using the estimate $\text{diam}(\mathcal{X}) \geq \log N / \log \Delta$, where $N = |\mathcal{X}|$ denotes the number of states. As a consequence, we readily deduce that non-negatively curved chains with reasonably good expansion exhibit cutoff, in the following sense.

Corollary 2 (Non-negative curvature and expansion imply cutoff). *A sufficient condition for a sequence of non-negatively curved chains with symmetric support to exhibit cutoff is that*

$$t_{\text{REL}} \ll \frac{(\log N)^{1/2}}{(\log \Delta)^{3/2}}. \tag{9}$$

This result applies to a variety of chains. To illustrate this, let us consider the important special case where P is the transition matrix of simple random walk on a Cayley graph $G = \text{Cay}(\mathcal{X}, S)$, where $(\mathcal{X}, +)$ is a finite abelian group and $S \subseteq \mathcal{X}$ a symmetric set of generators. Specifically,

$$P(x, y) = \frac{1}{|S|} \sum_{z \in S} \mathbf{1}_{(y=x+z)}.$$

Write $N = |\mathcal{X}|$ for the number of group elements, and $d = |S|$ for the number of generators.

Corollary 3 (Abelian graphs). *Random walk on abelian Cayley graphs exhibits cutoff whenever*

$$t_{\text{REL}} \ll \frac{(\log N)^{1/2}}{(\log d)^{3/2}}.$$

In particular, this applies to random instances. Indeed, a celebrated result of Alon & Roichman [3], refined by Pak [56], by Naor [52], and finally by Hermon & Olesker-Taylor [31], asserts that *almost all Cayley graphs with $d \geq (1 + \varepsilon) \log_2 N$ satisfy $t_{\text{REL}} \lesssim 1$* , leading to the following result.

Corollary 4 (Cutoff on almost all abelian Cayley graphs). *Let G be the Cayley graph obtained by choosing d generators uniformly at random in an abelian group of size N . Consider the regime*

$$d \geq a \log_2 N \quad \text{and} \quad \log d \ll (\log N)^{1/3},$$

where $a > 1$ is any fixed constant. Then the random walk on G exhibits cutoff with high probability.

Note that the requirement $a > 1$ cannot be improved in general, since the binary group \mathbb{Z}_2^d cannot be generated by less than $d = \log_2 N$ elements. The problem of establishing cutoff for random abelian Cayley graphs has a long history (see the survey [37]). It originates with a conjecture raised by Aldous and Diaconis in an extended version of [2]. The dense regime $d \gg \log N$ was settled several years ago by Dou and Hildebrand [23, 36]. The sparse regime $d \lesssim \log N$, in contrast, was tackled only very recently in an impressive series of works by Hermon and Olesker-Taylor [31–34], which additionally provides a very detailed picture of the geometry of random Cayley graphs. We emphasize that those two approaches crucially rely on the fact that the generators are chosen uniformly at random, whereas our Corollary 3 deterministically applies to *any* abelian Cayley graph with reasonably good expansion, without requiring any specific computation.

Remark 1 (Refinements). With some additional knowledge on the chain, our approach can easily be refined. For example, the “product-like” condition (8) can be replaced with

$$t_{\text{MIX}}(\varepsilon) \gg \frac{t_{\text{REL}} \log \Delta}{\sqrt{\kappa}},$$

which is strictly weaker as long as $1/\kappa \ll t_{\text{MIX}}(\varepsilon)$. See Remark 4 below for details.

Remark 2 (Characterization of the cutoff time). Our proof also provides an entropic characterization of the cutoff time; see the second part of Theorem 5 below.

Remark 3 (Bounded degrees). Condition (8) trivially holds in the bounded-degree expander regime

$$\max(\Delta, t_{\text{REL}}) \lesssim 1 \ll |\mathcal{X}|.$$

Unfortunately, such chains must have negative curvature [58].

Proof outline. The proof of Theorem 1 relies on a certain *entropic concentration phenomenon*, which we formulate in Section 2.1 below. In Section 2.2, we show that this phenomenon implies cutoff along *any* sequence of Markov chains, without any curvature or support assumption. This is inspired by a recent body of work establishing cutoff on random instances [8–10, 14–16, 19, 31, 35, 45], where entropy plays a crucial role (see also [50, 55]). Our entropic criterion can be understood as the common mechanism underlying these generic cutoff phenomena, and hopefully many more to come. Finally, in Section 2.3, we establish the entropic concentration phenomenon for non-negatively curved chains satisfying our “product-like” condition. This combines a new gradient estimate for the logarithm of the heat kernel with a local concentration inequality for Lipschitz observables. While connections between curvature and local concentration are well known (see, e.g., [24, 41, 61] and the references therein), their application to cutoff seems new and promising.

2. Proof

2.1. The entropic concentration phenomenon

Relative entropy (or *Kullback–Leibler divergence*) is a natural measure of discrepancy between a probability measure μ and a (fully supported) reference probability measure π . It is defined as

$$d_{\text{KL}}(\mu\|\pi) := \sum_{x \in \mathcal{X}} \mu(x) \log \frac{\mu(x)}{\pi(x)},$$

where \log denotes the natural logarithm. By the strict convexity of $u \mapsto u \log u$, we always have $d_{\text{KL}}(\mu\|\pi) \geq 0$, with equality if and only if $\mu = \pi$. The celebrated Pinsker inequality provides a one-sided quantitative version of this statement, by ensuring that $\|\mu - \pi\|_{\text{TV}}$ is small whenever $d_{\text{KL}}(\mu\|\pi)$ is small. This is the starting point of a variety of powerful upper bounds on mixing times (see, e.g., [12]). However, the converse relation – namely, that $\|\mu - \pi\|_{\text{TV}}$ has to be large whenever $d_{\text{KL}}(\mu\|\pi)$ is large – is much looser, because a very small region with μ/π large can boost the above sum, while being negligible from a total-variation viewpoint. A simple way to preclude this type of pathologies is to have some control on the typical fluctuations of μ/π . To do so, it is natural to consider the associated variance, beautifully called *varentropy* in a different context [27]:

$$\mathcal{V}_{\text{KL}}(\mu\|\pi) := \sum_{x \in \mathcal{X}} \mu(x) \left(\log \frac{\mu(x)}{\pi(x)} - d_{\text{KL}}(\mu\|\pi) \right)^2.$$

Coming back to Markov chains, we will consider the worst-case varentropy under the heat kernel:

$$\mathcal{V}_{\text{KL}}^*(t) := \max_{o \in \mathcal{X}} \mathcal{V}_{\text{KL}}(\mathcal{P}_t(o, \cdot)\|\pi).$$

This key quantity turns out to govern the cutoff phenomenon, as formalized in the following result (proved in Section 2.2 below). We emphasize that the latter applies to

any transition matrix P : neither non-negative curvature, nor symmetry of the support is required. To the best of our knowledge, the use of varentropy in relation with the cutoff phenomenon is new.

Theorem 5 (Entropic concentration implies cutoff). *For any stochastic matrix P and any $\varepsilon \in (0, 1/2)$,*

$$t_{\text{MIX}}(\varepsilon) - t_{\text{MIX}}(1 - \varepsilon) \leq \frac{2t_{\text{REL}}}{\varepsilon^2} \left[1 + \sqrt{\mathcal{V}_{\text{KL}}^*(t_{\text{MIX}}(1 - \varepsilon))} \right].$$

In particular, for a sequence of stochastic matrices to exhibit cutoff, it is enough that it satisfies

$$1 + \sqrt{\mathcal{V}_{\text{KL}}^*(t_{\text{MIX}}(\varepsilon))} \ll \frac{t_{\text{MIX}}(\varepsilon)}{t_{\text{REL}}} \quad (10)$$

for all $\varepsilon \in (1/2, 1)$. Moreover, in that case, we have $t_{\text{MIX}}(\varepsilon) \sim t$, where t solves the equation

$$d_{\text{KL}}^*(t) \asymp 1 + \sqrt{\mathcal{V}_{\text{KL}}^*(t)}$$

with $d_{\text{KL}}^(t) = \max_{o \in \mathcal{X}} d_{\text{KL}}(\mathcal{P}_t(o, \cdot) \| \pi)$, and with \asymp denoting equality up to any fixed prefactor.*

We naturally call (10) the *entropic concentration phenomenon*. Observe that the latter readily implies the product condition (2). Of course, to make our criterion effective, we need to complement it with an estimate on the varentropy $\mathcal{V}_{\text{KL}}^*$. This is precisely the aim of our second key result, established in Section 2.3 below, and which crucially exploits non-negative curvature.

Theorem 6 (Varentropy estimate). *Consider a sequence of non-negatively curved transitions matrices with symmetric support. Fix $\varepsilon \in (0, 1)$, and suppose that $\sqrt{t_{\text{REL}}} \ll t_{\text{MIX}}(\varepsilon)$. Then*

$$\mathcal{V}_{\text{KL}}^*(t_{\text{MIX}}(\varepsilon)) \lesssim t_{\text{MIX}}(\varepsilon) (\log \Delta)^2.$$

When combined together, Theorems 5 and 6 readily imply Theorem 1. The remainder of the paper is thus devoted to the proof of these two results.

2.2. Entropic concentration implies cutoff

In this section, we prove that entropic concentration implies cutoff, as stated in Theorem 5. To do so, we need a sharp, two-sided quantitative relation between entropy and mixing. We start with the following upper bound, which shows that mixing occurs quickly once relative entropy is small.

Lemma 7 (Entropic upper bound). *For all $t \in \mathbb{R}_+$ and $\varepsilon \in (0, 1)$,*

$$t_{\text{MIX}}(\varepsilon) \leq t + \frac{t_{\text{REL}}}{\varepsilon} (1 + d_{\text{KL}}^*(t)).$$

Proof. Recall that the relaxation time can be used to bound the total-variation distance to equilibrium via the following classical inequality: for any law μ on \mathcal{X} and any time $s \in \mathbb{R}_+$,

$$\|\mu \mathcal{P}_s - \pi\|_{\text{TV}} \leq \frac{e^{-s/t_{\text{REL}}}}{2} \sqrt{\left\| \frac{\mu}{\pi} \right\|_{\infty}}; \tag{11}$$

see, e.g., [51]. Now fix a law μ on \mathcal{X} , and consider the set $A \subseteq \mathcal{X}$ defined by

$$A := \left\{ x \in \mathcal{X} : \log \frac{\mu(x)}{\pi(x)} < 1 + \frac{2d_{\text{KL}}(\mu \parallel \pi)}{\varepsilon} \right\}.$$

Observe that by definition,

$$\begin{aligned} \left(1 + \frac{2d_{\text{KL}}(\mu \parallel \pi)}{\varepsilon} \right) \mu(A^c) &\leq \sum_{x \in A^c} \mu(x) \log \frac{\mu(x)}{\pi(x)} \\ &= d_{\text{KL}}(\mu \parallel \pi) + \sum_{x \in A} \mu(x) \log \frac{\pi(x)}{\mu(x)} \\ &\leq d_{\text{KL}}(\mu \parallel \pi) + \pi(A) - \mu(A) \\ &\leq d_{\text{KL}}(\mu \parallel \pi) + \mu(A^c), \end{aligned}$$

where at the third line we have used $\log u \leq u - 1$. After simplification, we are left with

$$\mu(A^c) \leq \varepsilon/2.$$

Now, let $\hat{\mu} := \mu(\cdot|A)$ be μ conditioned on A . Note that

$$\left\| \frac{\hat{\mu}}{\pi} \right\|_{\infty} = \frac{1}{\mu(A)} \max_{x \in A} \frac{\mu(x)}{\pi(x)} \leq \exp \left\{ 2 + \frac{2d_{\text{KL}}(\mu \parallel \pi)}{\varepsilon} \right\},$$

because $\mu(A) \geq 1/2 \geq 1/e$. Consequently, (11) applied to $\hat{\mu}$ yields

$$\|\hat{\mu} \mathcal{P}_s - \pi\|_{\text{TV}} \leq \frac{1}{2} \exp \left\{ 1 + \frac{d_{\text{KL}}(\mu \parallel \pi)}{\varepsilon} - \frac{s}{t_{\text{REL}}} \right\}$$

for all $s \geq 0$. To make the right-hand side less than $\varepsilon/2$, we choose

$$s := \frac{t_{\text{REL}}}{\varepsilon} (1 + d_{\text{KL}}(\mu \parallel \pi)). \tag{12}$$

On the other hand, we trivially have

$$\|\hat{\mu} \mathcal{P}_s - \mu \mathcal{P}_s\|_{\text{TV}} \leq \|\hat{\mu} - \mu\|_{\text{TV}} = \mu(A^c) \leq \varepsilon/2.$$

By the triangle inequality, we conclude that for s as in (12),

$$\|\mu \mathcal{P}_s - \pi\|_{\text{TV}} \leq \varepsilon.$$

Since μ is arbitrary, we may take $\mu = \mathcal{P}_t(o, \cdot)$, and then maximize over $o \in \mathcal{X}$ to conclude. ■

To complement the above estimate, we now provide a lower bound showing that mixing cannot occur until the relative entropy has reached a sufficiently low level.

Lemma 8 (Entropic lower bound). *For any probability measure μ on \mathcal{X} and any $\varepsilon \in (0, 1)$,*

$$\|\mu - \pi\|_{\text{TV}} \leq 1 - \varepsilon \implies d_{\text{KL}}(\mu\|\pi) \leq \frac{1 + \sqrt{\mathcal{V}_{\text{KL}}(\mu\|\pi)}}{\varepsilon}.$$

Proof. Consider the event $A \subseteq \mathcal{X}$ defined by

$$A := \{x \in \mathcal{X} : \mu(x) \geq \pi(x)e^\theta\} \quad \text{with} \quad \theta = d_{\text{KL}}(\mu\|\pi) - \frac{\sqrt{\mathcal{V}_{\text{KL}}(\mu\|\pi)}}{\varepsilon}.$$

Since $\log \frac{\mu}{\pi}$ has mean $d_{\text{KL}}(\mu\|\pi)$ and variance $\mathcal{V}_{\text{KL}}(\mu\|\pi)$ under μ , Chebyshev’s inequality implies

$$\mu(A) \geq 1 - \varepsilon^2.$$

On the other hand, the definition of A readily implies

$$\pi(A) \leq e^{-\theta} \mu(A).$$

Together, these two inequalities imply

$$\mu(A) - \pi(A) \geq (1 - \varepsilon^2)(1 - e^{-\theta}).$$

Assuming that $\|\mu - \pi\|_{\text{TV}} \leq 1 - \varepsilon$, we deduce that $1 - \varepsilon \geq (1 - \varepsilon^2)(1 - e^{-\theta})$, or equivalently,

$$\theta \leq \log(1 + 1/\varepsilon).$$

Since $\log(1 + u) \leq u$, this implies $\theta \leq 1/\varepsilon$, as desired. ■

With these lemmas at hand, the proof of Theorem 5 is now straightforward.

Proof of Theorem 5. Fix $\varepsilon \in (0, 1)$ and set $t = t_{\text{MIX}}(1 - \varepsilon)$. By Lemma 7, we have

$$t_{\text{MIX}}(\varepsilon) \leq t + \frac{t_{\text{REL}}}{\varepsilon} (1 + d_{\text{KL}}^*(t)).$$

On the other hand, Lemma 8 with $\mu = \mathcal{P}_t(o, \cdot)$ (followed by a maximization over o) forces

$$d_{\text{KL}}^*(t) \leq \frac{1 + \sqrt{\mathcal{V}_{\text{KL}}^*(t)}}{\varepsilon}.$$

Reinserting this above and using $\varepsilon \leq 1$ immediately yields the desired claim. ■

2.3. Non-negative curvature implies entropic concentration

In this section, we prove the general varentropy estimate for non-negatively curved chains, stated in Theorem 6. Our starting point is the following local concentration inequality for Lipschitz observables. The term *local* here refers to the fact that the underlying measure

is the heat kernel itself, rather than the equilibrium measure π . Let κ denote the curvature of the chain, in either the Ollivier–Ricci or the Bakry–Émery sense. This is the only place where curvature is used.

Lemma 9 (Local concentration inequality). *For any $f: \mathcal{X} \rightarrow \mathbb{R}$ and $t \in \mathbb{R}_+$,*

$$\|\mathcal{P}_t(f^2) - (\mathcal{P}_t f)^2\|_\infty \leq \frac{1 - e^{-2t\kappa}}{2\kappa} \|f\|_{\text{LIP}}^2,$$

where the fraction is interpreted as t if $\kappa = 0$. In particular, if $\kappa \geq 0$, then

$$\|\mathcal{P}_t(f^2) - (\mathcal{P}_t f)^2\|_\infty \leq t \|f\|_{\text{LIP}}^2.$$

Proof. Our starting point is the following well-known pointwise identity, which is easily checked by differentiating both sides with respect to t (see, e.g., [61, Problem 2.12.a]):

$$\mathcal{P}_t(f^2) - (\mathcal{P}_t f)^2 = 2 \int_0^t \mathcal{P}_{t-s} \Gamma(\mathcal{P}_s f, \mathcal{P}_s f) \, ds.$$

This reduces our task to proving that the integrand is at most $e^{-2\kappa s} \|f\|_{\text{LIP}}^2$. In the Bakry–Émery case, we can use the subcommutativity property (5) to write

$$2\Gamma(\mathcal{P}_s f, \mathcal{P}_s f) \leq 2e^{-2\kappa s} \mathcal{P}_s \Gamma(f, f) \leq e^{-2\kappa s} \|f\|_{\text{LIP}}^2,$$

where the last equality follows from the trivial bound $2\Gamma(f, f) \leq \|f\|_{\text{LIP}}^2$ and the stochasticity of the operator \mathcal{P}_s . On the other hand, in the Ollivier–Ricci case, we use (4) to write

$$2\Gamma(\mathcal{P}_s f, \mathcal{P}_s f) \leq \|\mathcal{P}_s f\|_{\text{LIP}}^2 \leq e^{-2\kappa s} \|f\|_{\text{LIP}}^2.$$

In either case, we obtain $2\Gamma(\mathcal{P}_s f, \mathcal{P}_s f) \leq e^{-2\kappa s} \|f\|_{\text{LIP}}^2$. Since this uniform bound is trivially preserved under the action of the stochastic operator \mathcal{P}_{t-s} , the claim is established. ■

Remark 4 (Refinement). The second part of the lemma uses the crude bound

$$\frac{1 - e^{-2t\kappa}}{2\kappa} \leq t,$$

which has the advantage of suppressing the dependency on the curvature. However, in situations where a quantitative lower bound on κ is known, the alternative bound

$$\frac{1 - e^{-2t\kappa}}{2\kappa} \leq \frac{1}{2\kappa}$$

might be preferable, and leads to the refined condition $t_{\text{MIX}}(\varepsilon) \gg \frac{t_{\text{REL}} \log \Delta}{\sqrt{\kappa}}$ mentioned in Remark 1.

Applying Lemma 9 to the observable $f(x) = \log \frac{\mathcal{P}_t(o, x)}{\pi(x)}$ readily yields the varentropy estimate

$$\forall t \geq 0, \quad \mathcal{V}_{\text{KL}}^*(t) \leq t \max_{o \in \mathcal{X}} \left\| \log \frac{\mathcal{P}_t(o, \cdot)}{\pi(\cdot)} \right\|_{\text{LIP}}^2. \tag{13}$$

This reduces our task to obtaining a gradient estimate on the logarithm of the heat kernel. While such estimates have been explored in the setting of diffusions on manifolds (see, e.g., [38]), we could not find any satisfactory analogue on discrete state spaces. Here is what we can prove.

Lemma 10 (Logarithmic gradient estimate). *If P has symmetric support, then*

$$\left\| \log \frac{\mathcal{P}_t(o, \cdot)}{\pi(\cdot)} \right\|_{\text{LIP}} \leq 3(1 + \log \Delta)$$

for any initial state $o \in \mathcal{X}$ and any time $t \geq \text{diam}(\mathcal{X})/4$.

Proof. Fix an initial state $o \in \mathcal{X}$ and a time $t > 0$, and define $f: \mathcal{X} \rightarrow (0, \infty)$ by

$$f(x) := \frac{\mathcal{P}_t(o, x)}{\pi(x)}.$$

Let also q denote the Poisson distribution with mean t , i.e.

$$\forall k \in \mathbb{Z}_+, \quad q(k) = \frac{t^k e^{-t}}{k!}.$$

It follows from the definitions of P^* , f , and \mathcal{P}_t that for all $x, y \in \mathcal{X}$,

$$\pi(x)P^*(x, y)f(y) = P(y, x)\mathcal{P}_t(o, y) = \sum_{k=0}^{\infty} q(k)P^k(o, y)P(y, x).$$

Summing over all $y \in \mathcal{X}$ and using the Poisson identity $tq(k) = (k + 1)q(k + 1)$, we arrive at

$$t\pi(x)(P^*f)(x) = \sum_{k=0}^{\infty} kq(k)P^k(o, x).$$

We may now divide both sides by $\mathcal{P}_t(o, x) = \sum_k q(k)P^k(o, x)$ to obtain

$$\frac{t(P^*f)(x)}{f(x)} = \frac{\sum_{k=0}^{\infty} kq(k)P^k(o, x)}{\sum_{k=0}^{\infty} q(k)P^k(o, x)}.$$

By Jensen’s inequality, we deduce from this expression that

$$\frac{t(P^*f)(x)}{f(x)} \leq \log \left(\frac{\sum_{k=0}^{\infty} e^k q(k)P^k(o, x)}{\sum_{k=0}^{\infty} q(k)P^k(o, x)} \right) \leq \log \left(\frac{e^{t(e-1)}}{\mathcal{P}_t(o, x)} \right),$$

where the second inequality simply uses the bound $P^k(o, x) \leq 1$ in the numerator. In other words,

$$\sum_{y \in \mathcal{X}} P^*(x, y) \frac{f(y)}{f(x)} \leq e - 1 + \frac{1}{t} \log \frac{1}{\mathcal{P}_t(o, x)}.$$

Using the notation $(\Delta(Q))^{-1}$ for the smallest non-zero entry of a matrix Q , this readily implies

$$\max_{y \sim x} \frac{f(y)}{f(x)} \leq \Delta(P^*) \left(e - 1 + \frac{\log \Delta(\mathcal{P}_t)}{t} \right).$$

Taking logarithms, we obtain the logarithmic gradient estimate

$$\|\log f\|_{\text{LIP}} \leq \log \Delta(P^*) + \log\left(e - 1 + \frac{\log \Delta(\mathcal{P}_t)}{t}\right). \tag{14}$$

It now only remains to bound $\Delta(P^*)$ and $\Delta(\mathcal{P}_t)$. For the former, we simply remark that

$$\Delta(P^*) \leq \Delta^2(P) = \Delta^2, \tag{15}$$

as is easily deduced from the identity $P^*(x, y)P^*(y, x) = P(x, y)P(y, x)$ and the symmetry of the support of P . To estimate $\Delta(\mathcal{P}_t)$, we consider the $\frac{3}{4}$ -idle transition matrix $\widehat{P} = \frac{3}{4}\text{Id} + \frac{1}{4}P$. Note that for $k = \text{diam}(\mathcal{X})$, all entries of \widehat{P}^k are at least $(\frac{1}{4\Delta})^k$. Consequently, for all $x, y \in \mathcal{X}$,

$$k \geq \text{diam}(\mathcal{X}) \implies \widehat{P}^k(x, y) \geq \left(\frac{1}{4\Delta}\right)^{\text{diam}(\mathcal{X})}.$$

Multiplying by $e^{-4t}(4t)^k/k!$ and summing over all $k \geq \text{diam}(\mathcal{X})$, we obtain

$$\mathcal{P}_t(x, y) \geq p\left(\frac{1}{4\Delta}\right)^{\text{diam}(\mathcal{X})},$$

where p denotes the probability that a Poisson variable with mean $4t$ is at least $\text{diam}(\mathcal{X})$. Choosing $t \geq \text{diam}(\mathcal{X})/4$ makes this probability at least $1/2$, and we deduce that

$$\Delta(\mathcal{P}_t) \leq 2(4\Delta)^{\text{diam}(\mathcal{X})}. \tag{16}$$

Inserting the estimates (15)–(16) into (14) easily yields the claim. ■

Our last ingredient is the following elementary diameter bound, in which the fact that ε may be taken arbitrarily close to 1 is crucial.

Lemma 11 (Diameter bound). *For any $\varepsilon \in (0, 1)$, we have*

$$\text{diam}(\mathcal{X}) \leq 2t_{\text{MIX}}(\varepsilon) + \sqrt{\frac{8t_{\text{MIX}}(\varepsilon)}{1-\varepsilon}} + \sqrt{\frac{8t_{\text{REL}}}{1-\varepsilon}}.$$

Proof. Fix $\varepsilon \in (0, 1)$, and set $t = t_{\text{MIX}}(\varepsilon)$. By definition, we have

$$\mathcal{P}_t(o, A) \leq \pi(A) + \varepsilon \tag{17}$$

for any initial state $o \in \mathcal{X}$ and any event $A \subseteq \mathcal{X}$. Let us consider the specific choice

$$A := \left\{x \in \mathcal{X} : \text{dist}(o, x) \leq t + \sqrt{\frac{2t}{1-\varepsilon}}\right\}.$$

Note that by Chebyshev’s inequality,

$$\mathcal{P}_t(o, A) > \frac{1 + \varepsilon}{2},$$

because the distance to the origin at time t is stochastically dominated by a Poisson random variable with mean t . In view of (17), we deduce that $\pi(A) > (1 - \varepsilon)/2$, i.e.

$$\mathbb{P}\left(\text{dist}(o, U) \leq t + \sqrt{\frac{2t}{1 - \varepsilon}}\right) > \frac{1 - \varepsilon}{2}, \quad (18)$$

where U denotes a random variable with distribution π . On the other hand, the function $f: x \mapsto \text{dist}(o, x)$ is trivially Lipschitz on \mathcal{X} , so the Poincaré inequality (6) implies

$$\text{Var}(\text{dist}(o, U)) \leq t_{\text{REL}}.$$

By Chebyshev's inequality again, we deduce that

$$\mathbb{P}\left(\text{dist}(o, U) - \mathbb{E}[\text{dist}(o, U)] \in \left[-\sqrt{\frac{2t_{\text{REL}}}{1 - \varepsilon}}, \sqrt{\frac{2t_{\text{REL}}}{1 - \varepsilon}}\right]\right) \geq \frac{1 - \varepsilon}{2}. \quad (19)$$

Thus, the events in (18)–(19) must intersect. In other words, for all $o \in \mathcal{X}$,

$$\mathbb{E}[\text{dist}(o, U)] \leq t + \sqrt{\frac{2t}{1 - \varepsilon}} + \sqrt{\frac{2t_{\text{REL}}}{1 - \varepsilon}}.$$

The triangle inequality $\text{dist}(o, o') \leq \mathbb{E}[\text{dist}(o, U) + \text{dist}(o', U)]$ completes the proof. ■

We may at last establish Theorem 6.

Proof of Theorem 6. Fix $\varepsilon \in (0, 1)$ and suppose that $\sqrt{t_{\text{REL}}} \ll t_{\text{MIX}}(\varepsilon)$. By Lemma 11, we have

$$\text{diam}(\mathcal{X}) \leq (2 + o(1))t_{\text{MIX}}(\varepsilon).$$

In particular, the condition $t_{\text{MIX}}(\varepsilon) \geq \text{diam}(\mathcal{X})/4$ is eventually satisfied. Consequently, Lemma 10 applies with $t = t_{\text{MIX}}(\varepsilon)$. Combining this with inequality (13), we obtain

$$\mathcal{V}_{\text{KL}}^*(t_{\text{MIX}}(\varepsilon)) \leq 9t_{\text{MIX}}(\varepsilon)(1 + \log \Delta)^2$$

for n sufficiently large. Since $\Delta \geq 2$, we have $18(1 + \log \Delta)^2 \lesssim \log^2 \Delta$, as desired. ■

Acknowledgments. The author warmly thanks Nathanaël Berestycki, Max Fathi, Jonathan Hermon and Sam Olesker-Taylor for their valuable comments on a first version of the paper.

Funding. This work was partly supported by Institut Universitaire de France.

References

- [1] Aldous, D.: [Random walks on finite groups and rapidly mixing Markov chains](#). In: Seminar on probability, XVII, Lecture Notes in Mathematics 986, Springer, Berlin, 243–297 (1983) Zbl 0514.60067 MR 770418
- [2] Aldous, D., Diaconis, P.: [Shuffling cards and stopping times](#). Amer. Math. Monthly **93**, 333–348 (1986) Zbl 0603.60006 MR 841111
- [3] Alon, N., Roichman, Y.: [Random Cayley graphs and expanders](#). Random Structures Algorithms **5**, 271–284 (1994) Zbl 0798.05048 MR 1262979

- [4] Bakry, D.: *Étude des transformations de Riesz dans les variétés riemanniennes à courbure de Ricci minorée*. In: Séminaire de Probabilités, XXI, Lecture Notes in Mathematics 1247, Springer, Berlin, 137–172 (1987) Zbl [0629.58018](#) MR [941980](#)
- [5] Bakry, D., Émery, M.: *Diffusions hypercontractives*. In: Séminaire de probabilités, XIX, 1983/84, Lecture Notes in Mathematics 1123, Springer, Berlin, 177–206 (1985) Zbl [0561.60080](#) MR [889476](#)
- [6] Bakry, D., Gentil, I., Ledoux, M.: *Analysis and geometry of Markov diffusion operators*. Grundlehren der mathematischen Wissenschaften 348, Springer, Cham (2014) Zbl [1376.60002](#) MR [3155209](#)
- [7] Basu, R., Hermon, J., Peres, Y.: *Characterization of cutoff for reversible Markov chains*. Ann. Probab. **45**, 1448–1487 (2017) Zbl [1374.60129](#) MR [3650406](#)
- [8] Ben-Hamou, A.: *A threshold for cutoff in two-community random graphs*. Ann. Appl. Probab. **30**, 1824–1846 (2020) Zbl [1472.60117](#) MR [4132638](#)
- [9] Ben-Hamou, A., Salez, J.: *Cutoff for nonbacktracking random walks on sparse random graphs*. Ann. Probab. **45**, 1752–1770 (2017) Zbl [1372.60101](#) MR [3650414](#)
- [10] Berestycki, N., Lubetzky, E., Peres, Y., Sly, A.: *Random walks on the random graph*. Ann. Probab. **46**, 456–490 (2018) Zbl [1393.60077](#) MR [3758735](#)
- [11] Berestycki, N., Şengül, B.: *Cutoff for conjugacy-invariant random walks on the permutation group*. Probab. Theory Related Fields **173**, 1197–1241 (2019) Zbl [1411.60012](#) MR [3936154](#)
- [12] Bobkov, S. G., Tetali, P.: *Modified logarithmic Sobolev inequalities in discrete settings*. J. Theoret. Probab. **19**, 289–336 (2006) Zbl [1113.60072](#) MR [2283379](#)
- [13] Bordenave, C.: *A new proof of Friedman’s second eigenvalue theorem and its extension to random lifts*. Ann. Sci. École Norm. Sup. (4) **53**, 1393–1439 (2020) Zbl [1462.05324](#) MR [4203039](#)
- [14] Bordenave, C., Caputo, P., Salez, J.: *Random walk on sparse random digraphs*. Probab. Theory Related Fields **170**, 933–960 (2018) Zbl [1383.05294](#) MR [3773804](#)
- [15] Bordenave, C., Caputo, P., Salez, J.: *Cutoff at the “entropic time” for sparse Markov chains*. Probab. Theory Related Fields **173**, 261–292 (2019) Zbl [1480.60202](#) MR [3916108](#)
- [16] Bordenave, C., Lacoïn, H.: *Cutoff at the entropic time for random walks on covered expander graphs*. J. Inst. Math. Jussieu **21**, 1571–1616 (2022) Zbl [1496.05177](#) MR [4476123](#)
- [17] Bordewich, M., Dyer, M.: *Path coupling without contraction*. J. Discrete Algorithms **5**, 280–292 (2007) Zbl [1137.60046](#) MR [2316551](#)
- [18] Chen, G.-Y., Saloff-Coste, L.: *The cutoff phenomenon for ergodic Markov processes*. Electron. J. Probab. **13**, no. 3, 26–78 (2008) Zbl [1190.60007](#) MR [2375599](#)
- [19] Conchon-Kerjan, G.: *Cutoff for random lifts of weighted graphs*. Ann. Probab. **50**, 304–338 (2022) Zbl [1486.05284](#)
- [20] Diaconis, P.: *The cutoff phenomenon in finite Markov chains*. Proc. Nat. Acad. Sci. U.S.A. **93**, 1659–1664 (1996) Zbl [0849.60070](#) MR [1374011](#)
- [21] Diaconis, P., Shahshahani, M.: *Generating a random permutation with random transpositions*. Z. Wahrsch. Verw. Gebiete **57**, 159–179 (1981) Zbl [0485.60006](#) MR [626813](#)
- [22] Ding, J., Lubetzky, E., Peres, Y.: *Total variation cutoff in birth-and-death chains*. Probab. Theory Related Fields **146**, 61–85 (2010) Zbl [1190.60005](#) MR [2550359](#)
- [23] Dou, C., Hildebrand, M.: *Enumeration and random random walks on finite groups*. Ann. Probab. **24**, 987–1000 (1996) Zbl [0862.60006](#) MR [1404540](#)
- [24] Eldan, R., Lee, J. R., Lehec, J.: *Transport-entropy inequalities and curvature in discrete-space Markov chains*. In: A journey through discrete mathematics, Springer, Cham, 391–406 (2017) Zbl [1391.60173](#) MR [3726607](#)
- [25] Erbar, M., Fathi, M.: *Poincaré, modified logarithmic Sobolev and isoperimetric inequalities for Markov chains with non-negative Ricci curvature*. J. Funct. Anal. **274**, 3056–3089 (2018) Zbl [1386.53020](#) MR [3782987](#)

- [26] Erbar, M., Henderson, C., Menz, G., Tetali, P.: [Ricci curvature bounds for weakly interacting Markov chains](#). *Electron. J. Probab.* **22**, art. 40, 23 pp. (2017) Zbl 1362.60084 MR 3646066
- [27] Fradelizi, M., Li, J., Madiman, M.: [Concentration of information content for convex measures](#). *Electron. J. Probab.* **25**, art. 20, 22 pp. (2020) Zbl 1445.60028 MR 4073681
- [28] Friedman, J.: [A proof of Alon's second eigenvalue conjecture and related problems](#). *Mem. Amer. Math. Soc.* **195**, no. 910, viii+100 pp. (2008) Zbl 1177.05070 MR 2437174
- [29] Ganguly, S., Seo, I.: [Information percolation and cutoff for the random-cluster model](#). *Random Structures Algorithms* **57**, 770–822 (2020) Zbl 1453.82077 MR 4144084
- [30] Hermon, J.: [Cutoff for Ramanujan graphs via degree inflation](#). *Electron. Commun. Probab.* **22**, art. 45, 10 pp. (2017) Zbl 1378.05188 MR 3693771
- [31] Hermon, J., Olesker-Taylor, S.: [Cutoff for almost all random walks on abelian groups](#). arXiv:2102.02809 (2021)
- [32] Hermon, J., Olesker-Taylor, S.: [Cutoff for random walks on upper triangular matrices](#). arXiv:1911.02974v2 (2021)
- [33] Hermon, J., Olesker-Taylor, S.: [Further results and discussions on random Cayley graphs](#). arXiv:1911.02975v2 (2021)
- [34] Hermon, J., Olesker-Taylor, S.: [Geometry of random Cayley graphs of abelian groups](#). arXiv:2102.02801 (2021)
- [35] Hermon, J., Sly, A., Sousi, P.: [Universality of cutoff for graphs with an added random matching](#). *Ann. Probab.* **50**, 203–240 (2022) Zbl 1486.05278 MR 4385126
- [36] Hildebrand, M.: [Random walks supported on random points of \$\mathbb{Z}/n\mathbb{Z}\$](#) . *Probab. Theory Related Fields* **100**, 191–203 (1994) MR 1296428
- [37] Hildebrand, M.: [A survey of results on random random walks on finite groups](#). *Probab. Surv.* **2**, 33–63 (2005) Zbl 1189.60016 MR 2121795
- [38] Hino, M.: [On short time asymptotic behavior of some symmetric diffusions on general state spaces](#). *Potential Anal.* **16**, 249–264 (2002) Zbl 1015.47029 MR 1885762
- [39] Jost, J.: [Riemannian geometry and geometric analysis](#). 7th ed., Universitext, Springer, Cham (2017) Zbl 1380.53001 MR 3726907
- [40] Joulin, A.: [Poisson-type deviation inequalities for curved continuous-time Markov chains](#). *Bernoulli* **13**, 782–798 (2007) Zbl 1131.60069 MR 2348750
- [41] Joulin, A., Ollivier, Y.: [Curvature, concentration and error estimates for Markov chain Monte Carlo](#). *Ann. Probab.* **38**, 2418–2442 (2010) Zbl 1207.65006 MR 2683634
- [42] Klartag, B., Kozma, G., Ralli, P., Tetali, P.: [Discrete curvature and abelian groups](#). *Canad. J. Math.* **68**, 655–674 (2016) Zbl 1341.53068 MR 3492631
- [43] Levin, D. A., Peres, Y.: [Markov chains and mixing times](#). American Mathematical Society, Providence, RI (2017) Zbl 1390.60001 MR 3726904
- [44] Lubetzky, E., Peres, Y.: [Cutoff on all Ramanujan graphs](#). *Geom. Funct. Anal.* **26**, 1190–1216 (2016) Zbl 1351.05208 MR 3558308
- [45] Lubetzky, E., Sly, A.: [Cutoff phenomena for random walks on random regular graphs](#). *Duke Math. J.* **153**, 475–510 (2010) Zbl 1202.60012 MR 2667423
- [46] Lubetzky, E., Sly, A.: [Cutoff for general spin systems with arbitrary boundary conditions](#). *Comm. Pure Appl. Math.* **67**, 982–1027 (2014) Zbl 1292.82025 MR 3193965
- [47] Lubetzky, E., Sly, A.: [An exposition to information percolation for the Ising model](#). *Ann. Fac. Sci. Toulouse Math. (6)* **24**, 745–761 (2015) Zbl 1333.60207 MR 3434254
- [48] Lubetzky, E., Sly, A.: [Information percolation and cutoff for the stochastic Ising model](#). *J. Amer. Math. Soc.* **29**, 729–774 (2016) Zbl 1342.60173 MR 3486171
- [49] Lubetzky, E., Sly, A.: [Universality of cutoff for the Ising model](#). *Ann. Probab.* **45**, 3664–3696 (2017) Zbl 1405.60148 MR 3729612
- [50] Merle, M., Salez, J.: [Cutoff for the mean-field zero-range process](#). *Ann. Probab.* **47**, 3170–3201 (2019) Zbl 1448.60159 MR 4021248

- [51] Montenegro, R., Tetali, P.: [Mathematical aspects of mixing times in Markov chains](#). *Found. Trends Theoret. Comput. Sci.* **1**, 237–354 (2006) Zbl [1193.68138](#) MR [2341319](#)
- [52] Naor, A.: [On the Banach-space-valued Azuma inequality and small-set isoperimetry of Alon–Roichman graphs](#). *Combin. Probab. Comput.* **21**, 623–634 (2012) Zbl [1247.05104](#) MR [2942733](#)
- [53] Ollivier, Y.: [Ricci curvature of Markov chains on metric spaces](#). *J. Funct. Anal.* **256**, 810–864 (2009) Zbl [1181.53015](#) MR [2484937](#)
- [54] Ollivier, Y.: [A survey of Ricci curvature for metric spaces and Markov chains](#). In: *Probabilistic approach to geometry*, *Advanced Studies in Pure Mathematics* 57, Mathematical Society of Japan, Tokyo, 343–381 (2010) Zbl [1204.53035](#) MR [2648269](#)
- [55] Ozawa, N.: [An entropic proof of cutoff on Ramanujan graphs](#). *Electron. Commun. Probab.* **25**, art. 77, 8 pp. (2020) Zbl [1462.05329](#) MR [4178418](#)
- [56] Pak, I.: [Random Cayley graphs with \$O\(\log |G|\)\$ generators are expanders](#). In: *Algorithms—ESA '99 (Prague)*, *Lecture Notes in Computer Science* 1643, Springer, Berlin, 521–526 (1999) Zbl [0942.05061](#) MR [1729149](#)
- [57] Peres, Y.: Talk at American Institute of Mathematics (AIM) research workshop “Sharp thresholds for mixing times” (Palo Alto, 2004). Summary available at <http://www.aimath.org/WWN/mixingtimes>
- [58] Salez, J.: [Sparse expanders have negative curvature](#). *Geom. Funct. Anal.* **32**, 1486–1513 (2022) Zbl [07648711](#) MR [4536468](#)
- [59] Saloff-Coste, L.: [Random walks on finite groups](#). In: *Probability on discrete structures*, *Encyclopaedia of Mathematical Sciences* 110, Springer, Berlin, 263–346 (2004) Zbl [1049.60006](#) MR [2023654](#)
- [60] Siconolfi, V.: [Ricci curvature, Bruhat graphs and Coxeter groups](#). *Proc. Amer. Math. Soc.* **151**, 17–27 (2023) Zbl [1503.05128](#) MR [4504604](#)
- [61] van Handel, R.: *Probability in high dimension*. Lecture notes (2016)
- [62] Villemonais, D.: [Lower bound for the coarse Ricci curvature of continuous-time pure-jump processes](#). *J. Theoret. Probab.* **33**, 954–991 (2020) Zbl [1451.60025](#) MR [4091580](#)