

Cutoff for non-negatively curved Markov chains

An addition to Justin Salez' lecture notes

J.M. Swart

December 1, 2022

Contents

1	Continuous-time Markov chains	2
2	The relaxation and mixing times	3
3	Reversibility	5
4	Covariance formulas	7
5	The Wasserstein distance	10
6	Curvature	16
7	Cut-off	18
8	Strategy of the proof	19

1 Continuous-time Markov chains

If A is any (finite, square, real or complex) matrix, then we define

$$e^A := \sum_{n=0}^{\infty} \frac{1}{n!} A^n, \quad (1.1)$$

where $A^0 := I$, the identity matrix. If A and B commute (i.e., $AB = BA$), then one can check that $e^A e^B = e^{A+B}$, but this formula does not hold in general if A and B do not commute. Let $s, t \geq 0$ be real numbers. Then certainly sA commutes with tA , so

$$e^{sA} e^{tA} = e^{(s+t)A} \quad (s, t \geq 0). \quad (1.2)$$

It is also clear that $e^{0A} = I$, so the operators $(e^{tA})_{t \geq 0}$ form a semigroup. It is easy to check that

$$\frac{\partial}{\partial t} e^{tA} = A e^{tA} \quad (t \geq 0). \quad (1.3)$$

Another useful formula, that is easy to prove, is

$$e^{tA} = \lim_{n \rightarrow \infty} (I + \frac{t}{n} A)^n. \quad (1.4)$$

Let \mathcal{X} be a finite set and let P be a probability kernel on \mathcal{X} . For any integer $n \geq 0$, we let P^n denote the n -th matrix power of P , which corresponds to the n -step transition kernel of the discrete-time Markov chain with transition kernel P . For any real $t \geq 0$, we define

$$P_t := e^{t(P-I)} \quad (t \geq 0). \quad (1.5)$$

Then clearly, $P_0 = I$ and $P_s P_t = P_{s+t}$. A simple calculation yields

$$\begin{aligned} P_t &= \sum_{n=0}^{\infty} \frac{1}{n!} t^n (P-I)^n = \sum_{n=0}^{\infty} \frac{t^n}{n!} \sum_{k=0}^n \binom{n}{k} P^k (-1)^{n-k} \\ &= \sum_{n=0}^{\infty} \sum_{k=0}^n \frac{t^n}{(n-k)! k!} P^k (-1)^{n-k} = \sum_{k=0}^{\infty} \sum_{n=k}^{\infty} \frac{(-t)^{n-k} t^k}{(n-k)! k!} P^k \\ &= \sum_{k=0}^{\infty} e^{-t} \frac{t^k}{k!} P^k = \sum_{k=0}^{\infty} p_t(k) P^k, \end{aligned} \quad (1.6)$$

where p_t is the Poisson distribution with parameter t . This means that P_t has the following interpretation. Let $(\sigma_k)_{k \geq 1}$ be i.i.d. exponentially distributed random variables with mean one. Set

$$\tau_n := \sum_{k=1}^n \sigma_k \quad (n \geq 0), \quad (1.7)$$

where $\tau_0 := 0$. Let $(X_n)_{n \geq 0}$ be a (discrete time) Markov chain with transition kernel P and an arbitrary initial law ν , and define $(\hat{X}_t)_{t \geq 0}$ by

$$\hat{X}_t := X_n \quad \text{for } \tau_n \leq t < \tau_{n+1}, \quad n \geq 0. \quad (1.8)$$

Then, for any $0 = t_0 < \dots < t_m$ and $x_0, \dots, x_m \in \mathcal{X}$, one has

$$\mathbb{P}[\hat{X}_{t_0} = x_0, \dots, \hat{X}_{t_m} = x_m] = \nu(x_0)P_{t_1-t_0}(x_0, x_1) \cdots P_{t_m-t_{m-1}}(x_{m-1}, x_m). \quad (1.9)$$

Indeed, the set $\{\tau_n : n \geq 1\}$ is a Poisson point set with intensity 1, which means that the number of jumps made by the process $(\hat{X}_t)_{t \geq 0}$ in the time interval $[t_{k-1}, t_k)$ has a Poisson distribution with parameter $t_k - t_{k-1}$, and disjoint time intervals are independent.

We recall that if P is a probability kernel, then we define its associated *lazy kernel* as $\frac{1}{2}(P + I)$. Applying (1.4) to $A = P - I$, we see that

$$P_t = \lim_{n \rightarrow \infty} \left(\frac{t}{n}P + \left(1 - \frac{t}{n}\right)I \right)^n. \quad (1.10)$$

This means that we may view continuous-time Markov chains as “extremely lazy” chains. Another way of interpreting (1.10) is as follows: we divide the interval $[0, t]$ into n pieces, and then independently in each time interval apply P with probability t/n and do nothing with the remaining probability. Letting $n \rightarrow \infty$, this means, of course, that we apply P at the times of a rate one Poisson point process.

An invariant law is a probability law π such that $\pi P = \pi$. If P is irreducible and aperiodic, then it has a unique invariant law π and $P^n(x, \cdot)$ converges to π as $n \rightarrow \infty$. Even without aperiodicity, it is true that $P_t(x, \cdot)$ converges to π as $t \rightarrow \infty$. This should not surprise us, since lazy chains are always aperiodic and continuous-time Markov chains are extremely lazy.

2 The relaxation and mixing times

If μ, ν are probability measures on \mathcal{X} , then we let $\Pi(\mu, \nu)$ denote the space of all probability measures γ on $\mathcal{X} \times \mathcal{X}$ whose first and second marginals are μ and ν , respectively:

$$\mu(x) = \sum_{y \in \mathcal{X}} \gamma(x, y) \quad \text{and} \quad \nu(y) = \sum_{x \in \mathcal{X}} \gamma(x, y). \quad (2.1)$$

We call γ a *coupling measure* for μ and ν . The total variation distance between μ and ν is given by

$$d_{\text{TV}}(\mu, \nu) := \inf_{\gamma \in \Pi(\mu, \nu)} \sum_{(x, y) \in \mathcal{X}^2} \gamma(x, y) 1_{\{x \neq y\}}. \quad (2.2)$$

More probabilistically, we can formulate this by saying that

$$d_{\text{TV}}(\mu, \nu) := \inf_{\substack{X \sim \mu \\ Y \sim \nu}} \mathbb{P}[X \neq Y], \quad (2.3)$$

where the infimum is over all possible *couplings* of random variables X and Y with laws μ and ν , respectively. One can check that the infimum is attained for a suitable coupling. There are other, simpler formulas for d_{TV} , for example

$$d_{\text{TV}}(\mu, \nu) = \frac{1}{2} \sum_{x \in \mathcal{X}} |\mu(x) - \nu(x)|. \quad (2.4)$$

If P is an irreducible probability kernel with invariant law π , then we set

$$\begin{aligned} \mathcal{D}_P(n) &:= \sup_{x \in \mathcal{X}} d_{\text{TV}}(\pi, P^n(x, \cdot)), \\ \hat{\mathcal{D}}_P(t) &:= \sup_{x \in \mathcal{X}} d_{\text{TV}}(\pi, P_t(x, \cdot)). \end{aligned} \quad (2.5)$$

It is shown in [Sal22, Lemma 11] that

$$\frac{1}{2} \tilde{\mathcal{D}}_P(n) \leq \mathcal{D}_P(n) \leq \tilde{\mathcal{D}}_P(n) \quad \text{with} \quad \tilde{\mathcal{D}}_P(n) := \sup_{x, y \in \mathcal{X}} d_{\text{TV}}(P^n(x, \cdot), P^n(y, \cdot)). \quad (2.6)$$

A similar claim holds in the continuous-time setting. We define the *mixing time* in the discrete and continuous-time settings as:

$$\begin{aligned} t_{\text{MIX}}^{(\varepsilon)}(P) &:= \inf \{n \in \mathbb{N} : \mathcal{D}_P(n) \leq \varepsilon\}, \\ \hat{t}_{\text{MIX}}^{(\varepsilon)}(P) &:= \inf \{t \geq 0 : \hat{\mathcal{D}}_P(t) \leq \varepsilon\}. \end{aligned} \quad (2.7)$$

Let $\text{Spec}(P)$ denote the spectrum of P (i.e., the set of its complex eigenvalues) and let

$$\lambda_*(P) := \max \{|\lambda| : \lambda \in \text{Spec}(P), \lambda \neq 1\}. \quad (2.8)$$

Then one can prove that

$$(\mathcal{D}_P(n))^{1/n} \xrightarrow{n \rightarrow \infty} \lambda_*(P). \quad (2.9)$$

In other words,

$$\mathcal{D}_P(n) = (\lambda_*(P) + o(1))^n, \quad (2.10)$$

where $o(1)$ is a term that tends to zero as $n \rightarrow \infty$. Defining the *relaxation time* by

$$\frac{-1}{t_{\text{REL}}(P)} := \log(\lambda_*(P)), \quad (2.11)$$

we can rewrite (2.10) as

$$\mathcal{D}_P(n) = e^{-n/t_{\text{REL}}(P)} + o(1). \quad (2.12)$$

In other words, for large n , the quantity $\mathcal{D}_P(n)$ decays exponentially fast, and the relaxation time $t_{\text{REL}}(P)$ is the time needed for this quantity to get a factor e^{-1} smaller (comparable to the half-life in nuclear decay).

In the continuous-time setting, we can define something similar. We start by observing that

$$\text{Spec}(P - I) = \{\lambda - 1 : \lambda \in \text{Spec}(P)\}. \quad (2.13)$$

We can order the (complex) eigenvalues of P according to their real parts, such that

$$1 = \lambda_1 \geq \Re(\lambda_2) \geq \dots \geq \Re(\lambda_n). \quad (2.14)$$

Letting $\lambda_2(P)$ denote the second eigenvalue in this order, we set

$$\hat{t}_{\text{REL}}(P) := \frac{1}{1 - \Re(\lambda_2(P))}. \quad (2.15)$$

(Note that this definition is unambiguous even though there may be several complex eigenvalues with the same real part as $\lambda_2(P)$.) One can prove that

$$-\frac{1}{t} \log \hat{\mathcal{D}}_P(t) \xrightarrow{t \rightarrow \infty} \frac{1}{\hat{t}_{\text{REL}}(P)}, \quad (2.16)$$

which can be rewritten as

$$\hat{\mathcal{D}}_P(t) = e^{-t/\hat{t}_{\text{REL}}(P)} + o(1). \quad (2.17)$$

For probability kernels that are irreducible but periodic, the discrete time relaxation time $t_{\text{REL}}(P)$ is infinite while $\hat{t}_{\text{REL}}(P)$ is finite. This is because in this case P has eigenvalues that are different from one, but whose absolute value is equal to one.

3 Reversibility

If π is an invariant law of P , and (X_0, \dots, X_n) is a Markov chain with transition kernel P and initial law π , then the *reversed chain* (X_n, \dots, X_0) is a Markov chain with transition kernel

$$P^*(x, y) := \pi(y)P(y, x)\pi(x)^{-1} \quad (x, y \in \mathcal{X}), \quad (3.1)$$

as follows by observing that

$$\pi(x_0) \prod_{k=1}^n P(x_{k-1}, x_k) = \pi(x_n) \prod_{k=1}^n P^*(x_k, x_{k-1}) \quad (3.2)$$

for all x_0, \dots, x_n . A Markov chain is *reversible* if $P = P^*$, i.e., if the *detailed balance equation*

$$\pi(x)P(x, y) = \pi(y)P(y, x) \quad (x, y \in \mathcal{X}) \quad (3.3)$$

holds. We can define an inner product on $\mathbb{R}^{\mathcal{X}}$ by

$$\langle f, g \rangle := \sum_{x \in \mathcal{X}} \pi(x) f(x) g(x). \quad (3.4)$$

Then it is easy to see that

$$\langle f, Pg \rangle = \langle P^* f, g \rangle, \quad (3.5)$$

so P^* is the adjoint of P with respect to the inner product $\langle \cdot, \cdot \rangle$ and reversibility is equivalent to P being self-adjoint.

If P is reversible, then there exist $\phi_1, \dots, \phi_n \in \mathbb{R}^{\mathcal{X}}$ that are eigenvectors of P , i.e.,

$$P\phi_i = \lambda_i \phi_i \quad (3.6)$$

for some $\lambda_1, \dots, \lambda_n \in \mathbb{R}$, and that are moreover *orthonormal* in the sense that

$$\langle \phi_i, \phi_j \rangle = \begin{cases} 1 & \text{if } i = j, \\ 0 & \text{if } i \neq j. \end{cases} \quad (3.7)$$

Without loss of generality we can assume that $1 = \lambda_1 \geq \dots \geq \lambda_n$. If P is irreducible, then $\lambda_1 > \lambda_2$. If P is moreover aperiodic, then also $\lambda_n > -1$ and hence $|\lambda_i| < 1$ for all $i = 2, \dots, n$. One has

$$\begin{aligned} P^n(x, y) &= \pi(y) + \pi(y) \sum_{i=2}^n \lambda_i^n \phi_i(x) \phi_i(y), \\ P_t(x, y) &= \pi(y) + \pi(y) \sum_{i=2}^n e^{(\lambda_i - 1)t} \phi_i(x) \phi_i(y). \end{aligned} \quad (3.8)$$

We observe that

$$\text{Spec}\left(\frac{1}{2}(P + I)\right) = \left\{ \frac{1}{2}(\lambda + 1) : \lambda \in \text{Spec}(P) \right\}, \quad (3.9)$$

which implies that the spectrum of a lazy, reversible kernel is contained in $[0, 1]$ and as a result

$$\lambda_2\left(\frac{1}{2}(P + 1)\right) = \lambda_\star\left(\frac{1}{2}(P + 1)\right). \quad (3.10)$$

4 Covariance formulas

For any probability law μ on \mathcal{X} and functions $f, g \in \mathbb{R}^{\mathcal{X}}$, we let

$$\text{Cov}_\mu(f, g) := \mu(fg) - (\mu f)(\mu g) \quad (4.1)$$

denote the covariance of f and g under μ . There is a nice way to calculate the covariance of two functions of a Markov chain, that is not as well-known as it should be. If P is a probability kernel on \mathcal{X} and $f, g \in \mathbb{R}^{\mathcal{X}}$, then we set

$$\Gamma_P(f, g) := \frac{1}{2}(P(fg) - (Pf)(Pg)). \quad (4.2)$$

The factor $\frac{1}{2}$ is there for historical reasons.

Lemma 4.1 (Covariance formula). *One has*

$$\text{Cov}_{\mu P^n}(f, g) = \text{Cov}_\mu(P^n f, P^n g) + 2 \sum_{k=1}^n \mu P^{n-k} \Gamma_P(P^{k-1} f, P^{k-1} g). \quad (4.3)$$

Proof The statement is trivial for $n = 0$. Fix $n \geq 1$ and for each $0 \leq k \leq n$ define a function $H_k : \mathcal{X} \rightarrow \mathbb{R}$ by

$$H_k := P^k((P^{n-k} f)(P^{n-k} g)) \quad (0 \leq k \leq n).$$

Then

$$\begin{aligned} \mu(H_n - H_0) &= \mu P^n(fg) - \mu((P^n f)(P^n g)) \\ &= [\mu P^n(fg) - (\mu P^n f)(\mu P^n g)] - [\mu((P^n f)(P^n g)) - (\mu P^n f)(\mu P^n g)] \\ &= \text{Cov}_{\mu P^n}(f, g) - \text{Cov}_\mu(P^n f, P^n g). \end{aligned}$$

It follows that

$$\begin{aligned} \text{Cov}_{\mu P^n}(f, g) - \text{Cov}_\mu(P^n f, P^n g) &= \sum_{k=1}^n \mu[H_k - H_{k-1}] \\ &= \sum_{k=1}^n \mu[P^k((P^{n-k} f)(P^{n-k} g)) - P^{k-1}((P^{n-k+1} f)(P^{n-k+1} g))] \\ &= 2 \sum_{k=1}^n \mu P^{k-1} \Gamma_P(P^{n-k} f, P^{n-k} g). \end{aligned}$$

Changing the summation order (setting $k' := n - k + 1$), we arrive at (4.3). ■

In the continuous-time setting, we have a similar formula. The *carré du champ* of f and g is the function $\hat{\Gamma}_P(f, g)$ defined as

$$\hat{\Gamma}_P(f, g)(x) := \frac{1}{2} \sum_{y \in \mathcal{X}} P(x, y) (f(y) - f(x))(g(y) - g(x)). \quad (4.4)$$

It seems there is no English name for this object. We claim that

$$2\hat{\Gamma}_P(f, g) = G(fg) - (Gf)g - f(Gg) \quad \text{with} \quad G := P - I. \quad (4.5)$$

To see this, we calculate

$$\begin{aligned} G(fg)(x) &= \sum_{y \in \mathcal{X}} P(x, y) (f(y)g(y) - f(x)g(x)) \\ &= \sum_{y \in \mathcal{X}} P(x, y) \left\{ (f(y) - f(x))(g(y) - g(x)) \right. \\ &\quad \left. + f(x)(g(y) - g(x)) + (f(y) - f(x))g(x) \right\} \\ &= 2\hat{\Gamma}_P(f, g)(x) + (Gf)(x)g(x) + f(x)(Gg)(x). \end{aligned} \quad (4.6)$$

Note that $I + \varepsilon G = (1 - \varepsilon)I + \varepsilon P$ is a probability kernel for all $\varepsilon \in [0, 1]$, and that

$$\begin{aligned} 2\Gamma_{I+\varepsilon G}(f, g) &= (I + \varepsilon G)(fg) - ((I + \varepsilon G)f)((I + \varepsilon G)g) \\ &= \varepsilon [G(fg) + (Gf)g + f(Gg)] + O(\varepsilon^2) = 2\varepsilon \hat{\Gamma}_P(f, g) + O(\varepsilon^2) \end{aligned} \quad (4.7)$$

as $\varepsilon \rightarrow 0$. This explains why $\hat{\Gamma}_P$ is the right continuous-time analogue of the object Γ_P . We state the following lemma without proof.

Lemma 4.2 (Covariance formula in continuous time). *One has*

$$\text{Cov}_{\mu_{P_t}}(f, g) = \text{Cov}_{\mu}(P_t f, P_t g) + 2 \int_0^t \mu_{P_{t-s}} \hat{\Gamma}_P(P_s f, P_s g) ds. \quad (4.8)$$

The *Dirichlet form* associated with an irreducible kernel P is the function $\hat{\mathcal{E}}_P : \mathbb{R}^{\mathcal{X}} \rightarrow \mathbb{R}$ defined as

$$\hat{\mathcal{E}}_P(f) := \sum_{x \in \mathcal{X}} \pi(x) \hat{\Gamma}_P(f, f), \quad (4.9)$$

where π is the invariant law of P . Contrary to the carré du champ, which is not very well-known outside the French literature, the Dirichlet form is a well-known and much studied object. We claim that

$$\hat{\mathcal{E}}_P(f) = -\langle f, Gf \rangle \quad \text{with} \quad G = P - I. \quad (4.10)$$

Indeed, this follows from (4.5) and the observation that for any function $h \in \mathbb{R}^{\mathcal{X}}$ (and hence in particular for $h = fg$)

$$\sum_{x \in \mathcal{X}} \pi(x) Gh(x) = \pi(P - I)h = \pi Ph - \pi h = 0. \quad (4.11)$$

Formula (4.10) is historically the oldest definition of the Dirichlet form, which explains the factor $\frac{1}{2}$ in the definition of the carré du champ.

Lemma 4.3 (Equilibrium variance). *Let P be an irreducible probability kernel with invariant law π . Then*

$$\mathrm{Var}_\pi(f) = 2 \int_0^\infty \hat{\mathcal{E}}_P(P_s f) \, ds. \quad (4.12)$$

Proof Since π is an invariant law, formula (4.8) simplifies to

$$\mathrm{Cov}_\pi(f, g) = \mathrm{Cov}_\mu(P_t f, P_t g) + 2 \int_0^t \pi \hat{\Gamma}_P(P_s f, P_s g) \, ds. \quad (4.13)$$

Setting $f = g$ and letting $t \rightarrow \infty$, using (4.9), we arrive at (4.12), where we use that since P is irreducible, $P_t f(x) \rightarrow \pi f$ for all $x \in \mathcal{X}$ and the variance of a constant function is zero. \blacksquare

Inserting (4.4) into (4.9) we see that

$$\hat{\mathcal{E}}_P(f) = \frac{1}{2} \sum_{x, y \in \mathcal{X}} \pi(x) P(x, y) (f(y) - f(x))^2. \quad (4.14)$$

We see from this formula that $\hat{\mathcal{E}}_P(f) \geq 0$, with equality if and only if f is constant. We recall that the *time-reversed* kernel P^* satisfies $\pi(x)P(x, y) = \pi(y)P^*(y, x)$. We see from (4.14) that

$$\hat{\mathcal{E}}_P(f) = \hat{\mathcal{E}}_{P^*}(f) = \hat{\mathcal{E}}_{(P+P^*)/2}(f). \quad (4.15)$$

By definition, the *Poincaré constant* is defined as

$$\gamma(P) := \inf_f \frac{\hat{\mathcal{E}}_P(f)}{\mathrm{Var}_\pi(f)}, \quad (4.16)$$

where we take the infimum over all f that are not constant. Since $\mathcal{E}_P(f)$ and $\mathrm{Var}_\pi(f)$ do not change if we add a constant to f , it suffices to take the infimum over all non-constant functions f with $\pi f = 0$. For such functions $\mathrm{Var}_\pi(f) = \langle f, f \rangle =: \|f\|_2^2$. Since the fraction does not change if we multiply f by a constant, we conclude that

$$\gamma(P) := \inf \{ \hat{\mathcal{E}}_P(f) : f \in \mathbb{R}^\mathcal{X}, \pi f = 0, \|f\|_2 = 1 \}. \quad (4.17)$$

In [Sal22, Def 19], a similar claim is made but it seems $\|f\|_2$ is replaced by the supremum norm. I do not see why this should hold. In view of (4.15),

$$\gamma(P) = \gamma(P^*) = \gamma((P + P^*)/2). \quad (4.18)$$

In [Sal22, Lemma 21], it is proved that

$$\gamma(P) = 1 - \lambda_2((P + P^*)/2). \quad (4.19)$$

In particular, by (2.15), this implies that

$$\hat{t}_{\mathrm{REL}}(P) = \frac{1}{\gamma(P)} \quad \text{if } P = P^*. \quad (4.20)$$

5 The Wasserstein distance

Assume that P is irreducible reversible. Then we can equip \mathcal{X} with the structure of a connected graph with set of edges E such that

$$\{x, y\} \in E \iff P(x, y) > 0 \iff P(y, x) > 0. \quad (5.1)$$

We define the *graph distance* by

$$\text{dist}(x, y) := \inf\{n \geq 0 : P^n(x, y) > 0\}. \quad (5.2)$$

Equivalently, $\text{dist}(x, y)$ is the length of the shortest path between x and y in the graph we have just defined. For probability measures μ, ν on \mathcal{X} , we define (compare (2.2))

$$\mathcal{W}(\mu, \nu) := \inf_{\gamma \in \Pi(\mu, \nu)} \sum_{(x, y) \in \mathcal{X}^2} \gamma(x, y) \text{dist}(x, y). \quad (5.3)$$

Note that since $\text{dist}(x, y) \leq 1_{\{x \neq y\}}$, an immediate consequence of (2.2) and (5.3) is that

$$d_{\text{TV}}(\mu, \nu) \leq \mathcal{W}(\mu, \nu). \quad (5.4)$$

The infimum in (5.3) is obtained, since it is the infimum of the continuous function

$$\gamma \mapsto \sum_{(x, y) \in \mathcal{X}^2} \gamma(x, y) \text{dist}(x, y) \quad (5.5)$$

over the compact and convex set $\Pi(\mu, \nu)$. Any γ for which the infimum is obtained is called an *optimal coupling*. We claim that the function

$$(\mu, \nu) \mapsto \mathcal{W}(\mu, \nu) \quad (5.6)$$

is convex. To see this, fix $0 < p < 1$ and (μ_i, ν_i) ($i = 1, 2$). Let γ_i be an optimal coupling for μ_i, ν_i . Then $p\gamma_1 + (1 - p)\gamma_2$ is a coupling measure for $\mu := p\mu_1 + (1 - p)\mu_2$ and $\nu := p\nu_1 + (1 - p)\nu_2$. Therefore,

$$\begin{aligned} \mathcal{W}(\mu, \nu) &\leq p \sum_{(x, y) \in \mathcal{X}^2} \gamma_1(x, y) \text{dist}(x, y) + (1 - p) \sum_{(x, y) \in \mathcal{X}^2} \gamma_2(x, y) \text{dist}(x, y) \\ &= p\mathcal{W}(\mu_1, \nu_1) + (1 - p)\mathcal{W}(\mu_2, \nu_2). \end{aligned} \quad (5.7)$$

Lemma 5.1 (Wasserstein metric). *The Wasserstein distance is a metric on the space of probability laws on \mathcal{X} .*

Proof If $\mathcal{W}(\mu, \nu) = 0$, then we can (optimally) couple random variables X and Y with laws μ and ν such that $\mathbb{E}[\text{dist}(X, Y)] = 0$. This implies $X = Y$ a.s. and hence $\mu = \nu$. It is also clear that $\mathcal{W}(\mu, \nu) = \mathcal{W}(\nu, \mu)$, so it remains to prove the triangle inequality. Let μ, ν, ρ be probability measures on \mathcal{X} . By Lemma 5.2 below, we can construct random variable X, Y, Z so that the law of (X, Y) is an optimal coupling for μ, ν , while (Y, Z) is an optimal coupling for ν, ρ . Now

$$\begin{aligned} \mathcal{W}(\mu, \rho) &= \mathbb{E}[\text{dist}(X, Z)] \\ &\leq \mathbb{E}[\text{dist}(X, Y) + \text{dist}(Y, Z)] = \mathcal{W}(\mu, \nu) + \mathcal{W}(\nu, \rho), \end{aligned} \quad (5.8)$$

where we have used the triangle inequality for the graph distance. \blacksquare

Lemma 5.2 (Combined coupling). *Let (X_1, Y_1) and (Y_2, Z_2) be random variables with values in \mathcal{X}^2 , so that Y_1 and Y_2 have the same law. Then it is possible to construct a random variable (X, Y, Z) with values in \mathcal{X}^3 such that (X, Y) has the same law as (X_1, Y_1) and (Y, Z) has the same law as (Y_2, Z_2) .*

Proof Let $\mu(y) := \mathbb{P}[Y_1 = y] = \mathbb{P}[Y_2 = y]$ ($y \in \mathcal{X}$). For each $y \in \mathcal{X}$ such that $\mu(y) > 0$, define $P_1(y, x) := \mathbb{P}[X_1 = x \mid Y_1 = y]$ and $P_2(y, z) := \mathbb{P}[Z_2 = z \mid Y_2 = y]$. If $\mu(y) = 0$, then define $P_1(y, x)$ and $P_2(y, z)$ in an arbitrary way. Then

$$\gamma(x, y, z) := \mu(y)P_1(y, x)P_2(y, z) \quad (x, y, z \in \mathcal{X}) \quad (5.9)$$

defines a probability measure on \mathcal{X}^3 , where it does not matter how $P_1(y, x)$ and $P_2(y, z)$ are defined when $\mu(y) = 0$. Let (X, Y, Z) be a random variable with law γ . Then (X, Y) has the same law as (X_1, Y_1) and (Y, Z) has the same law as (Y_2, Z_2) . Note that moreover, due to our construction, Z and X are conditionally independent given Y . \blacksquare

For any $f \in \mathbb{R}^{\mathcal{X}}$, we define the Lipschitz “norm” as

$$\|f\|_{\text{LIP}} := \sup_{\{x, y\} \in E} |f(x) - f(y)|, \quad (5.10)$$

where E is the set of edges defined in (5.1). It is not hard to see that $\|\cdot\|_{\text{LIP}}$ is a pseudonorm and that $\|f\|_{\text{LIP}} = 0$ if and only if f is constant. For any $x, y \in \mathcal{X}$, we can find $x = x_0, \dots, x_d = y$ with $d = d(x, y)$ and $\{x_{k-1}, x_k\} \in E$ for all $1 \leq k \leq d$. Using the triangle inequality for \mathcal{W} , we then get

$$|f(x) - f(y)| \leq \|f\|_{\text{LIP}} d(x, y) \quad (x, y \in \mathcal{X}). \quad (5.11)$$

The following theorem says that for any $f \in \mathbb{R}^{\mathcal{X}}$ and probability measures μ, ν on \mathcal{X} , one has

$$|\mu f - \nu f| \leq \mathcal{W}(\mu, \nu) \|f\|_{\text{LIP}}, \quad (5.12)$$

and $\mathcal{W}(\mu, \nu)$ is the optimal constant for which this inequality holds. This theorem is known under the name *Kantorovich duality*.

Theorem 5.3 (Wasserstein distance and Lipshitz functions). *For probability measures μ, ν on \mathcal{X} , one has*

$$\mathcal{W}(\mu, \nu) = \sup_{\|f\|_{\text{LIP}} \leq 1} |\mu f - \nu f|. \quad (5.13)$$

For better overview, we will split the proof of Theorem 5.3 into several steps. The proof of the inequality \geq in (5.13) is easy.

Lemma 5.4 (One-sided bound). *For probability measures μ, ν on \mathcal{X} , one has*

$$\mathcal{W}(\mu, \nu) \geq \sup_{\|f\|_{\text{LIP}} \leq 1} |\mu f - \nu f|. \quad (5.14)$$

Proof Let γ be an optimal coupling for μ and ν and let (X, Y) have law γ . Then for any $f \in \mathbb{R}^{\mathcal{X}}$,

$$\begin{aligned} \left| \mathbb{E}[f(X)] - \mathbb{E}[f(Y)] \right| &= \left| \mathbb{E}[f(X) - f(Y)] \right| \leq \mathbb{E}[|f(X) - f(Y)|] \\ &\leq \|f\|_{\text{LIP}} \mathbb{E}[\text{dist}(X, Y)] = \mathcal{W}(\mu, \nu) \|f\|_{\text{LIP}}. \end{aligned} \quad (5.15)$$

This proves (5.12) and the inequality \geq in (5.13). ■

The inequality \leq in (5.13) is a bit deeper. We will present a proof based on the strong duality theorem of linear programming (Theorem 5.7 below). In order for that theorem to be applicable, we first need to reformulate the definition of the Wasserstein distance in a suitable form.

Lemma 5.5 (Alternative definition). *For probability measures μ, ν on \mathcal{X} , the Wasserstein distance $\mathcal{W}(\mu, \nu)$ is the minimum of the function*

$$\mathbb{R}^{\mathcal{X}^2} \ni \gamma \mapsto \sum_{x, y} \gamma(x, y) \text{dist}(x, y) \quad (5.16)$$

subject to the constraints

$$\gamma(x, y) \geq 0 \quad \forall x, y, \quad \sum_y \gamma(x, y) \leq \mu(x) \quad \forall x, \quad \sum_x \gamma(x, y) \leq \nu(y) \quad \forall y. \quad (5.17)$$

Proof Any γ that satisfies the constraints (5.17) is a subprobability measure, i.e., $\sum_{x, y} \gamma(x, y) \leq 1$. Since the space of subprobability measures is compact and the function in (5.16) is continuous, there exists a γ in which the function in (5.16) assumes its minimum. To complete the proof, it suffices to show

that such a γ must necessarily be a probability measure with marginals μ and ν . Imagine that it is not. Set

$$\tilde{\mu}(x) := \mu(x) - \sum_y \gamma(x, y) \quad \text{and} \quad \tilde{\nu}(y) := \nu(y) - \sum_x \gamma(x, y). \quad (5.18)$$

By assumption, $p := \sum_x \tilde{\mu}(x) = 1 - \sum_{x,y} \gamma(x, y) = \sum_y \tilde{\nu}(y)$ satisfies $p > 0$ so we can define probability measures $\bar{\mu} := p^{-1}\tilde{\mu}$ and $\bar{\nu} := p^{-1}\tilde{\nu}$. We now see that

$$\gamma'(x, y) := \gamma(x, y) + p\bar{\mu}(x)\bar{\nu}(y) \quad (5.19)$$

satisfies (5.17) while the function in (5.16) has a strictly higher value in γ' than in γ , contradicting our assumption that the maximum is assumed in γ . \blacksquare

For our next lemma, we let

$$\Delta := \sup_{x,y \in \mathcal{X}} \text{dist}(x, y) \quad (5.20)$$

denote the diameter of the graph (\mathcal{X}, E) .

Lemma 5.6 (Dual formulation). *For probability measures μ, ν on \mathcal{X} , the Wasserstein distance $\mathcal{W}(\mu, \nu)$ is the maximum of the function*

$$\mathbb{R}^{\mathcal{X}} \times \mathbb{R}^{\mathcal{X}} \ni (f, g) \mapsto \sum_x \nu(x)f(x) - \sum_x \mu(x)g(x) \quad (5.21)$$

subject to the constraints

$$g(x) \geq 0 \quad \forall x, \quad f(x) \leq \Delta \quad \forall x, \quad f(y) \leq g(x) + \text{dist}(x, y) \quad \forall x, y. \quad (5.22)$$

Proof By Lemma 5.5, $\Delta - \mathcal{W}(\mu, \nu)$ is the maximum of the function

$$\gamma \mapsto \sum_{x,y} \gamma(x, y) [\Delta - \text{dist}(x, y)] \quad (5.23)$$

subject to the constraints (5.17). In order to apply the strong duality theorem of linear programming (Theorem 5.7 below), we rewrite (5.23) and (5.17) in a more abstract form. For each $x, y, z \in \mathcal{X}$ and $i = 1, 2$, we define

$$\rho(z, i) := \begin{cases} \mu(z) & \text{if } i = 1, \\ \nu(z) & \text{if } i = 2. \end{cases} \quad \text{and} \quad A(z, i; x, y) := \begin{cases} 1_{\{z=x\}} & \text{if } i = 1, \\ 1_{\{z=y\}} & \text{if } i = 2. \end{cases} \quad (5.24)$$

Then we can rewrite the constraints (5.17) as

$$\gamma \geq 0, \quad \sum_{x,y} A(z, i; x, y) \gamma(x, y) \leq \rho(z, i). \quad (5.25)$$

The strong duality theorem of linear programming (Theorem 5.7 below) now tells us that $\Delta - \mathcal{W}(\mu, \nu)$ is the minimum of the function

$$g \mapsto \sum_{z,i} g(z, i) \rho(z, i) \quad (5.26)$$

subject to the constraints

$$g \geq 0, \quad \sum_{z,i} A(z, i; x, y) g(z, i) \geq \Delta - \text{dist}(x, y) \quad \forall x, y. \quad (5.27)$$

Let us write $g_i(z) := g(z, i)$ and define $f(z) := \Delta - g_2(z)$. Then $\Delta - \mathcal{W}(\mu, \nu)$ is the minimum of the expression

$$\sum_x \mu(x) g_1(x) + \Delta - \sum_x \nu(x) f(x) \quad (5.28)$$

subject to the constraints

$$g_1 \geq 0, \quad f \leq \Delta, \quad f(y) - g_1(x) \leq \text{dist}(x, y) \quad \forall x, y. \quad (5.29)$$

Therefore, $\mathcal{W}(\mu, \nu)$ is the maximum of the expression (5.21) subject to the constraints (5.22). \blacksquare

Proof of Theorem 5.3 We first prove the statement under the additional assumptions that $\mu(x) > 0$ and $\nu(x) > 0$ for all $x \in \mathcal{X}$. Lemma 5.6 says that the Wasserstein distance $\mathcal{W}(\mu, \nu)$ is the maximum of the function in (5.21) subject to the constraints (5.22). Let (f, g) be a pair of functions for which the function in (5.21) assumes its maximum. (The formulation of Lemma 5.6 is meant to include the fact that the maximum is assumed. This follows from Theorem 5.7, on which Lemma 5.6 is based.) We claim that $\|g\|_{\text{LIP}} \leq 1$. To see this, define

$$\tilde{g}(x) := \inf_{y \in \mathcal{Y}} [g(y) + \text{dist}(x, y)]. \quad (5.30)$$

Setting $y = x$ in the right-hand side we see that $\tilde{g}(x) \leq g(x)$ for all x . Since $g(y) \geq 0$ and $\text{dist}(x, y) \geq 0$ for all y we also have $\tilde{g} \geq 0$. The triangle inequality and (5.22) imply that

$$\begin{aligned} \tilde{g}(x) + \text{dist}(x, y) &= \inf_z [g(z) + \text{dist}(x, z)] + \text{dist}(x, y) \\ &= \inf_z [g(z) + \text{dist}(x, z) + \text{dist}(x, y)] \\ &\geq \inf_z [g(z) + \text{dist}(y, z)] \geq f(y) \quad \forall x, y. \end{aligned} \quad (5.31)$$

This shows that (f, \tilde{g}) satisfies the constraints (5.22) while $\tilde{g} \leq g$. If $\tilde{g}(x) < g(x)$ for some $x \in \mathcal{X}$, then by our assumption that $\mu(x) > 0$ for all $x \in \mathcal{X}$, the value of the function in (5.21) in (f, \tilde{g}) is strictly larger than in (f, g) , which contradicts the assumption that the maximum is assumed in (f, g) . We therefore conclude that $\tilde{g} = g$. By (5.30), this means that

$$g(x) \leq g(y) + \text{dist}(x, y) \quad \forall x, y. \quad (5.32)$$

Thus,

$$g(x) - g(y) \leq \text{dist}(x, y) \quad \text{and} \quad g(y) - g(x) \leq \text{dist}(y, x), \quad (5.33)$$

which means that $|g(x) - g(y)| \leq \text{dist}(x, y)$ and hence $\|g\|_{\text{LIP}} \leq 1$.

By our arguments so far, there exists a pair of functions (f, g) for which the function in (5.21) assumes its maximum, and each such pair must satisfy $\|g\|_{\text{LIP}} \leq 1$. We claim that there exists a pair (f, g) for which the maximum is assumed such that moreover $\inf_x g(x) = 0$. Indeed, if this is not the case, then we can make it true by subtracting the same constant from both f and g , which does not change the value of the function in (5.21) and also does not change the fact that the constraints (5.22) are satisfied.

Let (f, g) be such a pair in which the maximum is achieved and for which $\inf_x g(x) = 0$ and $\|g\|_{\text{LIP}} \leq 1$. We claim that $f = g$. Setting $x = y$ in the constraint that $f(y) \leq g(x) + \text{dist}(x, y)$ for all x, y yields $f \leq g$. If $f(x) < g(x)$ for some $x \in \mathcal{X}$, then by our assumption that $\nu(x) > 0$ for all $x \in \mathcal{X}$, the value of the function in (5.21) in (g, g) is strictly larger than in (f, g) . Therefore, by maximality, we can conclude that $f = g$ provided we show that the pair (\tilde{f}, g) with $\tilde{f} := g$ satisfies the constraints (5.22). Indeed $\tilde{f}(y) \leq g(x) + \text{dist}(x, y)$ since $\tilde{f} = g$ and $\|g\|_{\text{LIP}} \leq 1$. Moreover, recalling the definition of Δ , we see that $\inf_x g(x) = 0$ and $\|g\|_{\text{LIP}} \leq 1$ imply $g \leq \Delta$, so $\tilde{f} \leq \Delta$ is also satisfied by $\tilde{f} = g$.

We have now proved that there exists a function f with $\|f\|_{\text{LIP}} \leq 1$ such that $\mathcal{W}(\mu, \nu) = \nu f - \mu f$. In particular, this proves that

$$\mathcal{W}(\mu, \nu) \leq \sup_{\|f\|_{\text{LIP}} \leq 1} |\mu f - \nu f|. \quad (5.34)$$

Together with Lemma 5.4, this proves Theorem 5.3 under the additional assumptions that $\mu(x) > 0$ and $\nu(x) > 0$ for all $x \in \mathcal{X}$. To remove these assumptions, it suffices to show that both the left- and right-hand sides of (5.13) are continuous as a function of (μ, ν) .

By Lemma 5.1, the Wasserstein distance is a metric on the space of probability laws on \mathcal{X} . A metric is always continuous with respect to the topology that it generates, which here is just the topology in which $\mu_n \rightarrow \mu$

if $\mu_n(x) \rightarrow \mu(x)$ for all $x \in \mathcal{X}$. This shows that $\mathcal{W}(\mu, \nu)$ is continuous in (μ, ν) . It is easy to see that the right-hand sides of (5.13) also defines a metric and hence is also continuous. \blacksquare

Below is the strong duality theorem of linear programming.

Theorem 5.7 (Strong duality). *Let $A(i, j)_{1 \leq i \leq n, 1 \leq j \leq m}$ be a real matrix, and let $(b(1), \dots, b(m))$ and $(c(1), \dots, c(m))$ be real vectors. Assume that the function $x \mapsto \sum_{j=1}^m c(j)x(j)$ assumes its maximum M_+ over the set*

$$\{x \in \mathbb{R}^m : x(j) \geq 0 \ \forall 1 \leq j \leq m, \sum_{j=1}^m A(i, j)x(j) \leq b(i) \ \forall 1 \leq i \leq n\}. \quad (5.35)$$

Then the function $y \mapsto \sum_{i=1}^n b(i)y(i)$ assumes its minimum M_- over the set

$$\{y \in \mathbb{R}^n : y(i) \geq 0 \ \forall 1 \leq i \leq n, \sum_{i=1}^n A(i, j)y(i) \geq c(j) \ \forall 1 \leq j \leq m\}, \quad (5.36)$$

and one has $M_- = M_+$.

6 Curvature

We continue to assume that P is irreducible and reversible. The (Ollivier-Ricci) *curvature* of P is the quantity $\kappa(P) \leq 1$ defined as

$$\kappa(P) := 1 - \sup_{\{x, y\} \in E} \mathcal{W}(P(x, \cdot), P(y, \cdot)), \quad (6.1)$$

where E is the set of edges defined in (5.1). A Markov chain is said to have *positive curvature* if $\kappa(P) > 0$. In view of Theorem 5.3 and (5.10),

$$\begin{aligned} 1 - \kappa(P) &= \sup_{\{x, y\} \in E} \sup_{\|f\|_{\text{LIP}} \leq 1} |Pf(x) - Pf(y)| \\ &= \sup_{\|f\|_{\text{LIP}} \leq 1} \|Pf\|_{\text{LIP}}. \end{aligned} \quad (6.2)$$

For positively curved Markov chains, in the light of (2.6) and (5.4), the following lemma gives a bound on the total variation distance to equilibrium.

Lemma 6.1 (Curvature bound). *For any probability measures μ, ν on \mathcal{X} , one has*

$$\begin{aligned} \text{(i)} \quad & \mathcal{W}(\mu P^n, \nu P^n) \leq (1 - \kappa(P))^n \mathcal{W}(\mu, \nu) \quad (n \geq 0), \\ \text{(ii)} \quad & \mathcal{W}(\mu P_t, \nu P_t) \leq e^{-\kappa(P)t} \mathcal{W}(\mu, \nu) \quad (n \geq 0) \end{aligned} \quad (6.3)$$

Proof In the discrete time setting this is proved in [Sal22, Lemma 20]. Note, however, that the notation there is a bit different. What we call $1 - \kappa(P)$ is called $e^{-\kappa(P)}$ there. Our definition is more suitable for the continuous-time setting and coincides with the definition in [Sal21].

To prove (6.3) (i), we use (6.2) and induction to obtain

$$\|P^n f\|_{\text{LIP}} \leq (1 - \kappa(P))^n \|f\|_{\text{LIP}} \quad (n \geq 0). \quad (6.4)$$

By Theorem 5.3 and (5.12), it follows that

$$\begin{aligned} \mathcal{W}(\mu P^n, \nu P^n) &= \sup_{\|f\|_{\text{LIP}} \leq 1} |\mu P^n f - \nu P^n f| \\ &\leq \sup_{\|f\|_{\text{LIP}} \leq 1} \mathcal{W}(\mu, \nu) \|P^n f\|_{\text{LIP}} \leq (1 - \kappa(P))^n \mathcal{W}(\mu, \nu), \end{aligned} \quad (6.5)$$

proving (6.3) (i). To prove also (6.3) (ii), we use (1.6), which says that

$$P_t = \sum_{n=0}^{\infty} p_t(n) P^n \quad \text{with} \quad p_t(n) := e^{-t} \frac{t^n}{n!} \quad (n \geq 0). \quad (6.6)$$

The Lipschitz “norm” is not really a norm but only a pseudonorm. Using the triangle inequality for this pseudonorm, as well as the continuity of the map $f \mapsto \|f\|_{\text{LIP}}$, we obtain

$$\begin{aligned} \|P_t f\|_{\text{LIP}} &= \left\| \sum_{n=0}^{\infty} p_t(n) P^n \right\|_{\text{LIP}} \leq \sum_{n=0}^{\infty} p_t(n) \|P^n\|_{\text{LIP}} \\ &\leq \sum_{n=0}^{\infty} p_t(n) (1 - \kappa(P))^n \|f\|_{\text{LIP}} = e^{-\kappa(P)t} \|f\|_{\text{LIP}}. \end{aligned} \quad (6.7)$$

The final equality here is not completely obvious. Recall that $\kappa(P) \leq 1$ by definition. If $\kappa(P) \geq 0$, then we can interpret $\kappa = \kappa(P)$ as a probability. Now $\sum_{n=0}^{\infty} p_t(n) (1 - \kappa)^n$ is the probability that if we have a random number of particles with a Poisson distribution with mean t , and we perform a random thinning of these particles, where each particle has an independent probability κ that we keep it, then no particles survive the thinning. Now it is well-known that after thinning, the number of particles that is left is Poisson distributed with mean κt . Thus $p_{\kappa t}(0) = e^{-\kappa t}$ is the probability that no particles survive the thinning. For general κ , one can verify the final equality in (6.7) by direct computation. We leave this as an exercise to the reader. Formula (6.3) (ii) follows from (6.7) in the same way we derived (6.3) (i) from (6.4). \blacksquare

7 Cut-off

For any probability kernel P , we define the *sparsity parameter* $\Delta(P)$ by

$$\Delta(P) := \max_{\{x,y\} \in E} \frac{1}{P(x,y)} \quad (7.1)$$

For positive functions f_n, g_n , we write $f_n \ll g_n$ as $n \rightarrow \infty$ if $f_n/g_n \rightarrow 0$. Recall the definitions of the mixing and relaxation times (in the continuous-time setting) in (2.7) and (2.15). We set $\hat{t}_{\text{MIX}}(P_n) := \hat{t}_{\text{MIX}}^{(1/4)}(P_n)$. Below is a simplified version of the main result of [Sal21].

Theorem 7.1 (Conditions for cut-off). *Let P_n be a sequence of irreducible reversible probability kernels on finite sets \mathcal{X}_n . Assume that $|\mathcal{X}_n| \geq 3$ and $\kappa(P_n) \geq 0$ for all n . Assume that for each $\varepsilon \in (0, 1)$,*

$$\left(\hat{t}_{\text{REL}}(P_n) \log \Delta(P_n)\right)^2 \ll \hat{t}_{\text{MIX}}^{(\varepsilon)}(P_n) \quad \text{as } n \rightarrow \infty. \quad (7.2)$$

Then for each $0 < \varepsilon < \frac{1}{2}$, there exists a constant C_ε such that

$$\hat{t}_{\text{MIX}}^{(\varepsilon)}(P_n) - \hat{t}_{\text{MIX}}^{(1-\varepsilon)}(P_n) \leq C_\varepsilon \sqrt{\hat{t}_{\text{MIX}}(P_n) \hat{t}_{\text{REL}}(P_n) \log \Delta(P_n)} \quad (7.3)$$

for all n large enough.

Remark 1 Assumption (7.2) implies that $\hat{t}_{\text{REL}}(P_n) \log \Delta(P_n) \ll \sqrt{\hat{t}_{\text{MIX}}(P_n)}$ so that the right-hand side of (7.3) satisfies

$$\sqrt{\hat{t}_{\text{MIX}}(P_n)} \hat{t}_{\text{REL}}(P_n) \log \Delta(P_n) \ll \hat{t}_{\text{MIX}}(P_n). \quad (7.4)$$

This shows that the sequence of Markov chains with transition kernels P_n exhibits *cut-off*: the total variation distance $\hat{\mathcal{D}}_{P_n}(t)$ changes from being close to one to close to zero in a time interval centered around $\hat{t}_{\text{MIX}}(P_n)$ whose duration is in the limit much shorter than $\hat{t}_{\text{MIX}}(P_n)$.

Remark 2 In these notes, we use a hat to indicate that a certain quantity belongs to the continuous-time setting. This is why our notation is a bit different from the notation in [Sal21].

Remark 3 The main result of [Sal21] is considerably more general than Theorem 7.1. The condition $\kappa(P_n) \geq 0$ says that P_n has non-negative Ollivier-Ricci curvature. In [Sal21], it is shown that the result remains true if Ollivier-Ricci curvature is replaced by Bakry-Émery curvature, a concept that we have not treated here. The main result of [Sal21] moreover applies

to a large class of non-reversible chains as well. In that case, however, the relaxation time from Theorem 7.1 has to be replaced by the relaxation time of the symmetrised kernel $(P + P^*)/2$ and one needs the additional assumption that $P(x, y) > 0$ if and only if $P(y, x) > 0$.

Remark 4 The assumption (7.2) is inspired by the so-called *product condition* $\hat{t}_{\text{REL}}(P_n) \ll \hat{t}_{\text{MIX}}^{(\varepsilon)}(P_n)$ ($\varepsilon \in (0, 1)$), that has in the past been conjectured to imply cut-off. This conjecture is known to be wrong, however. Typical counterexamples, however, are extremely non-sparse, in the sense that the Markov chain can jump from any point in the state space to any other point with positive probability. For these counterexamples, $\Delta(P_n)$ would grow very fast with n and (7.2) would not be satisfied. This justifies the occurrence of $\Delta(P_n)$ in (7.2).

Remark 5 The occurrence of the square in (7.2) is a less pleasant feature of the theorem. It seems natural that in order to prove a quantitative statement like (7.3) on the size of the “critical window” where $\hat{\mathcal{D}}_{P_n}(t)$ changes from being close to one to being close to zero, one needs a quantitative version of the product condition, i.e., one needs to say *how much larger* $\hat{t}_{\text{MIX}}^{(\varepsilon)}(P_n)$ is in comparison with $\hat{t}_{\text{REL}}(P_n)$. However, (7.2) is quite a strong assumption in the sense that even when $\Delta(P_n)$ is of order one, it requires $\hat{t}_{\text{MIX}}^{(\varepsilon)}(P_n)$ to be much larger than the *square* of $\hat{t}_{\text{REL}}(P_n)$ (for example, $\hat{t}_{\text{MIX}}^{(\varepsilon)}(P_n)$ could be of order $(\hat{t}_{\text{REL}}(P_n))^{2+\varepsilon}$ for any $\varepsilon > 0$). In [Sal22, Section 5] we have seen that the random walk on the hypercube has cut-off with $t_{\text{REL}}(P_n) \sim n$ and $t_{\text{MIX}}(P_n) \sim \frac{1}{2}n \log n$. In this example, the mixing time grows only a little bit faster than the relaxation time (the difference is the logarithmic term) so this example cannot be covered by Theorem 7.1. In future, one would hope to relax condition (7.2) in this respect..

Remark 6 In order to check assumption (7.2) of Theorem 7.1, one only needs upper bounds on $\hat{t}_{\text{REL}}(P_n)$ and lower bounds on $\hat{t}_{\text{MIX}}^{(\varepsilon)}(P_n)$. In particular, one does *not* need to determine the precise asymptotics of $\hat{t}_{\text{MIX}}^{(\varepsilon)}(P_n)$, including the preconstant, which is often very difficult. Theorem 7.1 is one of very few results of this sort and the first one that applies to a very general class of probability kernels, instead of only to very special classes such as birth-and-death chains etc. In this sense, Theorem 7.1 is a breakthrough result.

8 Strategy of the proof

The *relative entropy* (or *Kullback-Leibler divergence*) or a probability measure μ on a finite set \mathcal{X} with respect to another probability measure π is defined

as

$$d_{\text{KL}}(\mu\|\pi) := \sum_{x \in \mathcal{X}} \mu(x) \log \frac{\mu(x)}{\pi(x)}. \quad (8.1)$$

Here, for simplicity, we assume that $\pi(x) > 0$ everywhere. A related quantity, relatively unknown, the *varentropy*, is defined as

$$\mathcal{V}_{\text{KL}}(\mu\|\pi) := \sum_{x \in \mathcal{X}} \mu(x) \left(\log \frac{\mu(x)}{\pi(x)} - d_{\text{KL}}(\mu\|\pi) \right)^2. \quad (8.2)$$

Note that

$$\mathcal{V}_{\text{KL}}(\mu\|\pi) = \text{Var}_{\mu}(f) \quad \text{with} \quad f(x) := \log \frac{\mu(x)}{\pi(x)}, \quad (8.3)$$

since clearly $\mu f = d_{\text{KL}}(\mu\|\pi)$. Following [Sal21], we define

$$\mathcal{V}_{\text{KL}}^*(t) := \sup_{x \in \mathcal{X}} \mathcal{V}_{\text{KL}}(P_t(x, \cdot)\|\pi). \quad (8.4)$$

The following theorem, which is [Sal21, Thm 5], is key to the proof of Theorem 7.1.

Theorem 8.1 (Entropic concentration). *Let P be an irreducible probability kernel on a finite set \mathcal{X} and let $0 < \varepsilon < \frac{1}{2}$. Then*

$$\hat{t}_{\text{MIX}}^{(\varepsilon)}(P) - \hat{t}_{\text{MIX}}^{(1-\varepsilon)}(P) \leq 2\varepsilon^{-2} \hat{t}_{\text{REL}}(P) \left[1 + \sqrt{\mathcal{V}_{\text{KL}}^*(\hat{t}_{\text{MIX}}^{(1-\varepsilon)}(P))} \right]. \quad (8.5)$$

The second ingredient of the proof of Theorem 7.1 is [Sal21, Thm 5], which we cite here in simplified form.

Theorem 8.2 (Varentropy estimate). *Let P_n be a sequence of irreducible reversible probability kernels with $\kappa(P_n) \geq 0$. Fix $0 < \varepsilon < 1$ and assume that*

$$\sqrt{\hat{t}_{\text{REL}}(P_n)} \ll \hat{t}_{\text{MIX}}^{(\varepsilon)}(P_n) \quad \text{as } n \rightarrow \infty. \quad (8.6)$$

Then there exists a constant $C < \infty$ such that

$$\mathcal{V}_{\text{KL}}^*(\hat{t}_{\text{MIX}}^{(\varepsilon)}(P_n)) \leq C \hat{t}_{\text{MIX}}^{(\varepsilon)}(P_n) (\log \Delta(P_n))^2 \quad (8.7)$$

for all n large enough.

We now show how Theorems 8.1 and 8.2 imply Theorem 7.1.

Proof of Theorem 7.1 We first make some trivial observations. Since $|\mathcal{X}_n| \geq 3$ and P_n is irreducible and reversible, there must be at least one

$x \in \mathcal{X}$ that has degree two in the graph (\mathcal{X}_n, E_n) , where $E_n = \{\{x, y\} : P_n(x, y) > 0\}$. It follows that $P(x, y) \leq \frac{1}{2}$ for at least one $\{x, y\} \in E_n$ and hence $\Delta(P_n) \geq 2$ for all n . Since P_n is a probability measure, $|\lambda| \leq 1$ for all eigenvalues λ of P_n . Since P_n is irreducible, the multiplicity of the eigenvalue one is one, so $-1 \leq \Re(\lambda_2(P_n)) < 1$, which by (2.15) implies $\frac{1}{2} \leq \hat{t}_{\text{REL}}(P_n) < \infty$. Thus

$$\Delta(P_n) \geq 2 \quad \text{and} \quad \hat{t}_{\text{REL}}(P_n) \geq \frac{1}{2} \quad \forall n. \quad (8.8)$$

In view of this, (7.2) implies that

$$\hat{t}_{\text{MIX}}^{(\varepsilon)}(P_n) \rightarrow \infty \quad \text{as } n \rightarrow \infty \quad \forall 0 < \varepsilon < 1. \quad (8.9)$$

Fix $0 < \varepsilon < \frac{1}{2}$. Since $\Delta(P_n) \geq 2$, (7.2) implies that

$$\begin{aligned} (\hat{t}_{\text{REL}}(P_n))^2 &\ll \hat{t}_{\text{MIX}}^{(\varepsilon)}(P_n) \\ \Rightarrow \quad \sqrt{\hat{t}_{\text{REL}}(P_n)} &\ll \hat{t}_{\text{MIX}}^{(\varepsilon)}(P_n)^{1/4} \ll \hat{t}_{\text{MIX}}^{(\varepsilon)}(P_n), \end{aligned} \quad (8.10)$$

where in the last step we have used (8.9). This shows that Theorem 8.2 is applicable. Using moreover Theorem 8.1, inserting (8.7) into (8.5), we obtain

$$\begin{aligned} \hat{t}_{\text{MIX}}^{(\varepsilon)}(P_n) - \hat{t}_{\text{MIX}}^{(1-\varepsilon)}(P_n) &\leq 2\varepsilon^{-2} \hat{t}_{\text{REL}}(P_n) \left[1 + \sqrt{C \hat{t}_{\text{MIX}}^{(1-\varepsilon)}(P_n) (\log \Delta(P_n))^2} \right] \\ &\leq 2\varepsilon^{-2} \hat{t}_{\text{REL}}(P_n) \left[1 + \sqrt{C} \sqrt{\hat{t}_{\text{MIX}}^{(1-\varepsilon)}(P_n) \log \Delta(P_n)} \right]. \end{aligned} \quad (8.11)$$

In view of (8.8) and (8.9), the expression under the square root tends to infinity, so for sufficiently large n we can forget about the term that is 1, at the cost of making the constant \sqrt{C} a bit larger. Thus, we see that for some $C_\varepsilon < \infty$

$$\hat{t}_{\text{MIX}}^{(\varepsilon)}(P_n) - \hat{t}_{\text{MIX}}^{(1-\varepsilon)}(P_n) \leq C_\varepsilon \sqrt{\hat{t}_{\text{MIX}}^{(1-\varepsilon)}(P_n)} \hat{t}_{\text{REL}}(P_n) \log \Delta(P_n). \quad (8.12)$$

This is almost the same as (7.3), except that in the right-hand side we need to replace $\hat{t}_{\text{MIX}}^{(1-\varepsilon)}(P_n)$ by $\hat{t}_{\text{MIX}}(P_n)$, which is defined as $\hat{t}_{\text{MIX}}^{(1/4)}(P_n)$. Assumption (7.2) implies that $\hat{t}_{\text{REL}}(P_n) \log \Delta(P_n) \ll \sqrt{\hat{t}_{\text{MIX}}^{(1-\varepsilon)}(P_n)}$. Inserting this into (8.12) yields

$$\hat{t}_{\text{MIX}}^{(\varepsilon)}(P_n) - \hat{t}_{\text{MIX}}^{(1-\varepsilon)}(P_n) \ll \hat{t}_{\text{MIX}}^{(1-\varepsilon)}(P_n). \quad (8.13)$$

This shows that the chain exhibits cut-off, and hence, for each $0 < \varepsilon < \frac{1}{2}$,

$$\frac{\hat{t}_{\text{MIX}}^{(\varepsilon)}(P_n)}{\hat{t}_{\text{MIX}}^{(1-\varepsilon)}(P_n)} \xrightarrow{n \rightarrow \infty} 1 \quad (8.14)$$

As a result, in the right-hand side of (8.12), it asymptotically does not matter if we write $\hat{t}_{\text{MIX}}^{(1-\varepsilon)}(P_n)$ or $\hat{t}_{\text{MIX}}^{(1/4)}(P_n)$. We can replace one by the other and the inequality will remain true for large n , at the cost of perhaps having to change the constant C_ε a bit. ■

References

- [Sal21] Justin Salez. Cutoff for non-negatively curved Markov chains. Preprint, 2021, arXiv:2102.05597v2.
- [Sal22] Justin Salez. *Mixing Times of Markov Chains* Lecture notes, 2022.