Mean field conditions for coalescing random walks

Roberto Imbuzeiro Oliveira*

September 27, 2011

Abstract

The main results in this paper are about the *full coalescence time* C of a system of coalescing random walks over a finite graph G. Letting $\mathsf{m}(G)$ denote the mean meeting time of two such walkers, we give sufficient conditions under which $\mathbb{E}[\mathsf{C}] \approx 2\mathsf{m}(G)$ and $\mathsf{C}/\mathsf{m}(G)$ has approximatelly the same law as in the "mean field" setting of a large complete graph. One of our theorems is that mean field behavior occurs over all vertex-transitive graphs whose mixing times are much smaller than $\mathsf{m}(G)$; this nearly solves an open problem of Aldous and Fill and also generalizes results of Cox for discrete tori in $d \geq 2$ dimensions. Other results apply to non-reversible walks and also generalize previous theorems of Durrett and Cooper et al. Slight extensions of these results apply to voter model consensus times, which are related to coalescing random walks via duality.

Our main proof ideas are a strenghtening of the usual approximation of hitting times by exponential random variables, which give results for non-stationary initial states; and a new general set of conditions under which we can prove that the hitting time of a union of sets behaves like a minimum of independent exponentials. In particular, this will, show that the first meeting time among k random walkers has mean $\approx \mathsf{m}(G)/\binom{k}{2}$.

1 Introduction

Start a continuous-time random walk from each vertex of a finite, connected graph G. The walkers evolve independently, except that when two walkers *meet* – ie. lie on the same vertex at the same time –, they coalesce into one. One may easily show that there will almost surely be a finite time at which only one walk will remain in this system. The first such time is called the *full coalescence time* for G and is denoted by C.

The main goal of this paper is to show that one can estimate the law of C for a large family of graphs G, and that this law only depends on G through a single rescaling parameter. More precisely, we will prove results of the following form: if the *mixing time* $\mathrm{T}^G_{\mathrm{mix}}$ of G (defined in Section 2) is "small", then there exists a parameter $\mathsf{m}(G) > 0$ such that the law $\mathsf{C}/\mathsf{m}(G)$

*IMPA, Rio de Janeiro, RJ, Brazil, 22430-040.Supported by a *Projeto Universal* and a *Bolsa de Produtuvidade em Pesquisa* from CNPq, Brazil.

takes a universal shape. Slight extensions of these results will be used to study the so-called *voter model consensus time* on G.

The universal shape of $\mathsf{C}/\mathsf{m}(G)$ comes from a *mean field* computation over a large complete graph K_n. In this case the distribution of C can be computed exactly (cf. [4, Chapter 14]):

$$\frac{\mathsf{C}}{(n-1)/2} =_d \sum_{i=2}^{n} \mathsf{Z}_i,$$

where:

$$\mathsf{Z}_2, \mathsf{Z}_3, \mathsf{Z}_4 \ldots \text{are independent and } \forall i \geq 2,\, t \geq 0 \,:\, \mathbb{P}\left(\mathsf{Z}_i \geq t\right) = e^{-t\binom{i}{2}}. \tag{1}$$

In words, C is a rescaled sum of independent exponential random variables with means $1/\binom{i}{2}$, $2 \leq i \leq n$.

The scaling factor $(n-1)/2$ is the expected meeting time of two independent random walks over K_n, and we see that

$$\frac{\mathsf{C}}{(n-1)/2} \to_w \sum_{i\geq 2} \mathsf{Z}_i \text{ and } \frac{\mathbb{E}\left[\mathsf{C}\right]}{(n-1)/2} \to 2 \text{ when } n \text{ grows.}$$

Thus the general problem we address in this paper is:

> **General problem:** Given a graph G, let $\mathsf{m}(G)$ denote the expected meeting time of two independent random walks over G, both started from stationarity. Give sufficient conditons on G under which C has mean-field behavior, that is:
>
> $$\mathrm{Law}\left(\mathsf{C}/\mathsf{m}(G)\right) \approx \mathrm{Law}\left(\sum_{i\geq 2} \mathsf{Z}_i\right), \tag{2}$$
>
> and
>
> $$\mathbb{E}\left[\mathsf{C}\right] \approx \mathsf{m}(G)\, \mathbb{E}\left[\sum_{i\geq 2} \mathsf{Z}_i\right] = 2\, \mathsf{m}(G). \tag{3}$$

A version of this problem was posed in Aldous and Fill's 1994 draft [4, Chapter 14] and much more recently by Aldous [2]. However, as far as we know there are only two families of examples the problem has been fully solved. Discrete tori $G = (\mathbb{Z}/m\mathbb{Z})^d$ with with $d \geq 2$ fixed and $m \gg 1$ were considered in Cox's 1989 paper [6]. More recently, Cooper, Frieze and Radzik [5] proved mean field behavior in large random d-regular graphs (d bounded). Partial results were also obtained by Durrett [7, 8] for certain models of large networks.

We note that mean-field behaviour is not universal over all large graphs: counterexamples include one-dimensional tori [6] and large stars [1]

[1] In this case C is lower bounded by the time the last edge of the star is crossed by some walker, which is about $\log n$. This is much larger than $\mathsf{m}(G)$, which is bounded.

1.1 Main results

Our results in this paper address (2) and (3) simultaneously by proing approximation bounds in L_1 Wasserstein distance, which implies closeness of first moments (cf. Section 2.2). The first theorem implies that mean field behaviour occurs whenever G is vertex-transitive and its *mixing time* (defined in Section 2) is much smaller than $\mathsf{m}(G)$. This generalizes Cox's result for discrete tori and nearly solves an Open problem 12 in [4, Chapter 14], where it was conjectured that it would suffice for the *relaxation time* to be small relative to $\mathsf{m}(Q)$ [2].

The natural setting for this first theorem is that of walkers evolving according to the same reversible, transitive Markov chain (the definition of C easily generalizes to this case), where *transitive* means that for any two states x and y one can find a permutation of the state space mapping x to y and leaving the transition rates invariant. (Clearly, random walk on a vertex-transitive graph is transitive in this sense.)

Notational convention 1 *In this paper we will use "$b = O(a)$" in the following sense: there exist universal constants $C, \xi > 0$ such that $|a| \leq \xi \Rightarrow |b| \leq C|a|$.*

Theorem 1.1 (Mean field for transitive, reversible chains) *Let Q be the (generator of a) transitive, reversible, irreducible Markov chain over a finite state space $\mathbf{V}$, with mixing time $\mathrm{T}^Q_{\mathrm{mix}}$. Define $\mathsf{m}(Q)$ to be the expected meeting time of two independent continuous-time random walks over $\mathbf{V}$ that evolve according to Q, both started from the uniform (stationary) distribution. Denote by C be the full coalescence time , ie. the first time in the coalescing random walks process defined in termos of Q at which only one walker remains. Finally, let $\{\mathsf{Z}_i\}_{i=2}^{+\infty}$ be in (1). Then:*

$$d_W\left(\mathrm{Law}\left(\frac{\mathsf{C}}{\mathsf{m}(Q)}\right), \mathrm{Law}\left(\sum_{i\geq 2}\mathsf{Z}_i\right)\right) = O\left(\left[\rho(Q)\ln\left(\frac{1}{\rho(Q)}\right)\right]^{1/6}\right),$$

where:

$$\rho(Q) \equiv \frac{\mathrm{T}^Q_{\mathrm{mix}}}{\mathsf{m}(Q)}$$

and d_W denotes L_1 Wasserstein distance. In particular,

$$\mathbb{E}\left[\mathsf{C}\right] = \left\{2 + O\left(\left[\rho(Q)\ln\left(\frac{1}{\rho(Q)}\right)\right]^{1/6}\right)\right\}\mathsf{m}(Q).$$

We also have results on coalescing random walks evolving according to arbitrary generators Q on finite state spaces $\mathbf{V}$. Again, we only require that the mixing time $\mathrm{T}^Q_{\mathrm{mix}}$ of Q be sufficiently small relative to other parameters of the chain. This is partially motivated by recent attempts by Durrett [7, 8] to understand coalescing random walks and voter models over models of large networks, such as CHKWS random graphs and random graphs with prescribed degrees.

[2]Mixing and relaxation times have the same order of magnitude over discrete tori. In general, the relaxation time is at most $O(\ln|\mathbf{V}|)$ times smaller than the mixing time, where $|\mathbf{V}|$ is the number of vertices.

Theorem 1.2 (Mean field for general Markov chains) *Let Q denote (the generator of) a mixing Markov chain over a finite set $\mathbf{V}$, with unique stationary distribution π. Denote by $q_{\max}$ the maximum transition rate from any $x \in \mathbf{V}$ and by $\pi_{\max}$ the maximum stationary probability of an element of $\mathbf{V}$. Let $\mathsf{m}(Q)$ denote the expected meeting time of two random walks evolving according to Q, both started from π. Finally, let C denote the full coalescence time of random walks evolving in $\mathbf{V}$ according to Q. Then:*

$$d_W\left(\mathrm{Law}\left(\frac{\mathsf{C}}{\mathsf{m}(Q)}\right), \mathrm{Law}\left(\sum_{i\geq 2} \mathsf{Z}_i\right)\right) = O\left(\left(\alpha(Q)\,\ln\left(\frac{1}{\alpha(Q)}\right)\ln^4|\mathbf{V}|\right)^{1/6}\right)$$

where

$$\alpha(Q) = (1 + q_{\max}\,\mathrm{T}^{Q}_{\mathrm{mix}})\,\pi_{\max}$$

and d_W again denotes L^1 Wasserstein distance. In particular,

$$\mathbb{E}\left[\mathsf{C}\right] = \left\{2 + O\left(\left[\alpha(Q)\ln\left(\frac{1}{\alpha(Q)}\right)\,\ln^4|\mathbf{V}|\right]^{1/6}\right)\right\}\mathsf{m}(Q).$$

We note that this theorem does *not* imply Theorem 1.1: for instance, it does *not* work for two-dimensional discrete tori. However, it applies for discrete tori in any larger dimension, to supercritical percolation clusters in such tori, to expanders, and to all other graphs with $\mathrm{T}^G_{\mathrm{mix}} \ll |\mathbf{V}|/\ln^4|\mathbf{V}|\ln\ln|\mathbf{V}|$, where the ratio of maximal to average degree is bounded. More generally, we have the following cololary:

Corollary 1.1 *Assume G is a connected graph with vertex set $\mathbf{V}$ and that $\epsilon \in (|\mathbf{V}|^{-1}, 1)$ is such that:*

$$\left(\frac{\max_{x\in\mathbf{V}}\deg_G(x)}{|\mathbf{V}|^{-1}\sum_{x\in\mathbf{V}}\deg_G(x)}\right)\mathrm{T}^G_{\mathrm{mix}} \leq \frac{\epsilon|\mathbf{V}|}{\ln^4|\mathbf{V}|\ln\ln|\mathbf{V}|}.$$

Then

$$d_W\left(\mathrm{Law}\left(\frac{\mathsf{C}}{\mathsf{m}(G)}\right), \mathrm{Law}\left(\sum_{i\geq 2}\mathsf{Z}_i\right)\right) = O\left(\left[\epsilon\left(1 + \frac{\ln(1/\epsilon)}{\ln\ln|\mathbf{V}|}\right)\right]^{1/6}\right).$$

We do not prove this corollary explicitly, but note that it follows directly from the well-known formula for the stationary distribution of random walks on graphs. This corollary brings us back to the aforementioned results of Durrett in [7, 8]. In all of the examples of networks he considered one has $\mathrm{T}^Q_{\mathrm{mix}}$ of order polylogarithmic in $|\mathbf{V}|$ and the maximum degree is less than $|\mathbf{V}|^{1-\eta}$ for some $0 \leq \eta < 1$. Our result proves mean field behavior in all those cases with error bounds of the form $|\mathbf{V}|^{-\frac{\eta}{6}+o(1)}$. This is a stronger result than what Durrett obtained since he could only control a bounded number of coalescing random walkers on a given graph. More subtly, our result uses very little specific information about the network itself. This is important because recent measurements of real-life social networks [9] suggest that known models of large networks are very innacurate with respect to most network characteristics outside of degree distributions and conductance. In fairness, we should note that coalescing random walks and voter models over large networks are not particularly realistic either, but at the very least we know that mean field behavior is not an artifact of a particular class of models. We also observe that our Theorem 1.2 also works for non-reversible chains, eg. random walks on directed graphs.

1.2 Results for the voter model

The voter model is a very well-known process in the Interacting Particle Systems literature [10]. The configuration space for the voter model is the power set $\mathcal{O}^{\mathbf{V}}$ of functions $\eta : \mathbf{V} \to \mathcal{O}$, where $\mathbf{V}$ is some non-empty set and $\mathcal{O}$ is a non-empty set of possible opinions. The evolution of the process is determined by numbers $q(x, y)$ $(x, y \in V, x \neq y)$ and is informally described as follows: at rate $q(x, y)$, node x copies y's opinion. That is, there is a transition at rate $q(x, y)$ from any state $\eta : \mathbf{V} \to \mathcal{O}$ to the corresponding state $\eta^{x \leftarrow y}$, where:

$$\eta^{x \leftarrow y}(z) = \begin{cases} \eta(y), & x = y; \\ \eta(z), & z \in V \backslash \{x\}. \end{cases} .$$

A classical duality result relates this voter model to a system of coalescing random walks with transition rates $q(\cdot, \cdot\cdot)$ and corresponding generator Q. More precisely, suppose that $\mathbf{V} = \{x(1), \ldots, x(n)\}$ and that $(\overline{X}_t(i))_{t \geq 0, 1 \leq i \leq n}$ is a system of coalescing random walks evolving according to Q with $X_0(i) = x(i)$ for each $1 \leq i \leq n$.

Duality: Choose $\eta_0 \in \mathcal{O}^{\mathbf{V}}$. Then the configuration

$$\hat{\eta}_t : x(i) \in \mathbf{V} \mapsto \eta_0(\overline{X}_t(i)) \in \mathcal{O} \ \ (1 \leq i \leq n)$$

has the same distribution as the state η_t of the voter model at time t, when the initial state is η_0. In particular, the *consensus time* for the voter model:

$$\tau \equiv \inf\{t \geq 0 \, : \, \forall v, w \in \mathbf{V}, \, \eta_t(i) = \eta_t(j)\}$$

satisfies $\mathbb{E}[\tau] \leq \mathbb{E}[\mathsf{C}] < +\infty$.

Now assume that the initial state $\eta_0 \in \mathcal{O}^{\mathbf{V}}$ is random and that the random variables $\{\eta_0(x)\}_{x \in V}$ are iid and have common law μ which is not a point mass. In this case one can show via duality that the law of the consensus time τ is that of $\mathsf{C}_{K \wedge n}$, where K is a $\mathbb{N}$-valued random variable independent of the coalescing random walks, defined by:

$$K = \max\{i \in \mathbb{N} \, : \, U_i = U_1\}, \text{ where } U_1, U_2, U_3 \ldots, \text{ are iid draws from } \mu$$

and for each $1 \leq k \leq n$

$$\mathsf{C}_k \equiv \min\{t \geq 0 \, : \, |\{\overline{X}_t(i) \, : \, 1 \leq i \leq n\}| = k\}.$$

Thus the key step in analyzing the voter model via our techniques is to prove approximations for the distribution of C_k. Theorems 1.1 and 1.2 imply mean-field behavior for $\mathsf{C} = \mathsf{C}_1$. A quick inspection of the proofs reveals that the same bounds for Wasserstein distance can be obtained for C_k for any $1 \leq k \leq n$. It follows that:

Theorem 1.3 (Proof omitted) *Let* $\mathbf{V}, \mathcal{O}$ *and* μ *be as above, and consider the voter model defined by* $\mathbf{V}$, $\mathcal{O}$ *and by the generator* Q *corresponding to transition rates* $q(x,y)$*. Assume that the sequence* $\{\mathsf{Z}_i\}_{i\geq 2}$ *is defined as in (1), and also that* K *has the law described above and is independent from the* Z_i*. Define* $\rho(Q)$ *and* $\alpha(Q)$ *as in Theorems 1.1 and 1.2. Then the consensus time* τ *for this voter model satisfies:*

$$d_W\left(\mathrm{Law}\left(\frac{\tau}{\mathsf{m}(Q)}\right), \mathrm{Law}\left(\sum_{i\geq K+1} Z_i\right)\right) = \begin{cases} O\left(\left(\rho(Q)\,\ln(1/\rho(Q))\right)^{1/6}\right), \\ \textit{if } Q \textit{ is reversible and transitive;} \\ O\left(\left(\alpha(Q)\,\ln(1/\alpha(Q))\,\ln^6|\mathbf{V}|\right)^{1/6}\right), \\ \textit{otherwise.} \end{cases}$$

1.3 Main proof ideas

Our proofs of Theorems 1.1 and 1.2 both start from the formula (1) for the terms in the distribution of C over K_n. Crucially, each term Z_i has a specific meaning: Z_i is the time it takes for a system with i particles to evolve to a system with $i-1$ particles, rescaled by the expected meeting time of two walkers. For $i = 2$, this is just the (rescaled) meeting time of a pair of particles, which is an exponential random variable with mean 1. For $i > 2$, we are looking at the first meeting time amoung $\binom{i}{2}$ pairs of particles. It turns out that these pairwise meeting times are independent; since the minimum of k independent exponential random variables with mean μ is an exponential r.v. with mean μ/k, we deduce that Z_i is exponential with mean $1/\binom{i}{2}$.

The bulk of our proof consists of proving something similar for more general chains Q. Fix some such Q, with state space $\mathbf{V}$, and let C_i denote the time it takes for a system of coalescing random walks evolving acording to Q to have i uncoalesced particles. Clearly, $M \equiv \mathsf{C}_1 - \mathsf{C}_2$ is the meeting time of a pair of particles, which is the hitting time of the diagonal set:

$$\Delta \equiv \{(x,x) \,:\, x \in \mathbf{V}\}$$

by the Markov chain $Q^{(2)}$ given by a pair of independent realizations of Q. More generally, $M^{(i+1)} = \mathsf{C}_i - \mathsf{C}_{i+1}$ is the hitting time of

$$\Delta^{(i)} = \{(x(1), \ldots, x(j)) \,:\, \exists 1 \leq i_1 < i_2 \leq i+1,\, x(i_1) = x(i_2)\}.$$

The mean-field picture suggestes that each $M^{(i+1)}$ should be close in distribution to Z_i. Indeed, it is easy to show that $M^{(i+1)}$ is approximately exponentially distributed if $\mathbb{E}\left[M^{(i+1)}\right] \gg \mathrm{T}^Q_{\mathrm{mix}}$. This is a general meta-result for small subsets of the state space of a Markov chain; precise versions (with different quantitative bounds) are proven in [3, 1] when the chain starts from the stationary distribution. However, we face a few difficulties when trying to use these off-the-shelf results.

1. For each i, $M^{(i+1)}$ is the first hitting time of $\Delta^{(i+1)}$ after time C_{i+1}. The random walkers are *not* stationary at this random time, so we need to "do" exponential approximation from non-stationary starting points.

2. In order to get Wasserstein approximation, we need better control of the tail of $M^{(i+1)}$;

3. To prove that Z_i and $M^{(i+1)}/\mathsf{m}(Q)$ are close, we must show something like that $\mathbb{E}\left[M^{(i+1)}\right] \approx \mathbb{E}\left[M\right]/\binom{i+1}{2}$, ie that $M^{(i+1)}$ behaves like the minimum of $\binom{i+1}{2}$ independent exponentials.

4. Finally, we should not expect the exponential approximation to hold when $\Delta^{(i+1)}$ is too large. That means that the "big bang" phase (to use Durrett's phrase) at the beginning of the process has to be controlled by other means.

It turns out that we can deal with points 1 and 2 via a different kind exponential approximation result, stated as Theorem 3.1 . This result will give bounds of the following form:

$$\mathbb{P}_x\left(H_A > t\right) = (1+o\left(1\right))\exp\left(-\frac{t}{(1+o\left(1\right))\mathbb{E}\left[H_A\right]}\right) \text{ if } \mathrm{T}^Q_{\mathrm{mix}} \ll \mathbb{E}\left[H_A\right]. \quad (4)$$

This holds even for *non stationary* starting points x if $\mathbb{P}_x\left(H_A \leq \mathrm{T}^Q_{\mathrm{mix}}\right) \ll 1$. This is treated in Section 3 below. We also take some time in that section to develop a specific notion of "near exponential random variable". Although this takes up some space, we believe it provides a useful framework for tackling other problems. We note that a version of Theorem 3.1 for stationary initial states result is implicit in [1].

We now turn to point 3. The key difficulty in our setting is that, unlike Cox [6] or Cooper et al. [5], we do not have a good "local" description of the graphs under consideration which we could use to compute $\mathbb{E}\left[M^{(i+1)}\right]$ directly. We use instead a simple general idea, which we believe to be new, to address this point. Clearly, $M^{(i+1)}$ is a minimum $\binom{i+1}{2}$ hitting times. Let us consider the *general* problem of understanding the law of:

$$H_B = \min_{1\leq i\leq \ell} H_{B_i} \text{ where } B = \cup_{i=1}^{\ell} B_i,$$

under the assumption that $\mathbb{E}\left[H_{B_i}\right] = \mu$ does not depend on i when the initial distribution is stationary (this covers the case of $M^{(i+1)}$). Assume also that (4) holds for all $A \in \{B, B_1, B_2, \ldots, B_\ell\}$. Then the following holds for ϵ in a suitable range:

$$\forall A \in \{B, B_1, B_2, \ldots, B_\ell\} \,:\, \mathbb{P}\left(H_A \leq \epsilon\mathbb{E}\left[H_A\right]\right) \approx \epsilon.$$

Morally speaking, this means that $\epsilon\mathbb{E}\left[H_A\right]$ is the ϵ-quantile of H_A for all A as above; this is implicit in [1] and is made explicit in our own Theorem 3.1. Now apply this to $A = B$, with ϵ replaced by $\epsilon\mu/\mathbb{E}\left[H_B\right]$, and obtains:

$$\frac{\epsilon\mu}{\mathbb{E}\left[H_B\right]} \approx \mathbb{P}\left(H_B \leq \epsilon\mu\right) = \mathbb{P}\left(\cup_{i=1}^{\ell}\{H_{B_i} \leq \epsilon\mu\}\right).$$

If we can show that the pairwise correlations between the events $\{H_{B_i} \leq \epsilon\mu\}$ are sufficiently small, then we may obtain:

$$\frac{\epsilon\mu}{\mathbb{E}\left[H_B\right]} \approx \mathbb{P}\left(\cup_{i=1}^{\ell}\{H_{B_i} \leq \epsilon\mu\}\right) \approx \sum_{i=1}^{\ell}\mathbb{P}\left(H_{B_i} \leq \epsilon\mu\right) = \ell\epsilon,$$

This gives:

$$\mathbb{E}[H_B] \approx \frac{\mu}{\ell},$$

as if the times $H_{B_1}, \ldots, H_{B_\ell}$ were independent exponentials. The reasoning presented here is made rigorous and quantitative in Theorem 3.2 below.

Finally, we need to take care of point 4, ie. the “big bang” phase. The case of Theorem 1.1 is covered by a result in [11]. The setting of Theorem 1.2 can be dealt with just using our results on the coalescence times of smaller number of particles, at the cost of a logarithmic factor. Incidentally, the differences in th bounds of the two theorems come from the fact that in the reversible/transitive we have a better control of correlations of pairwise meeting times, as well as a better bound for the expected length of the “big bang phase”.

1.4 Outline

The remainder of the paper is organized as follows. Section 2 contains several preliminaries. Section 3 contains a general discussion of random variables with nearly exponential distribution and our general approximation results for hitting times. In Section 4 we apply these results to the first meeting time among k particles, after proving some technical estimates. Section 5 contains the formal definition of the coalescing random walks process and proves mean field behavior for a moderate initial number of walkers. Finally, Section 6 contains the proofs of Theorems 1.1 and 1.2. Related results and open problems are discussed in the final sections.

2 Preliminaries

2.1 Basic notation

We write $\mathbb{N}$ for non-negative integers and $[k] = \{1, 2, \ldots, k\}$ for any $k \in \mathbb{N}\backslash\{0\}$.

We will often speak of universal constants $C > 0$. These are numbers that do not depend on any of the parameters or mathematical objects under consideration in a given problem. We will also use the notation “$a = O(b)$” in the universal sense prescribed in Notational convention 1. In this way we can write down expressions such as:

$$e^b = 1 + b + O\left(b^2\right) \text{ and } \ln\left(\frac{1}{1-b}\right) = b + O\left(b^2\right) = O(b).$$

Given a finite set S, we let $M_1(S)$ denote the set of all probability measures over S. Given $p, q \in M_1(S)$, their total variation disance is defined as follows.

$$d_{\mathrm{TV}}(p,q) \equiv \frac{1}{2}\sum_{x \in S} |p(s) - q(s)| = \sup_{A \subset S} p(A) - q(A),$$

where $p(A) = \sum_{a \in A} p(a)$. For S not finite, $M_1(S)$ will denote the set of all probability measures over the “natural” σ-field over S. For instance, for $S = \mathbb{R}$ we consider the Borel

σ-field, and for $S = \mathbb{D}([0,+\infty), \mathbf{V})$ (see Section 2.3.1 for a definition) we use the σ-field generated by projections.

If X is a random variable taking values over S, we let $\mathrm{Law}\,(X) \in M_1(S)$ denote the distribution (or law) of X. Here we again assume that there is a "natural" σ-field to work with.

2.2 Wasserstein distance

The L_1 *Wasserstein distance* (or simply Wasserstein distance) is a metric over probability measures over $\mathbb{R}$ with finite first moments, given by:

$$d_W(\lambda_1, \lambda_2) = \int_{\mathbb{R}} |\lambda_1(x, +\infty] - \lambda_2(x, +\infty]|\, dx \;\; (\lambda_1, \lambda_2 \in M_1(\mathbb{R})).$$

A classical duality result gives:

$$d_W(\lambda_1, \lambda_2) = \sup\left\{ \int_{\mathbb{R}} f(x)\lambda_1(dx) - \int_{\mathbb{R}} f(x)\lambda_2(dx) \,:\, f : \mathbb{R} \to \mathbb{R} \text{ 1-Lipschitz.} \right\}$$

Notational convention 2 *Whenever we compute Wasserstein distances, we will* assume *that the distributions involved have first moments. This can be checked in each particular case.*

Remark 1 *If Z_1, Z_2 are random variables, we sometimes write*

$$d_W(Z_1, Z_2) \textit{ instead of } d_W(\mathrm{Law}\,(Z_1), \mathrm{Law}\,(Z_2)).$$

Note that:

$$d_W(Z_1, Z_2) = \int_{\mathbb{R}} |\mathbb{P}\,(Z_1 \geq t) - \mathbb{P}\,(Z_2 \geq t)\,|\, dx.$$

Also notice that:

$$|\mathbb{E}\,[Z_1] - \mathbb{E}\,[Z_2]\,| \leq d_W(Z_1, Z_2).$$

This is an equality if $Z_1 \geq 0$ a.s. and $Z_2 = C\,Z_1$ for some constant $C > 0$:

$$\forall C \in \mathbb{R},\, d_W(Z_1, C\,Z_1) = |C - 1|\mathbb{E}\,[Z_1]\,, \tag{5}$$

since $|f(C\,Z_1) - f(Z_1)| \leq |C-1|\,Z_1$ for every 1-Lipschitz function $f : \mathbb{R} \to \mathbb{R}$.

We note here three useful lemmas on Wasserstein distance. These are probably standard, but we could not find references for them, so we provide proofs for the latter two lemmas in Section A of the Appendix. The first lemma is immediate.

Lemma 2.1 (Sum lemma for Wasserstein distance; proof omitted) *For any two random variables X, Y with finite first moments and defined on the same probability space:*

$$d_W(X, X + Y) \leq \mathbb{E}\,[|Y|]\,.$$

For the next lemma, recall that, given two $\mathbb{R}$-valued random variables X, Y, we say that X is stochastically dominated by Y, and write $X \preceq_d Y$, if $\mathbb{P}(X > t) \le \mathbb{P}(Y > t)$ for all $t \in \mathbb{R}$.

Lemma 2.2 (Sandwich lemma for Wasserstein distance) *Let Z, Z_-, Z_+, W be $\mathbb{R}$-valued random variables and assume $Z_- \preceq_d Z \preceq_d Z_+$. Then:*

$$d_W(Z, W) \le d_W(Z_-, W) + d_W(Z_+, W).$$

Lemma 2.3 (Conditional Lemma for Wasserstein distance) *Let W_1, W_2, Z_1, Z_2 be real-valued random variables with finite first moments. Assume that Z_1 and Z_2 independent and that W_1 is $\mathcal{G}$-measurable for some sub-σ-field $\mathcal{G}$. Then:*

$$\begin{aligned} &d_W(\mathrm{Law}\,(W_1 + W_2)\,, \mathrm{Law}\,(Z_1 + Z_2)) \\ &\qquad\qquad \le d_W(\mathrm{Law}\,(W_1)\,, \mathrm{Law}\,(Z_1)) + \mathbb{E}\left[d_W(\mathrm{Law}\,(W_2 \mid \mathcal{G})\,, \mathrm{Law}\,(Z_2))\right] \end{aligned}$$

Remark 2 *Here we are implicitly assuming that* $\mathrm{Law}\,(W_2 \mid \mathcal{G})$ *is given by some regular conditional probability distribution*

2.3 Continuous-time Markov chains

2.3.1 State space and trajectories

Let $\mathbf{V}$ be some non-empty finite set, called the *state space*. We write $\mathbb{D} \equiv \mathbb{D}([0, +\infty), \mathbf{V})$ for the set of all paths:

$$\omega : t \ge 0 \mapsto \omega_t \in \mathbf{V}$$

for which there exist $0 = t_0 < t_1 < t_2 < \cdots < t_n < \dots$ with $t_n \nearrow +\infty$ and ω constant over each interval $[t_n, t_{n+1})$ $(n \in \mathbb{N})$. Such paths will sometimes be called *càdlàg*.

For each $t \ge 0$, we let $X_t : \mathbb{D} \to \mathbf{V}$ be the projection map sending ω to ω_t. We also define $X = (X_t)_{t \ge 0}$ as the identity map over $\mathbb{D}$. Whenever we speak about probability measures and events over $\mathbb{D}$, we will implicitly use the σ-field $\sigma(\mathbb{D})$ generated by the maps X_t, $t \ge 0$. We define an associated filtration as follows:

$$\mathcal{F}_t \equiv \sigma\{X_s \,:\, 0 \le s \le t\} \;\; (t \ge 0).$$

We also define the time-shift operators:

$$\Theta_T : \omega(\cdot) \in \mathbb{D} \mapsto \omega(\cdot + T) \in \mathbb{D} \; (T \ge 0).$$

2.3.2 Markov chains and their generators

Let $q(x, y)$ be non-negative real numbers for each pair $(x, y) \in \mathbf{V}^2$ with $x \ne y$. Define a linear operator $Q : \mathbb{R}^{\mathbf{V}} \to \mathbb{R}^{\mathbf{V}}$, which maps $f \in \mathbb{R}^{\mathbf{V}}$ to $Qf \in \mathbb{R}^{\mathbf{V}}$ satisfying:

$$(Qf)(x) \equiv \sum_{y \in \mathbf{V} \setminus \{x\}} q(x, y)(f(x) - f(y)) \; (x \in \mathbf{V}).$$

It is a well-known result that there exists a unique family of probability measures $\{\mathbb{P}_x\}_{x\in\mathbf{V}}$ with the properties listed below:

1. for all $x \in \mathbf{V}$, $\mathbb{P}_x(X_0 = x) = 1$;
2. for all distinct $x, y \in \mathbf{V}$, $\lim_{\epsilon\searrow 0} \frac{\mathbb{P}_x(X_\epsilon = y)}{\epsilon} = q(x,y)$;
3. **Markov property:** for any $x \in \mathbf{V}$ and $T \geq 0$, the conditional law of $X \circ \Theta_T$ given $\mathcal{F}_T$ under measure $\mathbb{P}_x$ is given by $\mathbb{P}_{X_T}$.

The family $\{\mathbb{P}_x\}_{x\in\mathbf{V}}$ satisfying these properties is the *Markov chain with generator Q*. We will often abuse notation and omit any distinction between a Markov chain and its generator in its notation.

For $\lambda \in M_1(\mathbf{V})$, $\mathbb{P}_\lambda$ denotes the mixture:

$$\mathbb{P}_\lambda \equiv \sum_{x\in\mathbf{V}} \lambda(x)\,\mathbb{P}_x.$$

This corresponds to starting the process from a random state distributed according to λ. For $x \in \mathbf{V}$ or $\lambda \in M_1(E)$ and $Y : \mathbb{D} \to S$ a random variable, we let $\mathrm{Law}_x(Y)$ or $\mathrm{Law}_\lambda(Y)$ denote the law (or distribution) of Y under $\mathbb{P}_x$ or $\mathbb{P}_\lambda$.

2.3.3 Stationary measures and mixing

Any Markov chain Q as above has at least one stationary measure $\pi \in M_1(\mathbf{V})$; this is a measure such that for any $T \geq 0$,

$$\mathrm{Law}_\pi(X \circ \Theta_T) = \mathrm{Law}_\pi(X).$$

We will be only interested in *mixing Markov chains*, which are those Q with a unique stationary measure that satisfy the following condition:

$$\forall \alpha \in (0,1),\ \exists T = \mathrm{T}^Q_{\mathrm{mix}}(\alpha) \geq 0, \forall x \in \mathbf{V},\ \forall S \subset \mathbf{V}\ :\ |\mathbb{P}_x(X_T \in S) - \pi(S)| \leq \alpha.$$

The value $\mathrm{T}^Q_{\mathrm{mix}}(\alpha)$ is called the α-mixing time of Q. By the Markov property, we also have:

$$\forall \alpha \in (0,1), \forall t \geq \mathrm{T}^Q_{\mathrm{mix}}(\alpha), \forall x \in \mathbf{V}, \forall S \subset \mathbf{V}\ :\ |\mathbb{P}_x(X_t \in S) - \pi(S)| \leq \alpha.$$

$\mathrm{T}^Q_{\mathrm{mix}} \equiv \mathrm{T}^Q_{\mathrm{mix}}(1/2e)$ is also called *the* mixing time of Q. We note that for all $\epsilon \in (0, 1/2)$:

$$\mathrm{T}^Q_{\mathrm{mix}}(\epsilon) \leq C\,\ln(1/\epsilon)\,\mathrm{T}^Q_{\mathrm{mix}}$$

where $C > 0$ is universal.

2.3.4 Product chains

Letting Q be as above, we may consider the joint trajectory of k independent realizations of Q:

$$X_t^{(k)} = (X_t(1), \ldots, X_t(k)) \quad (t \geq 0)$$

where each $(X_t(i))_{t\geq 0}$ has law $\mathbb{P}_{x(i)}$. It turns out that this corresponds to a Markov chain $Q^{(k)}$ on $\mathbf{V}^k$ with transition probabilities:

$$q^{(k)}(x^{(k)}, y^{(k)}) = \begin{cases} q(x(i), y(i)), & \text{if } x(i) \neq y(i) \wedge \forall j \in [k] \backslash \{i\},\, x(j) = y(j); \\ 0, & \text{otherwise.} \end{cases}$$

The following well-known results on $Q^{(k)}$ will often be useful:

Fact 1 *Assume Q is mixing and has (unique) stationary distribution π. Then $Q^{(k)}$ is also mixing, and the product measure $\pi^{\otimes k}$ is its (unique) stationary distribution. Moreover, the mixing times of $Q^{(k)}$ satisfy:*

$$\forall \alpha \in (0, 1/2),\, \mathrm{T}_{\mathrm{mix}}^{Q^{(k)}}(\alpha) \leq \mathrm{T}_{\mathrm{mix}}^{Q}(\alpha/k) \leq C\, \ln(k/\alpha)\, \mathrm{T}_{\mathrm{mix}}^{Q}$$

with $C > 0$ universal.

Remark 3 *In what follows we will always denote elements of $\mathbf{V}^k$ (resp. $M_1(\mathbf{V}^k)$ by symbols like $x^{(k)}, y^{(k)}, \ldots$ (resp. $\lambda^{(k)}, \rho^{(k)}, \ldots$.) We will then denote the distribution of $Q^{(k)}$ started from $x^{(k)}$ or $\lambda^{(k)}$ by $\mathbb{P}_{x^{(k)}}$ or $\mathbb{P}_{\lambda^{(k)}}$. This is a slight abouse of our convention for the Q chain, but the initial state/distribution will always make it clear that we are referring to the product chain.*

3 Nearly-exponential random variables and hitting times of Markov chains

3.1 Basic definitions

We first recall a standard definition: the *exponential distribution* with mean $m > 0$, denoted by $\mathsf{Exp}(m)$, is the unique probabilty dstribution $\mu \in M_1(\mathbb{R})$ such that, if Z is a random variable with law μ,

$$\mathbb{P}\left(Z \geq t\right) = e^{-t/m} \wedge 1 \ (t \geq 0).$$

We write $Z =_d \mathsf{Exp}(m)$ when Z is a random variable with $\mathrm{Law}\,(Z) = \mathsf{Exp}(m)$.

Similarly, given $m > 0$ as above and parameters $\alpha > 0, \beta \in (0,1)$, we say that a measure $\mu \in M_1(\mathbb{R})$ has distribution $\mathsf{Exp}(m, \alpha, \beta)$ if it is the law of a random variable $\widetilde{Z}$ with $\widetilde{Z} \geq 0$ almost surely and for all $t > 0$:

$$(1-\alpha)_+\, e^{-\frac{t}{(1-\beta)_+ m}} \leq \mathbb{P}\left(\widetilde{Z} \geq t\right) \leq (1+\alpha) e^{-\frac{t}{(1+\beta)m}}$$

We will write $\mu = \mathsf{Exp}(m, \alpha, \beta)$ or $\widetilde{Z} =_d \mathsf{Exp}(m, \alpha, \beta)$ as a shorthand for this. Notice that $\mathsf{Exp}(m, \alpha, \beta)$ does not denote a single distribution, but rather a family of distributions that obey the above property, but we will mostly neglect this minor issue.

Random variables with law $\mathsf{Exp}(m, \alpha, \beta)$ will naturally appear in our study of hitting times of Markov chains. We compile here some simple results about them. The first proposition is trivial and we omit its proof.

Proposition 3.1 (Proof omitted) *If $\mu \in M_1(\mathbb{R})$ satisfies $\mu = \mathsf{Exp}(m, \alpha, \beta)$ and $m' > 0, \gamma \in (0, 1)$ are such that $\beta + \gamma + \beta\gamma < 1$,*

$$(1 - \gamma)\, m' \leq m \leq (1 + \gamma)\, m',$$

then

$$\mu = \mathsf{Exp}(m', \alpha, \beta + \gamma + \beta\gamma)$$

We now show that random variables $\mathsf{Exp}(m, \alpha, \beta)$ are close to the corresponding exponentials.

Lemma 3.1 (Wasserstein distance error for $\mathsf{Exp}(m, \alpha, \beta)$) *We have the following inequality for all $\alpha, \beta > 0$:*

$$d_W(\mathsf{Exp}(m), \mathsf{Exp}(m, \alpha, \beta)) \leq 2(\alpha + \beta)\, m.$$

That is, if $\widetilde{Z} =_d \mathsf{Exp}(m, \alpha, \beta)$, the Wasserstein distance between $\mathrm{Law}\left(\widetilde{Z}\right)$ *and $\mathsf{Exp}(m)$ is at most $2\alpha\, m + 2\beta\, m$.*

Proof: Assume $\widetilde{Z} =_d \mathsf{Exp}(m, \alpha, \beta)$ and $Z =_d \mathsf{Exp}(m)$ are given. By convexity:

$$\begin{aligned} d_W(\widetilde{Z}, Z) = \int_0^{+\infty} |\mathbb{P}\left(\widetilde{Z} \geq t\right) - e^{-\frac{t}{m}}|\, dt \\ \leq \int_0^{\infty} \max_{\xi \in \{-1, +1\}} |(1 + \xi\alpha)_+ e^{-\frac{t}{(1+\xi\beta)m}} - e^{-\frac{t}{m}}|\, dt \\ \leq \int_0^{\infty} |(1 + \alpha) e^{-\frac{t}{(1+\beta)m}} - e^{-\frac{t}{m}}|\, dt + \int_0^{\infty} |(1 - \alpha)_+\, e^{-\frac{t}{(1-\beta)m}} - e^{-\frac{t}{m}}|\, dt \\ =: (I) + (II). \end{aligned} \tag{6}$$

For the first term in the RHS, we note that

$$\forall t \geq 0,\ (1 + \alpha) e^{-\frac{t}{(1+\beta)m}} - e^{-\frac{t}{m}} \geq 0,$$

hence:

$$(I) = \int_0^{\infty} \{(1 + \alpha) e^{-\frac{t}{(1+\beta)m}} - e^{-\frac{t}{m}}\}\, dt = [\alpha + \beta + \alpha\beta]\, m.$$

Similarly, for term (II) we have:

$$\forall t \geq 0,\ (1-\alpha)_+\, e^{-\frac{t}{(1-\beta)m}} - e^{-\frac{t}{m}} \leq 0$$

hence:

$$(II) = \int_0^\infty \{e^{-\frac{t}{m}} - (1-\alpha)_+ e^{-\frac{t}{(1-\beta)m}}\}\, dt \leq [\alpha + \beta - \alpha\beta] m.$$

Hence:

$$d_W(\widetilde{Z}, Z) \leq (I) + (II) = 2(\alpha + \beta)\, m.$$

□

3.2 Hitting times are nearly exponential

In this section we consider a mixing continuous-time Markov chain $\{\mathbb{P}_x\}_{x\in\mathbf{V}}$ taking values over a finite state space $\mathbf{V}$, with unique stationary distribution π. Given a nonempty $A \subset \mathbf{V}$ with $\pi(A) > 0$, we define the *hitting time of* A to be:

$$H_A(\omega) \equiv \inf\{t \geq 0 \,:\, \omega(t) \in A\} \ \ (\omega \in \mathbb{D}([0,+\infty), \mathbf{V})).$$

The condition $\pi(A) > 0$ ensures that $\mathbb{E}_x[H_A] < +\infty$ for all $x \in \mathbf{V}$.

Our first result in this section presents sufficient conditions on A and $\mu \in M_1(\mathbf{V})$ that ensure that H_A is approximately exponentially distributed.

Theorem 3.1 *There exists a universal constant $C_0 > 0$ such that the following holds. In the above Markov chain setting, assume that $0 < \epsilon < \delta < 1/5$ are such that:*

$$\mathbb{P}_\pi (H_A \leq \mathrm{T}_{\mathrm{mix}}(\delta\epsilon)) \leq \delta\epsilon.$$

Let $t_\epsilon(A)$ be the ϵ-quantile of $\mathrm{Law}_\pi(H_A)$, ie. the unique number $t_\epsilon(A) \in [0,+\infty)$ with $\mathbb{P}_\pi(H_A \leq t_\epsilon(A)) = \epsilon$ (this is well-defined since $\mathbb{P}_\pi(H_A \leq t)$ is a continuous and strictly increasing function of t in our setting).Given $\lambda \in M_1(\mathbf{V})$, write:

$$r_\lambda \equiv \mathbb{P}_\lambda (H_A \leq \mathrm{T}_{\mathrm{mix}}(\delta\epsilon)).$$

Then:

$$\mathrm{Law}_\lambda(H_A) = \mathsf{Exp}\left(\frac{t_\epsilon(A)}{\epsilon}, O(\epsilon) + 2r_\lambda, O(\delta)\right).$$

Moreover,

$$\left|\frac{\epsilon\mathbb{E}_\pi[H_A]}{t_\epsilon(A)} - 1\right| \leq O(\delta)$$

and:

$$\mathrm{Law}_\lambda(H_A) =_d \mathsf{Exp}(\mathbb{E}_\pi[H_A], O(\epsilon) + 2r_\lambda, O(\delta)).$$

We emphasize that results similar to this are not new in the literature [1, 3], but the lower-tail part of our result does not seem to be explicitly anywhere. The proof is strongly related to that in [1], but we wish to stress the relationship between the quantile $t_\epsilon(A)$ and the exponential approximation, which we will need below.

The second result considers what happens when we have an union of events

$$A = A_1 \cup A_2 \cup \cdots \cup A_\ell.$$

As described in the introduction, we give a sufficient condition under which the hitting time H_A behaves like a minimum of independent exponentials.

Theorem 3.2 *Assume that the set A considered above can be written as:*

$$A = \bigcup_{i=1}^{\ell} A_i$$

where the sets $A_1, \dots, A_\ell$ are non-empty and

$$m := \mathbb{E}_\pi\left[H_{A_1}\right] = \mathbb{E}_\pi\left[H_{A_2}\right] = \cdots = \mathbb{E}_\pi\left[H_{A_\ell}\right].$$

Assume $0 < \delta < 1/5$, $0 < \epsilon < \delta/\ell$ are such that for all $1 \le i \le \ell$:

$$\forall i \in [\ell],\ \mathbb{P}_\pi\left(H_{A_i} \le \mathrm{T}_{\mathrm{mix}}(\delta\epsilon/2)\right) \le \frac{\delta\epsilon}{3}.$$

Then for all $\lambda \in M_1(\mathbf{V})$,

$$\mathrm{Law}_\lambda\left(H_A\right) = \mathsf{Exp}\left(\frac{m}{\ell}, 2r_\lambda + O\left(\ell\epsilon\right), O\left(\delta + \xi\right)\right)$$

where

$$r_\lambda \equiv \mathbb{P}_\lambda\left(H_A \le \mathrm{T}_{\mathrm{mix}}(\delta\epsilon)\right),$$

and:

$$\xi \equiv \frac{1}{\ell\epsilon} \sum_{1 \le i < j \le \ell} \mathbb{P}_\pi\left(H_{A_i} \le \epsilon\, m,\ H_{A_j} \le \epsilon\, m\right).$$

Remark 4 *If the H_{A_i} are in fact independent, then $\xi = O\left(\epsilon\ell\right)$.*

The remainder of the section is devoted to the proof of these two results.

3.3 Hitting time of a single set: proofs

We first present the proof of Theorem 3.1 modulo two important Lemmas, and subsequently prove those Lemmas.

Proof: [of Theorem 3.1]

Let $\lambda \in M_1(\mathbf{V})$ be arbitrary. Throughout the proof we will assume implicitly that $\delta + r_\lambda + \epsilon$ is smaller than some sufficiently small absolute constant; the remaining case is easy to handle by increasing the value of C_0 if necessary.

We begin with an upper bound for $\mathbb{P}_\lambda\left(H_A \ge t\right)$ in terms of $t_\epsilon(A)$.

Lemma 3.2 (Proven in Section 3.3.1) *Under the assumptions of Theorem 3.1,*

$$\forall t \geq 0,\, \mathbb{P}_\lambda \left(H_A \geq t\right) \leq \left(1 + O\left(\epsilon\right)\right) e^{-\frac{\epsilon(1+O(\delta))\,t}{t_\epsilon(A)}}.$$

In particular, this implies:

$$\forall \mu \in M_1(\mathbf{V}),\, \mathbb{E}_\mu\left[H_A\right] = \int_0^{+\infty} \mathbb{P}_\mu\left(H_A \geq t\right)\, dt \leq \left(1 + O\left(\delta\right)\right) \frac{t_\epsilon(A)}{\epsilon}. \tag{7}$$

It turns out that the upper bound in the above Lemma can be nearly reversed if we start from some distribution that is "far" from A.

Lemma 3.3 (Proven in Section 3.3.2) *With the assumptions of Theorem 3.1, if* $2\epsilon + r_\lambda < 1/2$*:*

$$\forall t \geq 0,\, \mathbb{P}_\lambda\left(H_A \geq t\right) \geq \left(1 - O\left(\epsilon\right) - r_\lambda\right)_+ e^{-\frac{\epsilon(1+O(\delta))\,t}{t_\epsilon(A)}}.$$

Notice that the combination of these two lemmas already implies the first statement in the proof, as it shows that:

$$\forall t \geq 0,\, \mathbb{P}_\lambda\left(H_A \geq t\right) \in [(1 - O\left(\epsilon\right) - 2r_\lambda)_+ e^{-\frac{\epsilon t}{(1+O(\delta))\,t_\epsilon(A)}}, (1 + O\left(\epsilon\right))_+ e^{-\frac{\epsilon t}{(1+O(\delta))\,t_\epsilon(A)}}].$$

To see this, notice that the upper bound is always valid by Lemma 3.2. For the lower bound, we use Lemma 3.3 if $2\epsilon + r_\lambda \leq 1/2$, and note that the lower bound is 0 if $2\epsilon + r_\lambda > 1/2$ and the constant in the $O\left(\epsilon\right)$ term is at least 4.

We now prove the assertion about expectations in the Theorem. We use Lemma 3.1 and deduce:

$$|\mathbb{E}_\pi\left[H_A\right] - \epsilon^{-1}\, t_\epsilon(A)| \leq d_W(\mathrm{Law}_\pi\left(H_A\right), \mathsf{Exp}(\epsilon^{-1}\, t_\epsilon(A))) \leq O\left(\delta + r_\pi\right)\, \epsilon^{-1}\, t_\epsilon(A)$$

and the assertion follows from dividing by $\epsilon^{-1}\, t_\epsilon(A)$ and

$$r_\pi = \mathbb{P}_\pi\left(H_A \leq \mathrm{T}_{\mathrm{mix}}(\delta\epsilon)\right) \leq \delta\epsilon$$

by assumption. The final assertion in the Theorem then follows from Proposition 3.1. $\square$

3.3.1 Proof of Lemma 3.2

Proof: Set $T = \mathrm{T}_{\mathrm{mix}}(\delta\epsilon)$. We note for later reference that $T < t_\epsilon(A)$, since

$$\mathbb{P}_\pi\left(H_A \leq T\right) \leq \delta\epsilon < \epsilon = \mathbb{P}_\pi\left(H_A \leq t_\epsilon(A)\right).$$

Our main goal will be to show the following inequality:

$$\textbf{Goal}: \forall k \in \mathbb{N},\, \mathbb{P}_\lambda\left(H_A > (k+1)t_\epsilon(A)\right) \leq \left(1 - \epsilon + 2\delta\epsilon\right) \mathbb{P}_\lambda\left(H_A > k t_\epsilon(A)\right). \tag{8}$$

Once established, this goal will imply:

$$\forall k \in \mathbb{N},\, \mathbb{P}_\lambda\left(H_A \geq k t_\epsilon(A)\right) \leq (1-\epsilon+2\delta\epsilon)^k$$

and

$$\forall t \geq 0,\, \mathbb{P}_\lambda\left(H_A \geq t\right) \leq e^{-\epsilon(1+O(\delta))\left\lfloor \frac{t}{t_\epsilon(A)}\right\rfloor} = (1+O(\epsilon))\, e^{-\frac{\epsilon\, t}{(1+O(\delta))t_\epsilon(A)}},$$

which is the desired result. To achieve the goal, we fix some $k\in\mathbb{N}$ and use $T\leq t_\epsilon(A)$ to bound:

$$\begin{aligned}\mathbb{P}_\lambda\left(H_A > (k+1)t_\epsilon(A)\right) &\leq \mathbb{P}_\lambda\left(H_A > k t_\epsilon(A),\, H_A\circ\Theta_{kt_\epsilon(A)+T} > t_\epsilon(A)-T\right) \qquad (9)\\ \text{(Markov prop.)} &= \mathbb{E}_\lambda\left[\mathbb{I}_{\{H_A>kt_\epsilon(A)\}}\,\mathbb{P}_{X_{kt_\epsilon(A)+T}}\left(H_A > t_\epsilon(A)-T\right)\right].\end{aligned}$$

Since $T=\mathrm{T}_{\mathrm{mix}}(\delta\epsilon)$, the conditional law of $X_{kt_\epsilon(A)+T}$ given $\mathcal{F}_{kt_\epsilon(A)}$ is $\delta\epsilon$-close to π. Since $\{H_A > kt_\epsilon(A)\}\in\mathcal{F}_{kt_\epsilon(A)}$ that:

$$\begin{aligned}&\mathbb{E}_\lambda\left[\mathbb{I}_{\{H_A>kt_\epsilon(A)\}}\,\mathbb{P}_{X_{kt_\epsilon(A)+T}}\left(H_A > t_\epsilon(A)-T\right)\right]\\ &\qquad\leq \mathbb{E}_\lambda\left[\mathbb{I}_{\{H_A>kt_\epsilon(A)\}}\left(\mathbb{P}_\pi\left(H_A>t_\epsilon(A)-T\right)+\delta\epsilon\right)\right]\\ &\qquad= \mathbb{P}_\lambda\left(H_A\geq kt_\epsilon(A)\right)\left(\mathbb{P}_\pi\left(H_A>t_\epsilon(A)-T\right)+\delta\epsilon\right).\end{aligned}$$

Now:

$$\begin{aligned}\mathbb{P}_\pi\left(H_A>t_\epsilon(A)-T\right) &\leq \mathbb{P}_\pi\left(H_A>t_\epsilon(A)\right)+\mathbb{P}_\pi\left(H_A\in(t_\epsilon(A)-T,t_\epsilon(A)]\right)\\ &\leq \mathbb{P}_\pi\left(H_A>t_\epsilon(A)\right)+\mathbb{P}_\pi\left(H_A\circ\Theta_{t_\epsilon(A)-T}\leq T\right)\\ &\leq \epsilon+\mathbb{P}_\pi\left(H_A\leq T\right)\ \text{(using π stationary + defn. of $t_\epsilon(A)$)}\\ &\leq \epsilon+\delta\epsilon\ \text{(by the assumption on $\mathbb{P}_\pi\left(H_A\leq T\right)$)}\end{aligned}$$

and plugging this into the previous equation gives:

$$\mathbb{P}_\lambda\left(H_A\geq (k+1)\,t_\epsilon(A)\right)\leq (1-\epsilon(1-2\delta))\,\mathbb{P}_\lambda\left(H_A\geq k\,t_\epsilon(A)\right),$$

as desired. □

3.3.2 Proof of Lemma 3.3

Proof: The general scheme of the proof is similar to that of Lemma 3.2, but we will need to be a bit more careful in our estimates. In particular, we will need that $(1+5\delta)\epsilon<1/2$ and $2\epsilon+r_\lambda<1/2$.

Define $T\equiv\mathrm{T}_{\mathrm{mix}}(\delta\epsilon)$ as in the proof of Lemma 3.2 in Section 3.3.1. Again observe that $T<t_\epsilon(A)$. Define:

$$f(k)\equiv\mathbb{P}_\lambda\left(H_A\geq kt_\epsilon(A)\right)\quad(k\in\mathbb{N}).$$

Clearly, $f(0)=1$ and:

$$f(1)\geq\mathbb{P}_\lambda\left(H_A\circ\Theta_T\geq t_\epsilon(A)\right)-\mathbb{P}_\lambda\left(H_A\leq T\right)\geq 1-\epsilon-\delta\epsilon-r_\lambda\geq 1-2\epsilon-r_\lambda \qquad (10)$$

since $T = \mathrm{T_{mix}}(\delta\epsilon)$ and by the properties of mixing times:

$$\mathbb{P}_\lambda\left(H_A \circ \Theta_T \geq t_\epsilon(A)\right) \geq \mathbb{P}_\pi\left(H_A \geq t_\epsilon(A)\right) - \delta\epsilon.$$

We now claim that:

Claim 3.1 *For all $k \in \mathbb{N}\backslash\{0\}$,*

$$\frac{f(k+1)}{f(k)} \geq (1 - \epsilon - 5\delta\epsilon).$$

Notice that the Claim and (10) imply:

$$\begin{aligned}
\forall t \geq 0,\ \mathbb{P}_\lambda\left(H_A \geq t\right) &\geq f(\lceil t/t_\epsilon(A)\rceil) \\
&\geq (1 - 2\epsilon - r_\lambda)\,(1 - \epsilon - 5\delta\epsilon)^{\lceil \frac{t}{t_\epsilon(A)}\rceil - 1} \\
&= (1 - O(\epsilon) - r_\lambda)\,(1 - \epsilon - 5\delta\epsilon)^{\frac{t}{t_\epsilon(A)}} \geq (1 - O(\epsilon) - r_\lambda)\, e^{-(1+O(\delta))\frac{\epsilon\, t}{t_\epsilon(A)}},
\end{aligned}$$

which is precisely the bound we wish to prove. We spend the rest of this proof proving the Claim.

Fix some $k \geq 1$ and notice that:

$$\begin{aligned}
f(k+1) &= \mathbb{P}_\lambda\left(H_A \geq kt_\epsilon(A),\, H_A \circ \Theta_{kt_\epsilon(A)} \geq T,\, H_A \circ \Theta_{kt_\epsilon(A)+T} \geq t_\epsilon(A) - T\right) \\
&\geq \mathbb{P}_\lambda\left(H_A \geq kt_\epsilon(A),\, H_A \circ \Theta_{kt_\epsilon(A)+T} \geq t_\epsilon(A) - T\right) \\
&\quad - \mathbb{P}_\lambda\left(H_A \geq kt_\epsilon(A),\, H_A \circ \Theta_{kt_\epsilon(A)} < T\right). \qquad (11)
\end{aligned}$$

We bound the two terms in the RHS of (11) separately. The first term is lower bounded by:

$$\mathbb{P}_\lambda\left(H_A \geq kt_\epsilon(A),\, H_A \circ \Theta_{kt_\epsilon(A)+T} \geq t_\epsilon(A)\right) = \mathbb{E}_\lambda\left[\mathbb{I}_{\{H_A \geq kt_\epsilon(A)\}}\mathbb{P}_{X_{kt_\epsilon(A)+T}}\left(H_A \geq t_\epsilon(A)\right)\right];$$

we used the Markov property for the second inequality. Since $T = \mathrm{T_{mix}}(\epsilon\delta)$, the law of $X_{kt_\epsilon(A)+T}$ conditioned $\mathcal{F}_{kt_\epsilon(A)}$ is within distance $\delta\epsilon$ from π. Since $\{H_A \geq kt_\epsilon(A)\} \in \mathcal{F}_{kt_\epsilon(A)}$, we deduce:

$$\begin{aligned}
&\mathbb{P}_\lambda\left(H_A \geq kt_\epsilon(A),\, H_A \circ \Theta_{kt_\epsilon(A)+T} \geq t_\epsilon(A) - T\right) \\
&\qquad \geq \mathbb{E}_\lambda\left[\mathbb{I}_{\{H_A \geq kt_\epsilon(A)\}}(\mathbb{P}_\pi\left(H_A \geq t_\epsilon(A)\right) - \delta\epsilon)\right] = (\epsilon - \delta\epsilon)\, f(k). \qquad (12)
\end{aligned}$$

We now upper bound the second term in (11). Notice that (again because of the Markov property):

$$\begin{aligned}
\mathbb{P}_\pi\left(H_A \geq kt_\epsilon(A),\, H_A \circ \Theta_{kt_\epsilon(A)} < T\right) &\leq \mathbb{P}_\pi\left(H_A \geq (k-1)t_\epsilon(A),\, H_A \circ \Theta_{kt_\epsilon(A)} < T\right) \\
&= \mathbb{E}_\pi\left[\mathbb{I}_{\{H_A \geq (k-1)t_\epsilon(A)\}}\mathbb{P}_{X_{kt_\epsilon(A)}}\left(H_A < T\right)\right].
\end{aligned}$$

Recalling that $t_\epsilon(A) \geq T = \mathrm{T_{mix}}(\delta\epsilon)$, we see that the law of $X_{kt_\epsilon(A)}$ given $\mathcal{F}_{(k-1)t_\epsilon(A)}$ is $\delta\epsilon$-far from π. Moreover, we have assumed that $\mathbb{P}_\pi\left(H_A \leq T\right) \leq \delta\epsilon$. We deduce:

$$\begin{aligned}\mathbb{P}_\pi\left(H_A \geq kt_\epsilon(A),\, H_A \circ \Theta_{kt_\epsilon(A)} < T\right) &\leq \mathbb{E}_\pi\left[\mathbb{I}_{\{H_A \geq (k-1)t_\epsilon(A)\}}(\mathbb{P}_\pi\left(H_A < T\right) + \delta\epsilon)\right] \\ &\leq 2\delta\epsilon\, f(k-1).\end{aligned}$$

Combining (12) and (11), we obtain:

$$\forall k \in \mathbb{N}\backslash\{0,1\}\; f(k+1) \geq f(k)(1-\epsilon-\delta\epsilon) - f(k-1)(2\delta\epsilon).$$

One can argue inductively that $f(k)/f(k-1) \geq 1/2$ for all $k \geq 1$. Indeed, this holds for $k \geq 2$ by the Claim applied to $k-1$. For $k=1$ we may use (10) and the assumption on $2\epsilon + r_\lambda$ to deduce the same result. Applying this to the previous inequality we obtain

$$\forall k \in \mathbb{N}\backslash\{0\}\; f(k+1) \geq f(k)(1-\epsilon-5\delta\epsilon),$$

which finishes the proof of the Claim and of the Lemma. $\square$

3.4 Hitting times of a union of sets: proofs

We present the proof of Theorem 3.2 below.
Proof: [of Theorem 3.2] There are three main steps in the proof, here outlined in a slightly oversimplified way.

1. We show that Theorem 3.1 is applicable to the hitting times of $A_1, \ldots, A_\ell$. In particular, this shows that $\mathbb{P}_\pi\left(H_{A_i} \leq \epsilon\, m\right) \approx \epsilon$.

2. We show that
$$\mathbb{P}_\pi\left(H_A \leq \epsilon\, m\right) \approx \sum_{i=1}^{\ell} \mathbb{P}_\pi\left(H_{A_i} \leq \epsilon\, m\right) \approx \ell\epsilon,$$
so that $t_{\ell\epsilon}(A) \approx \epsilon\, m$.

3. Finally, we apply Theorem 3.1 to H_A and deduce that this random variable is approximately exponential with mean:
$$\mathbb{E}_\pi\left[H_A\right] \approx t_{\epsilon\ell}(A)/\epsilon\ell \approx m/\ell.$$

The actual proof is only slightly more complicated than this outline. We begin with a claim corresponding to step 1 above.

Claim 3.2 *For all* $1 \leq i \leq \ell$,

$$\epsilon_i \equiv \mathbb{P}_\pi\left(H_{A_i} \leq \epsilon\, m\right) = \left(1 + O\left(\delta\right)\right)\epsilon.$$

Proof: [of the Claim] Consider some $\epsilon' \in [\epsilon/2, 2\epsilon]$. Notice that $\mathrm{T}_{\mathrm{mix}}(\delta\epsilon') \leq \mathrm{T}_{\mathrm{mix}}(\delta\epsilon/2)$ and therefore:

$$\mathbb{P}_\pi\left(H_{A_i} \leq \mathrm{T}_{\mathrm{mix}}(\delta\epsilon')\right) \leq \frac{\delta\epsilon}{2} \leq \delta\epsilon'.$$

This shows that Theorem 3.1 is applicable with A_i replacing A and ϵ' replacing ϵ. We deduce in particular that:

$$\forall \frac{\epsilon}{2} \leq \epsilon' \leq 2\epsilon, \left|\frac{\epsilon'\mathbb{E}_\pi\left[H_{A_i}\right]}{t_{\epsilon'}(A_i)} - 1\right| \leq O\left(\delta + \epsilon'\right) = O\left(\delta\right).$$

In particular, there exists a universal constant $c > 0$ such that if $\epsilon' \leq (1 - c\delta)\,\epsilon$, then $t_{\epsilon'}(A_i) < \epsilon\mathbb{E}_\pi\left[H_{A_i}\right]$, whereas if $\epsilon' > (1 + c\delta)\,\epsilon$, $t_{\epsilon'}(A_i) > \epsilon\mathbb{E}_\pi\left[H_{A_i}\right]$. In other words,

$$(1 - c\delta)\,\epsilon \leq \mathbb{P}_\pi\left(H_{A_i} \leq \epsilon\mathbb{E}_\pi\left[H_{A_i}\right]\right) \leq (1 + c\delta)\epsilon.$$

□

We now come to the second part of the proof.

Claim 3.3 *Let ξ be as in the statement of Theorem 3.2. Then:*

$$\mathbb{P}_\pi\left(H_A \leq \epsilon\, m\right) = \left(1 + O\left(\delta + \xi\right)\right)\ell\,\epsilon.$$

In particular, there exists a number $\eta = (1 + O\left(\delta + \xi\right))\ell\,\epsilon$ with $\epsilon\, m = t_\eta(A)$.

Proof: To see this, we note that:

$$\{H_A \leq \epsilon\, m\} = \bigcup_{i=1}^{\ell}\{H_{A_i} \leq \epsilon\, m\}.$$

The union bound gives:

$$\mathbb{P}_\pi\left(H_A \leq \epsilon\, m\right) \leq \sum_{i=1}^{\ell} \mathbb{P}_\pi\left(H_{A_i} \leq \epsilon\, m\right) \leq \left(1 + O\left(\delta\right)\right)\ell\epsilon.$$

A lower bound can be obtained via the Bonferroni inequality:

$$\begin{aligned}\mathbb{P}_\pi\left(H_A \leq \epsilon\, m\right) \geq \sum_{i=1}^{\ell} \mathbb{P}_\pi\left(H_{A_i} \leq \epsilon\, m\right) - \sum_{1\leq i<j\leq \ell} \mathbb{P}_\pi\left(H_{A_i} \leq \epsilon\, m, H_{A_j} \leq \epsilon\, m\right) \\ = \left(1 + O\left(\delta + \xi\right)\right)\ell\epsilon\end{aligned}$$

using the definition of ξ. □

We now need to show that the assumptions of Theorem 3.1 are applicable to H_A, with the value of η in Claim 3.3 replacing ϵ. Indeed, we note that $\eta \geq \epsilon/2\ell$ and $\mathrm{T}^Q_{\mathrm{mix}}(\delta\eta) \leq \mathrm{T}^Q_{\mathrm{mix}}(\delta\epsilon/2)$ (here we are assuming that $\xi + \delta$ is small enough). Therefore,

$$\mathbb{P}_\pi\left(H_A \leq \mathrm{T}_{\mathrm{mix}}(\delta\eta)\right) \leq \sum_{i=1}^{\ell} \mathbb{P}_\pi\left(H_{A_i} \leq \mathrm{T}_{\mathrm{mix}}(\delta\epsilon/2)\right) \leq \ell\delta\epsilon/2 \leq \delta\eta.$$

We deduce from Theorem 3.1 that for any $\lambda \in M_1(\mathbf{V})$,

$$\mathrm{Law}_\lambda\left(H_A\right) = \mathsf{Exp}(t_\eta(A)/\eta, O\left(\eta\right) + r_\lambda, O\left(\delta + \xi\right)).$$

To finish the proof, we note that $\eta = O\left(\ell\epsilon\right)$,

$$t_\eta(A)/\eta = \epsilon\, m/(1 + O\left(\delta + \xi\right))\,\ell\epsilon = (1 + O\left(\delta + \xi\right))\,\frac{m}{\ell}$$

and apply Proposition 3.1. $\square$

4 Meeting times of multiple random walks

We now put our two exponential approximation results to use, showing that the meeting times we are interested in are well-approximated by exponential random variables. Much of the work needed for this is contained in technical estimates whose proofs can be safely skipped in a first reading.

4.1 Basic definitions

For the remainder of this section, $\mathbf{V}$ is a finite set and Q is the generator of a mixing Markov chain over $\mathbf{V}$ with mixing times $\mathrm{T}^Q_{\mathrm{mix}}(\cdot)$ and stationary measure π. For each $k \in \mathbb{N}\backslash\{0, 1\}$ we will also consider the Markov chains $Q^{(k)}$ over $\mathbf{V}^k$ that correspond to k independent realizations of Q from prescribed initial states, as defined in Section 2.3.4. We will also follow the notation from that section.

For $k = 2$, we define the first meeting time:

$$M \equiv \inf\{t \geq 0\ :\ X_t(1) = X_t(2)\} \tag{13}$$

and the parameters:

$$\mathsf{m}(Q) \equiv \mathbb{E}_{\pi^{\otimes 2}}\left[M\right], \tag{14}$$

$$\rho(Q) \equiv \frac{\mathrm{T}^Q_{\mathrm{mix}}}{\mathsf{m}(Q)}. \tag{15}$$

We also define an extra prameter $\mathsf{err}(Q)$ which will appear as an error term in our Lemmas. Ths parameter $\mathsf{err}(Q)$ is defined as

$$\mathsf{err}(Q) = c_0\sqrt{\rho(Q)\ln(1/\rho(Q))} \text{ if } Q \text{ is reversible and transitive.} \tag{16}$$

For other Q, we define it as:

$$\mathsf{err}(Q) = c_1 \sqrt{(1 + q_{\max} \mathrm{T}^Q_{\mathrm{mix}})\, \pi_{\max}\, \ln\left(\frac{1}{(1 + q_{\max} \mathrm{T}^Q_{\mathrm{mix}})\, \pi_{\max}}\right)}. \tag{17}$$

The numbers $c_0, c_1 > 0$ are universal constants that we do not specify explicitly. We choose them so as to satisfy Propositions 4.1, 4.4 and 4.5 below.

We now take $k > 2$ and consider the process $Q^{(k)}$, with trajectories

$$(X^{(k)}_t = (X_t(1), X_t(2), \ldots, X_t(k)))_{t \geq 0}$$

corresponding to k independent realizations of Q (cf. Section 2.3.4). This has stationary distribution $\pi^{\otimes k}$.

We write $M^{(k)}$ for the first meeting time among these random walks:

$$M^{(k)} \equiv \inf \{t \geq 0 \,:\, \exists 1 \leq i < j \leq k,\, X_t(i) = X_t(j)\}\,. \tag{18}$$

One may note that

$$M^{(k)} = \min_{\{i,j\} \in \binom{[k]}{2}} M_{i,j}$$

where for $1 \leq i < j \leq k$:

$$M_{i,j} = M_{j,i} \equiv \inf\{t \geq 0 \,:\, X_t(i) = X_t(j)\}. \tag{19}$$

is distributed as M for a realization of $Q^{(2)}$ starting from $(X_0(i), X_0(j))$.

4.2 Technical estimates for reversible and transitive chains

In this subsection we collect the estimates that we will use in the case of chains that are reversible and transitive.

Proposition 4.1 *Assume Q is reversible and transitive and define* $\mathsf{err}(Q)$ *accordingly. If* $\mathsf{err}(Q) \leq 1/4$*, then:*

$$\mathbb{P}_{\pi^{\otimes 2}}\left(M \leq \mathrm{T}^Q_{\mathrm{mix}}(\mathsf{err}(Q)^2)\right) \leq \mathsf{err}(Q)^2.$$

Remark 5 *The proof is entirely general, but we will only use this estimate in the transitive/reversible case.*

Proof: We will prove a result in contrapositive form: if $0 < \beta < 1/4$ is such that:

$$\mathbb{P}_{\pi^{\otimes 2}}\left(M \leq \mathrm{T}^Q_{\mathrm{mix}}(\beta)\right) > \beta,$$

then $\beta > c_0^2 \rho(Q) \ln(1/\rho(Q))$ for some universal $c_0 > 0$.

Notice that for any $x^{(2)} \in \mathbf{V}^2$,

$$\begin{aligned}\mathbb{P}_{x^{(2)}}\left(M > \mathrm{T}^{Q}_{\mathrm{mix}}(\beta/4) + \mathrm{T}^{Q}_{\mathrm{mix}}(\beta)\right) &\le \mathbb{P}_{x^{(2)}}\left(M \circ \Theta_{\mathrm{T}^{Q}_{\mathrm{mix}}(\beta/4)} > \mathrm{T}^{Q}_{\mathrm{mix}}(\beta)\right)\\ &\le \mathbb{P}_{\pi^{\otimes 2}}\left(M > \mathrm{T}^{Q}_{\mathrm{mix}}(\beta)\right) + \beta/2\\ &< (1-\beta/2),\end{aligned}$$

where the last inequality follows from the fact that $\mathrm{T}^{Q}_{\mathrm{mix}}(\beta/4)$ is an upper bound for the $\beta/2$-mixing time of $Q^{(2)}$ (cf. Remark **??**). A standard argument using the Markov property implies that for any $k \in \mathbb{N}$,

$$\mathbb{P}_{x^{(2)}}\left(M > k\,(\mathrm{T}^{Q}_{\mathrm{mix}}(\beta/4) + \mathrm{T}^{Q}_{\mathrm{mix}}(\beta))\right) < (1-\beta/2)^k,$$

so that:

$$\mathsf{m}(Q) = \mathbb{E}_{\pi^{\otimes 2}}[M] \le C\,\frac{\mathrm{T}^{Q}_{\mathrm{mix}}(\beta) + \mathrm{T}^{Q}_{\mathrm{mix}}(\beta/4)}{\beta}.$$

Since $\mathrm{T}^{Q}_{\mathrm{mix}}(\alpha) \le C\ln(1/\alpha)\,\mathrm{T}^{Q}_{\mathrm{mix}}$, we deduce that:

$$\frac{\beta}{c\,\ln(1/\beta)} > \rho(Q),$$

which implies the desired result. $\Box$

We now prove an estimate on correlations.

Proposition 4.2 *Assume Q is transitive. Then for all $t, s \ge 0$ and $\{i,j\}, \{\ell, r\} \subset \mathbf{V}$ with $\{i,j\} \ne \{r,\ell\}$,*

$$\mathbb{P}_{\pi^{\otimes k}}\left(M_{i,j} \le t,\ M_{\ell,r} \le s\right) \le 2\mathbb{P}_{\pi^{\otimes 2}}\left(M \le s\right)\mathbb{P}_{\pi^{\otimes 2}}\left(M \le t\right).$$

Proof: If $\{i,j\} \cap \{\ell,r\} = \emptyset$, the events $\{M_{i,j} \le t\}$ and $\{M_{\ell,r} \le t\}$ are independent. Since the laws of both $M_{i,j}$ and $M_{\ell,r}$ under $\pi^{\otimes k}$ are equal to the law of M under $\pi^{\otimes 2}$, we obtain:

$$\mathbb{P}_{\pi^{\otimes k}}\left(M_{i,j} \le t,\ M_{\ell,r} \le s\right) \le \mathbb{P}_{\pi^{\otimes 2}}\left(M \le s\right)\mathbb{P}_{\pi^{\otimes 2}}\left(M \le t\right)$$

in this case. Assume now $\{i,j\} \cap \{\ell,r\}$ has one element. Without loss of generality we may assume $k = 3$, $\{i,j\} = \{1,2\}$ and $\{\ell,r\} = \{1,3\}$. We have:

$$\begin{aligned}\mathbb{P}_{\pi^{\otimes 3}}\left(M_{1,2} \le t,\ M_{1,3} \le s\right) \le \mathbb{P}_{\pi^{\otimes 3}}&\left(M_{1,2} \le t, M_{1,3} \circ \Theta_{M_{1,2}} \le s\right)\\ &+ \mathbb{P}_{\pi^{\otimes 3}}\left(M_{1,3} \le s, M_{1,2} \circ \Theta_{M_{1,3}} \le t\right). \quad (20)\end{aligned}$$

Consider the first term in the RHS. By the Markov property:

$$\mathbb{P}_{\pi^{\otimes 3}}\left(M_{1,2} \le t, M_{1,3} \circ \Theta_{M_{1,2}} \le s\right) = \mathbb{P}_{\pi^{\otimes 2}}\left(M \le t\right)\,\mathbb{P}_{\lambda^{(2)}}\left(M \le s\right)$$

where $\lambda^{(2)}$ is the law of $X_{M_{1,2}}(1), X_{M_{1,2}}(3)$ conditionally on $M_{1,2} \leq t$. Since $(X_t(3))_t$ is stationary and independent from this event, $\lambda^{(2)} = \lambda \otimes \pi$ for some $\lambda \in M_1(\mathbf{V})$ which is the law of $X_{M_{1,2}}(1)$ under $\mathbb{P}_{\pi^{\otimes 2}}$. The transitivity of Q (which implies that π is uniform) implies that $\lambda = \pi$ and therefore:

$$\mathbb{P}_{\pi^{\otimes 3}}\left(M_{1,2} \leq t, M_{1,3} \circ \Theta_{M_{1,2}} \leq s\right) = \mathbb{P}_{\pi^{\otimes 2}}\left(M \leq t\right) \mathbb{P}_{\pi^{\otimes 2}}\left(M \leq s\right).$$

The same bound can be shown for the other term in the RHS of (20), and this implies the Proposition. $\square$

4.3 Technical estimates for the general case

We will need the following general result:

Proposition 4.3 *For any $\lambda \in M_1(\mathbf{V})$ and $T \geq 0$*

$$\mathbb{P}_{\lambda\otimes\pi}\left(M \leq T\right) \leq (1 + 2T\, q_{\max})\, \pi_{\max}.$$

Proof: Let $(X_t)_t$ be a single realization of Q. One may imagine that the trajectory of $(X_t)_{t\geq 0}$ is sampled as follows. First, let $\mathcal{P}$ be a Poisson process with intensity $q_{\max}$ independent from the initial state X_0. At each time $t \in \mathcal{P}$, one updates the value of X_t as follows: if $X_s = x$ for s immediately before t, one sets:

$$X_t = y \text{ with probability } \frac{q(x,y)}{q_{\max}}\ (y \in \mathbf{V}\backslash\{x\})$$

and $X_t = x$ with the remaining probability. This implies that, at the points of the Poisson process, X_t is updated as in the discrete-time Markov chain with matrix $P = (I + Q/q_{\max})$, and it is easy to see that π is stationary for this chain.

Now let $X_t(1), X_t(2)$ be independent trajectories of Q, with $X_t(1)$ started from λ and $X_t(2)$ started from the stationary distribution π. We will imagine that each $X_t(i)$ has its own Poisson process $\mathcal{P}(i)$ and was generated in the way described above. It then follows that:

$$\begin{aligned}\mathbb{P}_{\lambda\otimes\pi}\left(M \leq T \mid \mathcal{P}(1), \mathcal{P}(2)\right) \leq\ & \mathbb{P}_{\lambda\otimes\pi}\left(X_0(1) = X_0(2)\right) \\ & + \sum_{t \in (\mathcal{P}(1)\cup\mathcal{P}(2))\cap(0,T]} \mathbb{P}_{\lambda\otimes\pi}\left(X_t(1) = X_t(2) \mid \mathcal{P}(1), \mathcal{P}(2)\right)\end{aligned}$$

since the processes can only change values at the times of the two Poisson processes. At time 0 we have:

$$\mathbb{P}_{\lambda\otimes\pi}\left(X_0(1) = X_0(2)\right) = \sum_{x\in\mathbf{V}} \lambda(x)\pi(x) \leq \sum_{x\in\mathbf{V}} \lambda(x)\pi_{\max} \leq \pi_{\max}.$$

For $t \in \mathcal{P}(1) \cup \mathcal{P}(2)$, the law of $X_t(1), X_t(2)$ equals:

$$(\lambda\, P^{k_1}) \otimes (\pi\, P^{k_2})$$

where $k_i = |\mathcal{P}(i) \cap (0,t]|$ $(i = 1,2)$. Crucially, π is stationary for P, hence $\pi\, P^{k_2} = \pi$ and we obtain:

$$\mathbb{P}_{\lambda\otimes\pi}\left(X_t(1) = X_t(2) \mid \mathcal{P}(1), \mathcal{P}(2)\right) = \sum_{x\in\mathbf{V}} (\lambda\, P^{k_1})(x)\pi(x) \leq \pi_{\max}$$

as for $t = 0$. We deduce:

$$\mathbb{P}_{\lambda\otimes\pi}\left(M \leq T \mid \mathcal{P}(1), \mathcal{P}(2)\right) \leq (1 + |(\mathcal{P}(1) \cup \mathcal{P}(2)) \cap (0,T]|)\, \pi_{\max}.$$

The Proposition follows from taking expectations on both sides and noticing that:

$$\mathbb{E}\left[|(\mathcal{P}(1) \cup \mathcal{P}(2)) \cap (0,T]|\right] = 2T\, q_{\max}.$$

$\Box$

We now prove an estimate corresponding to Proposition 4.1 in this general setting.

Proposition 4.4 *Assume* $\mathsf{err}(Q)$ *is as defined in (17). Then:*

$$\mathbb{P}_{\pi^{\otimes 2}}\left(M \leq \mathrm{T}^Q_{\mathrm{mix}}(\mathsf{err}(Q)^2)\right) \leq \mathsf{err}(Q)^2.$$

Proof: The previous proposition implies:

$$\begin{aligned}\mathbb{P}_{\pi^{\otimes 2}}\left(M \leq \mathrm{T}^Q_{\mathrm{mix}}(\mathsf{err}(Q)^2)\right) &\leq (1 + 2\mathrm{T}^Q_{\mathrm{mix}}(\mathsf{err}(Q))\, q_{\max})\, \pi_{\max} \\ &\leq C\,(1 + 2\mathrm{T}^Q_{\mathrm{mix}}\, q_{\max})\, \pi_{\max}\, \ln(1/\mathsf{err}(Q)).\end{aligned}$$

This is $\leq \mathsf{err}(Q)^2$ by definition of this quantity, if we choose c_1 in (17) to be large enough. $\Box$

We now prove an estimate on correlations that is similar to Proposition 4.2, but with an extra term.

Proposition 4.5 *For any mixing Markov chain Q, if one defines* $\mathsf{err}(Q)$ *as in (17), we have the following inequality for $k \geq 3$ and all distinct pairs $\{i,j\}, \{\ell, r\} \in \binom{[k]}{2}$:*

$$\mathbb{P}_{\pi^{\otimes k}}\left(M_{i,j} \leq t, M_{\ell,r} \leq s\right) \quad \leq \quad 2\mathbb{P}_{\pi^{\otimes 2}}\left(M \leq t\right)\mathbb{P}_{\pi^{\otimes 2}}\left(M \leq s\right) \quad + \quad O\left(\mathsf{err}(Q)^2\right).$$

Proof: The case $\{i,j\} \cap \{\ell, r\} = \emptyset$ follows as in the proof of Proposition 4.2. In case $\{i,j\} \cap \{\ell, r\}$ has one element, we may again assume that $i = \ell = 1$, $j = 2$ and $r = 3$. Equation (20) still applies, so we proceed to bound:

$$\mathbb{P}_{\pi^{\otimes 3}}\left(M_{1,2} \leq t, M_{1,3} \circ \Theta_{M_{1,2}} \leq s\right),$$

which is upper bounded by:

$$\begin{aligned}
\mathbb{P}_{\pi^{\otimes 3}}\left(M_{1,2} \le t, M_{1,3} \circ \Theta_{M_{1,2}} \le s\right) & \\
\le \mathbb{P}_{\pi^{\otimes 3}} & \left(M_{1,2} \le t, M_{1,3} \circ \Theta_{M_{1,2}} \le \mathrm{T}^{Q}_{\mathrm{mix}}(\eta)\right) \\
& + \mathbb{P}_{\pi^{\otimes 3}}\left(M_{1,2} \le t, M_{1,3} \circ \Theta_{M_{1,2}+\mathrm{T}^{Q}_{\mathrm{mix}}(\eta)} \le s\right) = (I) + (II) \quad (21)
\end{aligned}$$

for some $\eta \in (0, 1/4)$ to be chosen later.

Term (I) is equal to:

$$\mathbb{P}_{\pi^{\otimes 3}}\left(M_{1,2} \le t\right) \mathbb{P}_{\lambda^{(2)}}\left(M \le \mathrm{T}^{Q}_{\mathrm{mix}}(\eta)\right)$$

where $\lambda^{(2)}$ is the law of $(X_{M_{1,2}}(2), X_{M_{1,2}}(3))$ conditionally on $\{M_{1,2} \le t\}$. As in the previous proof, $(X_t(3))_t$ is stationary and independent from the conditioning, hence $\lambda^{(2)} = \lambda \otimes \pi$ for some $\lambda \in M_1(\mathbf{V})$. We use Proposition 4.3 to deduce:

$$(I) \le \mathbb{P}_{\pi^{\otimes 3}}\left(M_{1,2} \le t\right) O\left((1 + \mathrm{T}^{Q}_{\mathrm{mix}}(\eta)\, q_{\max})\, \pi_{\max}\right)$$

The analysis of term (II) is simpler: we have

$$(II) = \mathbb{P}_{\pi^{\otimes 3}}\left(M_{1,2} \le t\right) \mathbb{P}_{\lambda_* \otimes \pi}\left(M \le s\right)$$

for some $\lambda_* \in M_1(\mathbf{V})$ which is the law of $X_{M_{1,2}+\mathrm{T}^{Q}_{\mathrm{mix}}(\eta)}$ conditionally on $\{M_{1,2} \le t\}$. The time shift by $\mathrm{T}^{Q}_{\mathrm{mix}}(\eta)$ implies that λ_* is η-close to stationary, hence:

$$(II) \le \mathbb{P}_{\pi^{\otimes 3}}\left(M_{1,2} \le t\right) (\eta + \mathbb{P}_{\pi^{\otimes 2}}\left(M \le s\right)).$$

We deduce from (21) that:

$$\begin{aligned}
\mathbb{P}_{\pi^{\otimes 3}}\left(M_{1,2} \le t, M_{1,3} \circ \Theta_{M_{1,2}} \le s\right) & \le (I) + (II) \\
& \le \mathbb{P}_{\pi^{\otimes 1}}\left(M \le t\right) \mathbb{P}_{\pi^{\otimes 2}}\left(M \le s\right) \\
& \quad + \mathbb{P}_{\pi^{\otimes 2}}\left(M \le t\right) O\left(\eta + (1 + \mathrm{T}^{Q}_{\mathrm{mix}}(\eta)\, q_{\max})\, \pi_{\max}\right). \quad (22)
\end{aligned}$$

Recall $\mathrm{T}^{Q}_{\mathrm{mix}}(\eta) \le C\, \mathrm{T}^{Q}_{\mathrm{mix}} \ln(1/\eta)$ for some universal $C > 0$. If

$$\eta_0 \equiv (1 + q_{\max}\, \mathrm{T}^{Q}_{\mathrm{mix}})\, \pi_{\max} \le 1/2$$

we may take $\eta = \eta_0$ to obtain:

$$\begin{aligned}
\mathbb{P}_{\pi^{\otimes 3}}\left(M_{1,2} \le t, M_{1,3} \circ \Theta_{M_{1,2}} \le s\right) & \\
\le \mathbb{P}_{\pi^{\otimes 1}}\left(M \le t\right) & \mathbb{P}_{\pi^{\otimes 2}}\left(M \le s\right) \\
+ \mathbb{P}_{\pi^{\otimes 2}}\left(M \le t\right) O & \left((1 + \mathrm{T}^{Q}_{\mathrm{mix}}\, q_{\max})\, \pi_{\max} \ln\left(\frac{1}{(1 + \mathrm{T}^{Q}_{\mathrm{mix}}\, q_{\max})\, \pi_{\max}}\right)\right). \quad (23)
\end{aligned}$$

The case of $\eta_0 \ge 1/2$ is covered "automatically" by the big-oh notation.

An analogous bound can be obtained with the roles of $(t, 2)$ and $(s, 3)$ reversed. Plugging these into (20) gives the desired bound. $\square$

4.4 Exponential approximation for a pair of particles

We now come back to the setting of Section 4.1 and show M is approximatelly exponentially distributed.

Lemma 4.1 *Define* $\mathsf{err}(Q)$ *as in (16) (if* Q *is reversible and transitive) or as in (17) (if not) and assume that* $\mathsf{err}(Q) < 1/10$*. Then* $\forall \lambda^{(2)} \in M_1(\mathbf{V}^{(2)})$*:*

$$\mathrm{Law}_{\lambda^{(2)}}(M) = \mathsf{Exp}(\mathsf{m}(Q), O(\mathsf{err}(Q)) + 2\mathbb{P}_{\lambda^{(2)}}\left(M \leq \mathrm{T}_{\mathrm{mix}}^{Q}(\mathsf{err}(Q)^2)\right), O(\mathsf{err}(Q))).$$

Proof: This is a direct application of Theorem 3.1 to the hitting time of the diagonal set

$$\Delta \equiv \{(x,x) \,:\, x \in \mathbf{V}\} \subset \mathbf{V}^2$$

by the chain with generator $Q^{(2)}$ defined in Section 2 and with $\epsilon = \mathsf{err}(Q)$, $\delta = 2\mathsf{err}(Q)$. All we need to show is that:

$$\mathbb{P}_{\pi^{\otimes 2}}\left(M \leq \mathrm{T}_{\mathrm{mix}}^{Q^{(2)}}(\delta\epsilon)\right) \leq \epsilon\delta$$

where $\mathrm{T}_{\mathrm{mix}}^{Q^{(2)}}(\cdot)$ denotes the mixing times of $Q^{(2)}$. This inequality follows from

$$\mathrm{T}_{\mathrm{mix}}^{Q^{(2)}}(2\mathsf{err}(Q)^2) \leq \mathrm{T}_{\mathrm{mix}}^{Q}(\mathsf{err}(Q)^2) \text{ (Fact 1)}$$

and

$$\mathbb{P}_{\pi^{\otimes 2}}\left(M \leq \mathrm{T}_{\mathrm{mix}}^{Q}(\mathsf{err}(Q)^2)\right) \leq \mathsf{err}(Q)^2 < \delta\epsilon,$$

which follows from Proposition 4.1 in the reversible/transitive case and Proposition 4.4 in the general case. □

4.5 Exponential approximation for many random walkers

We now consider the more complex problem of bounding the meeting times among $k \geq 2$ particles. We take the notation in Section 4.1 for granted.

Lemma 4.2 *Let* $\ell = \binom{k}{2} > 0$ *and assume that the quantity* $\mathsf{err}(Q)$ *defined in (16) (if* Q *is reversible and transitive) or as in (17) (if not) satisfies* $\ell \leq 1/10\ell$*. Then for all* $\lambda^{(k)} \in M_1(\mathbf{V}^k)$*,*

$$\mathrm{Law}_{\lambda^{(k)}}\left(M^{(k)}\right) = \mathsf{Exp}\left(\frac{\mathsf{m}(Q)}{\ell}, O\left(k^2\mathsf{err}(Q)\right) + 2r_{\lambda^{(k)}}), O\left(k^2\mathsf{err}(Q)\right)\right)$$

where $r_{\lambda^{(k)}} = \mathbb{P}_{\lambda^{(k)}}\left(M^{(k)} \leq \mathrm{T}_{\mathrm{mix}}^{Q}(\mathsf{err}(Q)^2)\right)$ *and*

$$\xi_0 \equiv k^2\,\mathsf{err}(Q).$$

Proof: $M^{(k)}$ is the hitting time of a union of ℓ sets.

$$\Delta^{(k)} \equiv \bigcup_{\{i,j\}\in\binom{k}{2}} \Delta_{\{i,j\}} \text{ where } \Delta_{\{i,j\}} \equiv \{x^{(k)} \in \mathbf{V}^k \,:\, x^{(k)}(i) = x^{(k)}(j)\}.$$

We will apply Theorem 3.2 applied to the product chain $Q^{(k)}$ to show that this hitting time is approximately exponential. We set $\delta = 2k\mathsf{err}(Q)$, $\epsilon = \mathsf{err}(Q)$ and verify the conditions of the Theorem:

- $0 < \delta < 1/5$, $0 < \epsilon < 1/5\ell$: These conditions follow from $\mathsf{err}(Q) < 1/10\ell$.
- $\mathbb{P}_{\pi^{\otimes k}}\left(M_{i,j} \leq \mathrm{T}_{\mathrm{mix}}^{Q^{(k)}}(\delta\epsilon/2)\right) \leq \delta\epsilon/3$. To prove this we simply observe that:

  $$\mathrm{T}_{\mathrm{mix}}^{Q^{(k)}}(\delta\epsilon/2) \leq \mathrm{T}_{\mathrm{mix}}^{Q}(\mathsf{err}(Q)^2) \text{ (Fact 1 and defn. of } \epsilon, \delta)$$

  and that:
  $$\mathbb{P}_{\pi^{\otimes 2}}\left(M \leq \mathrm{T}_{\mathrm{mix}}^{Q}(\mathsf{err}(Q)^2)\right) \leq \mathsf{err}(Q)^2 \ = \frac{\delta\epsilon}{2k} < \frac{\delta\epsilon}{3}$$
  by Proposition 4.1 (in the reversible/transitive case) or by Proposition 4.4 (in general).
- $\mathbb{E}_{\pi^{\otimes k}}\left[H_{\Delta_{\{i,j\}}}\right] = \mathsf{m}(Q)$ is the same for all $\{i,j\} \in \binom{[k]}{2}$: this is obvious.

The Lemma will then follow once we show that the ξ quantity in Theorem 3.2 satisfies:

$$\xi \equiv \frac{1}{\ell\epsilon} \sum_{\{i,j\}\neq\{\ell,r\} \text{ in } \binom{[k]}{2}} \mathbb{P}_{\pi^{\otimes k}}\left(M_{\{i,j\}} \leq \epsilon\mathsf{m}(Q),\, M_{\{\ell,r\}} \leq \epsilon\mathsf{m}(Q)\right) = O\left(\xi_0\right).$$

To start, we go back to Claim 3.2 in the proof of Theorem 3.2 and observe that, whenever the assumptions of that Theorem hold,

$$\mathbb{P}_{\pi^{\otimes k}}\left(M_{\{i,j\}} \leq \epsilon\mathsf{m}(Q)\right) = O\left(\epsilon\right). \tag{24}$$

Now note that Proposition 4.2 (in the reversible/transitive case) and Proposition 4.5 (in the general case) imply that each term in the sum defining ξ is $O\left(\mathsf{err}(Q)^2\right)$. We deduce:

$$\xi_0 \leq \frac{O\left(\epsilon^2\right)\, \binom{\ell}{2}}{\ell\epsilon} \leq O\left(\ell\epsilon\right) = O\left(k^2\,\mathsf{err}(Q)\right).$$

□

5 Coalescing random walks: basics

In this section we formally define the coalesing random walks process. We then show that, if the initial number of particles is not large, mean field behaviour follows from the exponential approximation of meeting times.

5.1 Definitions

Fix a Markov chain Q on a finite state space $\mathbf{V}$. Given a number $k \in [|\mathbf{V}|]\backslash\{1\}$ and an initial state $x^{(k)} \in \mathbf{V}^k$, consider a realization of $Q(k)$

$$(X^{(k)})_{t\geq 0} \equiv (X_t(1), \ldots, X_t(k))_{t\geq 0}.$$

We build the coalescing random walks process from $X^{(k)}$ by defining the trajectories of the k walkers one by one. We first set:

$$\overline{X}_t(1) = X_t(1),\, t \geq 0.$$

Given $j \in [k]\backslash\{1\}$, assume that $\overline{X}_t(i)$ has bee defined for all $1 \leq i < j$ and $t \geq 0$. We let T_j be the first time $t \geq 0$ at which $X_t(j) = \overline{X}_t(I_j)$ for some $1 \leq I_j < i$, and then set:

$$\overline{X}_t(j) \equiv \begin{cases} X_t(j), & t < T_j; \\ \overline{X}_t(I_j), & t \geq T_j. \end{cases}$$

Intuitively, this says that as soon as j encounters a walker with lower index, it starts moving along with it. The process:

$$(\overline{X}^{(k)}_t)_{t\geq 0} \equiv (\overline{X}_t(j))_{t\geq 0}$$

is what we call the *coalescing random walks process* based on Q, with initial state $x^{(k)}$.

Remark 6 *For any $j \geq 3$, there might be more than one index $i < j$ such that $\overline{X}_{T_j}(i) = X_{T_j}(j)$. However, it is easy to see that all such i will have the same trajectory after time T_j, because they must have met by that time. This implies that there is no ambiguity in the definition of $\overline{X}_t(j)$ for any j.*

We also define:

$$\mathsf{C}_i \equiv \inf\{t \geq 0 \,:\, |\{\overline{X}_t(j) \,:\, j \in [k]\}| \leq i\}$$

and $\mathsf{C} \equiv \mathsf{C}_1$. The fact that we are working in continous time implies:

Proposition 5.1 (Proof omitted) *Assume that the initial state $x^{(k)} = (x(1), x(2), \ldots, x(k))$ is such that $x(i) \neq x(j)$ for all $1 \leq i < j \leq k$. Then $\mathsf{C}_k = 0 < \mathsf{C}_{k-1} < \mathsf{C}_{k-2} < \cdots < \mathsf{C}_1$ almost surely.*

It is sometimes useful to view the coalescing random walks process as a process with *killings*. Define a random $2^{[k]}$-valued process $(A_t)_{t\geq 0}$ as follows.

- $1 \in A_t$ for all t;
- proceeding recursively, for each $j \in [k]\backslash\{1\}$, we have $j \in A_t$ if and only if $\tau_j < t$, where τ_j is the first time t at which $X_t(i) = X_t(j)$ for some $i < j$ with $i \in A_t$.

Intuitively, A_t is the set of all walkers that are "alive" at time $t \geq 0$, and a walker dies at the first time it meets an alive walker with smaller index. One may check that coalescing random walks is equivalent to the killed process in the following sense.

Proposition 5.2 (Proof omitted) *We have $\tau_j = T_j$ for all $j \in [k]\backslash\{1\}$. Moreover, for all $t \geq 0$ we have:*

$$\{X_t(j) \,:\, j \in A_t\} = \{\overline{X}_t(j) \,:\, j \in [k]\}.$$

Finally, for all $i \in [k-1]$

$$\mathsf{C}_i = \inf\{t \geq 0 \,:\, |A_t| \leq i\}.$$

Recall that $M_{i,j}$ is the meeting time between walkers i and j (cf. (19)). We have the following simple proposition.

Proposition 5.3 (Proof omitted) *Assume that the initial state*

$$x^{(k)} = (x(1), x(2), \ldots, x(k))$$

is such that $x(i) = x(j)$ for all $1 \leq i < j \leq k$. Then for each $1 \leq p \leq k-1$,

$$\mathsf{C}_p - \mathsf{C}_{p+1} = \min_{\{i,j\} \subset A_{\mathsf{C}_{p+1}}} M_{i,j} \circ \Theta_{\mathsf{C}_{p+1}}.$$

Moreover, each time C_p equals $M_{i,j}$ for some $\{i,j\} \in \binom{[k]}{2}$.

5.2 Mean-field behavior for moderately large k

We now prove a mean-field-like result for an initial number of particles k that is not too large, assuming that meeting times of up to k walkers sasitfy our exponential appoximation property

Lemma 5.1 *There exists a universal constant $C > 0$ such that the following holds. Assume that Q, $\mathsf{err}(Q)$ and k satisfy the assumptions of Lemma 4.2. Let $x^{(k)} \in \mathbf{V}^k$. Then for all $p \in [k-1]$*

$$d_W\left(\mathrm{Law}_{x^{(k)}}\left(\frac{\mathsf{C}_p}{\mathsf{m}(Q)}\right), \mathrm{Law}\left(\sum_{i=p+1}^{k} \mathsf{Z}_i\right)\right) = \frac{O\left(k^2 \mathsf{err}(Q)\right) + 12\,\eta(x^{(k)})}{p},$$

where

$$\eta(x^{(k)}) = \mathbb{P}_{x^{(k)}}\left(M^{(k)} \leq \mathrm{T}_{\mathrm{mix}}(\mathsf{err}(Q)^2)\right) + \mathbb{P}_{x^{(k)}}\left(\exists \{i,j\}, \{\ell, r\} \in \binom{[k]}{2},\ \{\ell,r\} \neq \{i,j\} \text{ but } M_{i,j} \circ \Theta_{M_{\ell,r}} \leq \mathrm{T}_{\mathrm{mix}}(\mathsf{err}(Q)^2)\right)$$

and the Z_i are the random variables described in (1).

Proof: Write $x^{(k)} = (x(1), \ldots, x(k))$. We will prove the similar bound:

$$\text{“}\forall 1 \leq i < j \leq k\, x(i) \neq x(j)\text{”} \Rightarrow$$
$$d_W\left(\mathrm{Law}_{x^{(k)}}\left(\frac{\mathsf{C}_p}{\mathsf{m}(Q)}\right), \mathrm{Law}\left(\sum_{i=p+1}^{k} \mathsf{Z}_i\right)\right) = \frac{O\left(k^2 \mathsf{err}(Q)\right) + 4\,\eta(x^{(k)})}{p}, \quad (25)$$

To see how this implies the general result, consider some $x^{(k)}$ such that some of its coordinates are equal, so that in particular $\eta(x^{(k)}) \geq 1$. One still has the trivial bound:

$$d_W\left(\mathrm{Law}_{x^{(k)}}\left(\frac{\mathsf{C}_p}{\mathsf{m}(Q)}\right), \mathrm{Law}\left(\sum_{i=p+1}^{k} \mathsf{Z}_i\right)\right) \leq \mathbb{E}_{x^{(k)}}\left[\frac{\mathsf{C}_p}{\mathsf{m}(Q)}\right] + \mathbb{E}\left[\sum_{i=p+1}^{k} \mathsf{Z}_i\right].$$

The second term in the RHS is $\leq 2/p$. For the first term, let $y^{(j)} = (y(1), \ldots, y(j)) \in \mathbf{V}^j$ have distinct coordinates with:

$$\{y(1), \ldots, y(j)\} = \{x(1), \ldots, x(k)\}.$$

Then clearly:

$$\mathbb{E}_{x^{(k)}}\left[\frac{\mathsf{C}_p}{\mathsf{m}(Q)}\right] = \mathbb{E}_{y^{(j)}}\left[\frac{\mathsf{C}_p}{\mathsf{m}(Q)}\right].$$

If $p \geq j$ the RHS is 0. If not, it can be upper bounded using the Lemma:

$$\mathbb{E}_{y^{(j)}}\left[\frac{\mathsf{C}_p}{\mathsf{m}(Q)}\right] \leq \mathbb{E}\left[\sum_{i=p+1}^{k} \mathsf{Z}_i\right] + d_W\left(\mathrm{Law}_{y^{(j)}}\left(\frac{\mathsf{C}_p}{\mathsf{m}(Q)}\right), \mathrm{Law}\left(\sum_{i=p+1}^{j} \mathsf{Z}_i\right)\right)$$
$$\leq \frac{2 + 4\eta(y^{(j)}) + O\left(k^2 \mathsf{err}(Q)\right)}{p}.$$

Since $\eta(x^{(k)}) \geq 1 \geq \eta(y^{(j)})/2$ in this case, we obtain

$$d_W\left(\mathrm{Law}_{x^{(k)}}\left(\frac{\mathsf{C}_p}{\mathsf{m}(Q)}\right), \mathrm{Law}\left(\sum_{i=p+1}^{k} \mathsf{Z}_i\right)\right) \leq \frac{12\,\eta(x^{(k)}) + O\left(k^2 \mathsf{err}(Q)\right)}{p}$$

for such $x^{(k)}$ with repetitions, which gives the Lemma in general.

We prove (25) by reverse induction on p. The case $p = k-1$ is trivial: C_{k-1} is simply $M^{(k)}$ and $\eta(x^{(k)})$ is an upper bound for $r_{\delta_{x^{(k)}}}$, so we may apply Lemma 4.2 to deduce the desired bound.

For the inductive step, consider $p_0 < k-1$ and assume the result is true for all $p_0 < p \leq k-1$. We will use the easily proven fact that C_{p_0+1} is a stopping time for the process

$(X_t^{(k)})_{t\geq 0}$ process. Consider the corresponding σ-field $\mathcal{F}_{\mathsf{C}_{p_0+1}}$. We will apply Lemma 2.3 with:

$$\begin{aligned} Z_1 &= \sum_{i=p_0+2}^{k} \mathsf{Z}_i, \\ Z_2 &= \mathsf{Z}_{p_0+1}, \\ W_1 &= \frac{\mathsf{C}_{p_0+1}}{\mathsf{m}(Q)}, \\ W_2 &= \mathsf{C}_{p_0} - \mathsf{C}_{p_0+1} = \frac{\min_{\{i,j\}\subset A_{\mathsf{C}_{p+1}}} M_{i,j}\circ\Theta_{\mathsf{C}_{p+1}}}{\mathsf{m}(Q)}, \\ \mathcal{G} &= \mathcal{F}_{\mathsf{C}_{p_0+1}}. \end{aligned}$$

(We used Proposition 5.3 to obtain the second expression for W_2 above.) Applying the Lemma in conjunction with the induction hypothesis gives:

$$\begin{aligned} d_W\left(\mathrm{Law}_{x^{(k)}}\left(\frac{\mathsf{C}_{p_0}}{\mathsf{m}(Q)}\right), \sum_{i=p_0+1}\mathsf{Z}_i\right) &\leq \frac{O\left(k^2\mathsf{err}(Q)\right)+4\eta(x^{(k)})}{p_0+1} \\ &+ \mathbb{E}_{x^{(k)}}\left[d_W\left(\mathrm{Law}_{x^{(k)}}\left(\frac{\min_{\{i,j\}\subset A_{\mathsf{C}_{p_0+1}}} M_{i,j}\circ\Theta_{\mathsf{C}_{p_0+1}}}{\mathsf{m}(Q)}\middle|\mathcal{F}_{C_{p_0+1}}\right), \mathsf{Z}_{p_0+1}\right)\right]. \quad (26) \end{aligned}$$

Note that $A_{\mathsf{C}_{p_0+1}}$ is $\mathcal{F}_{\mathsf{C}_{p_0+1}}$-measurable. Therefore the strong Markov property for $Q^{(k)}$ implies:

$$\mathrm{Law}_{x^{(k)}}\left(\frac{\min_{\{i,j\}\subset A_{\mathsf{C}_{p_0+1}}} M_{i,j}\circ\Theta_{\mathsf{C}_{p_0+1}}}{\mathsf{m}(Q)}\middle|\mathcal{F}_{C_{p_0+1}}\right) = \mathrm{Law}_{X^{(k)}_{\mathsf{C}_{p_0+1}}}\left(\frac{\min_{\{i,j\}\subset A_{\mathsf{C}_{p_0+1}}} M_{i,j}}{\mathsf{m}(Q)}\right).$$

Now define $Y^{(p_0+1)}$ as the vectors whose coordinates are the p_0+1 distinct points $X_{\mathsf{C}_{p_0+1}}(i)$ with $i\in A_{p_0+1}$ (the order of the coordinates does not matter). Clearly,

$$\mathrm{Law}_{X^{(k)}_{\mathsf{C}_{p_0+1}}}\left(\frac{\min_{\{i,j\}\subset A_{\mathsf{C}_{p_0+1}}} M_{i,j}}{\mathsf{m}(Q)}\right) = \mathrm{Law}_{Y^{(p_0+1)}}\left(\frac{M^{(p_0+1)}}{\mathsf{m}(Q)}\right). \quad (27)$$

By assumption, this last law satisfies:

$$\mathrm{Law}_{Y^{(p_0+1)}}\left(\frac{M^{(p_0+1)}}{\mathsf{m}(Q)}\right) = \mathsf{Exp}\left(\frac{1}{\binom{p_0+1}{2}}, O\left(k^2\mathsf{err}(Q)\right)+2\,r_{\delta_{Y^{(p_0+1)}}}, O\left(k^2\,\mathsf{err}(Q)\right)\right),$$

and Lemma 3.1 gives:

$$d_W\left(\mathrm{Law}_{Y^{(p_0+1)}}\left(\frac{M^{(p_0+1)}}{\mathsf{m}(Q)}\right), \mathsf{Z}_{p_0+1}\right) \leq \frac{O\left(k^2\mathsf{err}(Q)\right)+4r_{\delta_{Y^{(p_0+1)}}}}{p_0(p_0+1)}.$$

Using the definition of $r_{\delta_{Y^{(p_0+1)}}}$, we obtain from (26) the following inequality:

$$d_W\left(\mathrm{Law}_{x^{(k)}}\left(\frac{\mathsf{C}_{p_0}}{\mathsf{m}(Q)}\right), \sum_{i=p_0+1} \mathsf{Z}_i\right) \leq \frac{O\left(k^2\mathsf{err}(Q)\right) + 4\eta(x^{(k)})}{p_0+1}$$
$$+ \frac{O\left(k^2\mathsf{err}(Q)\right) + 4\mathbb{E}_{x^{(k)}}\left[\mathbb{P}_{Y^{(p_0+1)}}\left(M^{(p_0+1)} \leq \mathrm{T}^Q_{\mathrm{mix}}(\mathsf{err}(Q)^2)\right)\right])}{p_0(p_0+1)}. \quad (28)$$

To finish, we need to show that the expected value in the RHS is $\leq \eta(x^{(k)})$. For this we recall (27) note that

$$\mathbb{P}_{Y^{(p_0+1)}}\left(M^{(p_0+1)} \leq \mathrm{T}^Q_{\mathrm{mix}}(\mathsf{err}(Q)^2)\right)$$
$$= \mathbb{P}_{X^{(k)}_{\mathsf{C}_{p_0+1}}}\left(\min_{\{i,j\}\subset A_{\mathsf{C}_{p_0+1}}} M_{i,j}\circ\Theta_{\mathsf{C}_{p_0+1}} \leq \mathrm{T}^Q_{\mathrm{mix}}(\mathsf{err}(Q)^2)\right)$$
$$(\text{use Prop. 5.3 + strong Markov}) = \mathbb{P}_{x^{(k)}}\left(\mathsf{C}_{p_0} - \mathsf{C}_{p_0+1} \leq \mathrm{T}^Q_{\mathrm{mix}}(\mathsf{err}(Q)^2) \mid \mathcal{F}_{\mathsf{C}_{p_0+1}}\right).$$

Averaging gives:

$$\mathbb{E}_{x^{(k)}}\left[\mathbb{P}_{Y^{(p_0+1)}}\left(M^{(p_0+1)} \leq \mathrm{T}^Q_{\mathrm{mix}}(\mathsf{err}(Q)^2)\right)\right] = \mathbb{P}_{x^{(k)}}\left(\mathsf{C}_{p_0} - \mathsf{C}_{p_0+1} \leq \mathrm{T}^Q_{\mathrm{mix}}(\mathsf{err}(Q)^2)\right)$$

and Proposition 5.3 implies:

$$\mathbb{P}_{x^{(k)}}\left(\mathsf{C}_{p_0} - \mathsf{C}_{p_0+1} \leq \mathrm{T}^Q_{\mathrm{mix}}(\mathsf{err}(Q)^2)\right) \leq \mathbb{P}_{x^{(k)}}\left(\bigcup_{\{i,j\}\neq\{\ell,r\}} \{M_{i,j}\circ\Theta_{M_{\ell,r}} \leq \mathrm{T}^Q_{\mathrm{mix}}(\mathsf{err}(Q)^2)\}\right).$$

SInce the RHS is $\leq \eta(x^{(k)})$, we are done. $\square$

6 Proofs of the main theorems

6.1 The full coalescence time in the transitive case

In this section we prove Theorem 1.1.
Proof: [of Theorem 1.1] Recall that $\mathsf{C} = \mathsf{C}_1$ by definition. Lemma 4.2 gives the following bound for any $k \leq \sqrt{1/4\mathsf{err}(Q)} \wedge |\mathbf{V}|$ and $x^{(k)} \in \mathbf{V}^k$

$$d_W\left(\mathrm{Law}_{x^{(k)}}\left(\frac{\mathsf{C}}{\mathsf{m}(Q)}\right), \sum_{i=2}^{k} \mathsf{Z}_i\right) \leq 12\,\eta(x^{(k)}) + O\left(k^2\mathsf{err}(Q)\right).$$

Notice that:

$$d_W\left(\sum_{i=2}^{k} \mathsf{Z}_i, \sum_{i=2}^{+\infty} \mathsf{Z}_i\right) \leq \mathbb{E}\left[\sum_{j\geq k+1} \mathsf{Z}_j\right] = \frac{2}{k+1},$$

hence:

$$d_W\left(\mathrm{Law}_{x^{(k)}}\left(\frac{\mathsf{C}_1}{\mathsf{m}(Q)}\right), \sum_{i=2}^{+\infty}\mathsf{Z}_i\right) = 12\,\eta(x^{(k)}) + O\left(k^2\mathsf{err}(Q) + \frac{1}{k}\right).$$

Convexity of d_W implies:

Proposition 6.1 *Under the assumptions of Theorem 1.1, if* $\mathsf{err}(Q) \le 1/4$*, the following holds for* $k \le \sqrt{1/4\mathsf{err}(Q)} \wedge |\mathbf{V}|$ *and* $\lambda^{(k)} \in M_1(\mathbf{V}^k)$*:*

$$d_W\left(\mathrm{Law}_{\lambda^{(k)}}\left(\frac{\mathsf{C}_1}{\mathsf{m}(Q)}\right), \sum_{i=2}^{+\infty}\mathsf{Z}_i\right) \le 12\int \eta(x^{(k)})\,d\lambda^{(k)}(x^{(k)}) + O\left(k^2\mathsf{err}(Q) + \frac{1}{k}\right). \quad (29)$$

Notice that our control of C_1 gets worse as k increases, and we cannot use the above bound to approximate the law of C_1 started with one particle at each vertex of $\mathbf{V}$. What we use instead is a truncation argument combined with the Sandwich Lemma for d_W (Lemma 2.2 above). For this we need to find two random variables

$$\mathsf{C}_- \preceq_d \mathsf{C}_1 \preceq_d \mathsf{C}_+$$

such that both $\mathsf{C}_-/\mathsf{m}(Q)$ and $\mathsf{C}_+/\mathsf{m}(Q)$ are close to $\sum_{i=2}^{+\infty}\mathsf{Z}_i$. More specifically, we will show that:

$$d_W\left(\frac{\mathsf{C}_\pm}{\mathsf{m}(Q)}, \sum_{i\ge 2}\mathsf{Z}_i\right) = O\left(k^2\,\mathsf{err}(Q) + k^4\mathsf{err}(Q)^2 + \frac{1}{k+1} + \rho(Q)\,\ln(1/\rho(Q))\right). \quad (30)$$

Before we continue, let us show how this last bound implies our result. Lemma 2.2 then gives:

$$d_W\left(\frac{\mathsf{C}_1}{\mathsf{m}(Q)}, \sum_{i\ge 2}\mathsf{Z}_i\right) = O\left(k^2\,\mathsf{err}(Q) + k^4\mathsf{err}(Q)^2 + \frac{1}{k+1} + \rho(Q)\,\ln(1/\rho(Q))\right).$$

Since $\rho(Q)\ln(1/\rho(Q)) = O\left(\mathsf{err}(Q)\right)$, we may choose $k = (\mathsf{err}(Q))^{-1/3}$ (which works for $\mathsf{err}(Q)$ sufficiently small) to obtain:

$$d_W\left(\frac{\mathsf{C}_1}{\mathsf{m}(Q)}, \sum_{i\ge 2}\mathsf{Z}_i\right) = O\left(\mathsf{err}(Q)^{1/3}\right),$$

and this is precisely the bound we seek because $\mathsf{err}(Q) = O\left(\sqrt{\rho(Q)\ln(1/\rho(Q))}\right)$. We now construct $\mathsf{C}_-, \mathsf{C}_+$ and prove that they have the required properties.

Construction of C_-: pick $x(1), \ldots, x(k) \in \mathbf{V}$ from distribution π, independently and with replacement. Let C_- denote the full coalescence time for k walkers started from these positions. This might be degenerate: there might be more than one walker starting from some element of $\mathbf{V}$, but this only means those particles will coalesce instantly.

Clearly, $\mathsf{C}_- \preceq_d \mathsf{C}_1$. Moreover,

$$\mathrm{Law}\left(\frac{\mathsf{C}_-}{\mathsf{m}(Q)}\right) = \mathrm{Law}_{\pi^{\otimes k}}\left(\frac{\mathsf{C}_1}{\mathsf{m}(Q)}\right).$$

Therefore by Proposition 6.1,

$$\begin{aligned} d_W\left(\mathrm{Law}\left(\frac{\mathsf{C}_-}{\mathsf{m}(Q)}\right), \sum_{i=2}^{+\infty} \mathsf{Z}_i\right) &= d_W\left(\mathrm{Law}_{\pi^{\otimes k}}\left(\frac{\mathsf{C}_1}{\mathsf{m}(Q)}\right), \sum_{i=2}^{+\infty} \mathsf{Z}_i\right) \\ &= O\left(\int \eta(x^{(k)})\, d\pi^{\otimes k}(x^{(k)}) + k^2\, \mathsf{err}(Q) + \frac{1}{k}\right). \end{aligned}$$

Notice that the integral in the RHS is at most:

$$\begin{aligned} &\sum_{\{i,j\}\in\binom{[k]}{2}} \mathbb{P}_{\pi^{\otimes k}}\left(M_{i,j} \leq \mathrm{T}^Q_{\mathrm{mix}}(\mathsf{err}(Q)^2)\right) \\ &\quad + \sum_{\{i,j\},\{\ell,r\}\in\binom{[k]}{2}\,:\,\{i,j\}\neq\{\ell,r\}} \mathbb{P}_{\pi^{\otimes k}}\left(M_{i,j}\circ\Theta_{M_{\ell,r}} \leq \mathrm{T}^Q_{\mathrm{mix}}(\mathsf{err}(Q)^2)\right) = O\left(k^4 \mathsf{err}(Q)^2\right). \end{aligned} \tag{31}$$

as can be deduced from the proofs of Propositions 4.2 and 4.1. We conclude that:

$$d_W\left(\mathrm{Law}\left(\frac{\mathsf{C}_-}{\mathsf{m}(Q)}\right), \mathrm{Law}\left(\sum_{i\geq 2} \mathsf{Z}_i\right)\right) = O\left(k^4 \mathsf{err}(Q)^2 + k^2\, \mathsf{err}(Q) + \frac{1}{k+1}\right). \tag{32}$$

Construction of C_+: we will use the following simple stochastic domination result, which we describe in the language of the process with killings. Let $\tau \leq \sigma$ be stopping times for the $X^{(k)}$ process. If all killings are supressed between time τ and σ, the resulting full coalescence time C_+ stochastically dominates C_1. We will use this result, whose proof we omit, with the following choice of τ and σ:

$$\tau = \mathsf{C}_k \text{ and } \sigma = \mathsf{C}_k + \mathrm{T}^Q_{\mathrm{mix}}(\mathsf{err}(Q)^2).$$

Lemma 2.1 implies:

$$d_W\left(\frac{\mathsf{C}_+}{\mathsf{m}(Q)}, \frac{\mathsf{C}_1\circ\Theta_\sigma}{\mathsf{m}(Q)}\right) \leq \frac{\mathbb{E}\left[\sigma\right]}{\mathsf{m}(Q)} = \frac{\mathbb{E}\left[\mathsf{C}_k\right]}{\mathsf{m}(Q)} + \frac{\mathrm{T}^Q_{\mathrm{mix}}(\mathsf{err}(Q)^2)}{\mathsf{m}(Q)}.$$

Since Q is transitive, $\mathsf{m}(Q)$, can be bounded from below in terms of the maximal hitting time in Q [4, Chapter 14]. Theorem 1.2 in [11] implies:

$$\mathbb{E}\left[\mathsf{C}_k\right] \leq \frac{C\,\mathsf{m}(Q)}{k} + C\,\mathrm{T}^Q_{\mathrm{mix}}$$

for some universal $C > 0$. Recalling the definition of $\rho(Q)$ in (15), we obtain:

$$\frac{\mathbb{E}[\mathsf{C}_k]}{\mathsf{m}(Q)} = O\left(\frac{1}{k} + \rho(Q)\right).$$

Moreover, we alse have

$$\mathrm{T}^Q_{\mathrm{mix}}(\mathsf{err}(Q)^2) = O\left(\ln(1/\mathsf{err}(Q))\,\mathrm{T}^Q_{\mathrm{mix}}\right) = O\left(\mathrm{T}^Q_{\mathrm{mix}}\ln(1/\rho(Q))\right),$$

hence:

$$d_W\left(\frac{\mathsf{C}_+}{\mathsf{m}(Q)}, \frac{\mathsf{C}_1\circ\Theta_\sigma}{\mathsf{m}(Q)}\right) = O\left(\frac{1}{k} + \rho(Q)\,\ln(1/\rho(Q))\right).$$

This shows:

$$d_W\left(\frac{\mathsf{C}_+}{\mathsf{m}(Q)}, \sum_{i=2}^k \mathsf{Z}_i\right) = O\left(\frac{1}{k} + \rho(Q)\,\ln(1/\rho(Q))\right) + d_W\left(\frac{\mathsf{C}_1\circ\Theta_\sigma}{\mathsf{m}(Q)}, \sum_{i=2}^k \mathsf{Z}_i\right).$$

Now consider the time $\mathsf{C}_1\circ\Theta_\sigma$. Since all killings were supressed between times $\tau = \mathsf{C}_k$ and $\sigma = \mathsf{C}_k + \mathrm{T}^Q_{\mathrm{mix}}(\mathsf{err}(Q)^2)$, there are k alive particles at time σ_-. Letting $\lambda^{(k)}$ denote their law, we have:

$$\mathrm{Law}\left(\frac{\mathsf{C}_1\circ\Theta_\sigma}{\mathsf{m}(Q)}\right) = \mathrm{Law}_{\lambda^{(k)}}\left(\frac{\mathsf{C}_1}{\mathsf{m}(Q)}\right)$$

and Proposition 6.1 implies:

$$d_W\left(\frac{\mathsf{C}_1\circ\Theta_\sigma}{\mathsf{m}(Q)}, \sum_{i=2}^k \mathsf{Z}_i\right) = O\left(\int \eta(x^{(k)})\,d\lambda^{(k)}(x^{(k)}) + k^2\mathsf{err}(Q) + \frac{1}{k+1}\right)$$

Now observe that

$$\mathrm{T}^Q_{\mathrm{mix}}(\mathsf{err}(Q)^2) \le \mathrm{T}^{Q^{(k)}}_{\mathrm{mix}}(k\mathsf{err}(Q)^2) \text{ (cf. Fact 1)},$$

hence the law of the k particles at time $\mathsf{C}_k + \mathrm{T}^Q_{\mathrm{mix}}(\mathsf{err}(Q)^2)$ is $k\mathsf{err}(Q)^2$-close to stationary, irrespective of their states at time C_k. We deduce that $\lambda^{(k)}$ is $k\mathsf{err}(Q)^2$-close to stationary, and deduce:

$$d_W\left(\frac{\mathsf{C}_1\circ\Theta_\sigma}{\mathsf{m}(Q)}, \sum_{i=2}^k \mathsf{Z}_i\right) = O\left(\int \eta(x^{(k)})\,d\pi^{\otimes k}(x^{(k)}) + k^2\mathsf{err}(Q) + k\mathsf{err}(Q)^2 + \frac{1}{k+1}\right).$$

The integral in the RHS was estimated in (31), and we deduce:

$$d_W\left(\frac{\mathsf{C}_1\circ\Theta_\sigma}{\mathsf{m}(Q)}, \sum_{i=2}^k \mathsf{Z}_i\right) = O\left(k^2\,\mathsf{err}(Q) + k^4\mathsf{err}(Q)^2 + \frac{1}{k}\right)$$

and we deduce:

$$d_W\left(\mathsf{C}_+, \sum_{i=2}^k \mathsf{Z}_i\right) = O\left(k^2\,\mathsf{err}(Q) + k^4\mathsf{err}(Q)^2 + \frac{1}{k+1} + \rho(Q)\,\ln(1/\rho(Q))\right).$$

□

6.2 The general setting

Proof: The proof is essentially the same as in the reversible/transitive case, with the definition of $\mathsf{err}(Q)$ given in (17). The only point at which we need a different strategy is that, in the analysis of C_+, we need to bound $\mathbb{E}[\mathsf{C}_k]$ by different means.

We use the following simple strategy: $\mathsf{C}_k \geq t$ if and only if there exist distinct $y(1), \dots, y(k) \in \mathbf{V}$ such that there is *no coalescence* among the walkers started from these vertices. The probability of this "no coalescence event" for a given choice of $y(i)$'s is $\mathbb{P}_{y^{(k)}}\left(M^{(k)} \geq t\right)$ for $y^{(k)} = (y(1), \dots, y(k))$. Therefore,

$$\mathbb{P}(\mathsf{C}_k \geq t) \leq \left(\sum_{y^{(k)} \in \mathbf{V}^k} \mathbb{P}_{y^{(k)}}\left(M^{(k)} \geq t\right)\right) \wedge 1.$$

By Lemma 4.2, each term in the RHS satisfies:

$$\mathbb{P}_{y^{(k)}}\left(M^{(k)} \geq t\right) \leq C\, e^{-\frac{t\binom{k}{2}}{(1+O(\mathsf{err}(Q)))\,\mathsf{m}(Q)}}$$

for some universal $C > 0$. SInce there are $\leq |\mathbf{V}|^k$ terms in the sum, we have:

$$\mathbb{P}(\mathsf{C}_k \geq t) \leq \left(C\,|\mathbf{V}|^k\, e^{-\frac{t\binom{k}{2}}{(1+O(\mathsf{err}(Q)))\,\mathsf{m}(Q)}}\right) \wedge 1.$$

Integrating the RHS gives:

$$\frac{\mathbb{E}[\mathsf{C}_k]}{\mathsf{m}(Q)} \leq C\,\frac{\ln|\mathbf{V}|}{k}$$

for a potentially different, but still universal C. One may check that this new bound changes (30) into:

$$d_W\left(\frac{\mathsf{C}_\pm}{\mathsf{m}(Q)}, \sum_{i\geq 2} \mathsf{Z}_i\right) = O\left(k^2\,\mathsf{err}(Q) + k^4\mathsf{err}(Q)^2 + \frac{\ln|\mathbf{V}|}{k+1}\right). \tag{33}$$

We deduce as in the previous proof that:

$$d_W\left(\frac{\mathsf{C}_1}{\mathsf{m}(Q)}, \sum_{i\geq 2} \mathsf{Z}_i\right) = O\left(k^2\,\mathsf{err}(Q) + k^4\mathsf{err}(Q)^2 + \frac{\ln|\mathbf{V}|}{k+1}\right).$$

We choose $k = O\left((\ln|\mathbf{V}|/\mathsf{err}(Q))^{1/3}\right)$ to finish the proof, at least if this is smaller than $1/5\sqrt{\mathsf{err}(Q)}$. But the bound in the Theorem is trivial if that is not the case, so we are done. $\square$

7 Final remarks

- Recall that the original Open Problem 12 in [4, Chapter 14] was to prove an analogue of Theorem 1.1 with the relaxation time replacing $\mathrm{T}_{\mathrm{mix}}^{Q}$. Is this true? In fact, we do not even know of an explicit vertex-transitive graph with relaxation time much smaller than the mixing times.
- Cooper et al. [5] consider many other processes besides coalescing random walks. It is not hard to modify our analysis to study those processes over more general graphs, at least when the initial number of random walks is not too large (this restriction is also present in [5]).
- Our Theorems 3.1 and 3.2 can be used to study other problems related to hitting times. We are in the process of writing a preprint where we prove that the fluctuations of cover times over certain vertex-transitive graphs follow the Gumbel law. This is another instance of "mean-field behavior", as the Gumbel law is the limiting distribution over large complete graphs as well.

A Proofs of techncal results on L_1 Wasserstein distance

A.1 Proof of Sandwich Lemma (Lemma 2.2)

Notice that for all $t \in \mathbb{R}$,

$$\mathbb{P}\left(Z_- \geq t\right) \leq \mathbb{P}\left(Z \geq t\right) \leq \mathbb{P}\left(Z_+ \geq t\right).$$

By convexity, this implies:

$$|\mathbb{P}\left(Z \geq t\right) - \mathbb{P}\left(W \geq t\right)| \leq |\mathbb{P}\left(Z_- \geq t\right) - \mathbb{P}\left(W \geq t\right)| + |\mathbb{P}\left(Z_+ \geq t\right) - \mathbb{P}\left(W \geq t\right)|.$$

Integrate both sides to obtain the result.

A.2 Proof of Conditional Lemma (Lemma 2.3)

First notice that the sigma field $\sigma(W_1)$ generated by W_1 is contained in $\mathcal{G}$. This implies that for all $t \in \mathbb{R}$:

$$\begin{aligned}\mathbb{E}\left[|\mathbb{P}\left(W_2 \geq t \mid \mathcal{G}\right) - \mathbb{P}\left(Z_2 \geq t\right)|\right] &= \mathbb{E}\left[\mathbb{E}\left[|\mathbb{P}\left(W_2 \geq t \mid \mathcal{G}\right) - \mathbb{P}\left(Z_2 \geq t\right)\right] \mid \sigma(W_1)\right] \\ &\geq \mathbb{E}\left[|\mathbb{P}\left(W_2 \geq t \mid \sigma(W_1)\right) - \mathbb{P}\left(Z_2 \geq t\right)|\right].\end{aligned}$$

Integrating both sides in t and applying Fubini-Tonelli gives:

$$\mathbb{E}\left[d_W(\mathrm{Law}\left(W_2 \mid \mathcal{G}\right), \mathrm{Law}\left(Z_2\right))\right] \geq \mathbb{E}\left[d_W(\mathrm{Law}\left(W_2 \mid \sigma(W_1)\right), \mathrm{Law}\left(Z_2\right))\right].$$

Therefore it suffices to prove the Theorem in the case $\mathcal{G} = \sigma(W_1)$. For simplicity, we will assume that (Z_1, Z_2, W_1, W_2) are all defined in the same probability space, with (Z_1, Z_2) independent from (W_1, W_2). Let $f : \mathbb{R} \to \mathbb{R}$ be 1-Lipschitz. We have:

$$\mathbb{E}\left[f(W_1 + W_2) \mid W_1 = w_1\right] = \int f(w_1 + w_2)\, \mathbb{P}\left(W_2 \in dw_2 \mid W_1 = w_1\right).$$

By the duality version of d_W, we have:

$$\begin{aligned} &\int f(w_1 + w_2)\, \mathbb{P}\left(W_2 \in dw_2 \mid W_1 = w_1\right) \\ &\qquad\qquad \leq \int f(w_1 + z_2)\, \mathbb{P}\left(Z_2 \in dz_2\right) + d_W(\mathrm{Law}\left(W_2 \mid W_1 = w_1\right), \mathrm{Law}\left(Z_2\right)) \end{aligned}$$

Integrating over $W_1 = w_1$ and using the fact that Z_2 is independent from W_1, we obtain:

$$\mathbb{E}\left[f(W_1 + W_2)\right] \leq \mathbb{E}\left[f(W_1 + Z_2)\right] + d_W(\mathrm{Law}\left(W_2 \mid W_1\right), \mathrm{Law}\left(Z_2\right)).$$

But we also have:

$$\mathbb{E}\left[f(W_1 + Z_2) \mid Z_2 = z_2\right] = \mathbb{E}\left[f(W_1 + z_2)\right] \leq \mathbb{E}\left[f(Z_1 + z_2)\right] + d_W(W_1, Z_1),$$

and the independence of Z_1, Z_2 implies:

$$\mathbb{E}\left[f(W_1 + Z_2)\right] \leq \mathbb{E}\left[f(Z_1 + Z_2)\right] + d_W(W_1, Z_1).$$

We conclude:

$$\mathbb{E}\left[f(W_1 + W_2)\right] \leq \mathbb{E}\left[f(Z_1 + Z_2)\right] + d_W(W_1, Z_1) + d_W(\mathrm{Law}\left(W_2 \mid W_1\right), \mathrm{Law}\left(Z_2\right)).$$

Since f is an arbitrary 1-Lipschitz function, we are done.

References

[1] David Aldous. Markov chains with almost exponential hitting times. *Stochastic Processes and Applications*, 13(3):305–310, 1982.

[2] David Aldous. Mixing times and hitting times. AIM talk slides available from `http://www.stat.berkeley.edu/~aldous/Talks/`., 2010.

[3] David Aldous and Mark Brown. Inequalities for rare events in time-reversible markov chains. i. In *Stochastic Inequalities (Seattle, WA, 1991)*, pages 1–16. IMS Lecture Notes Series, 1992.

[4] David Aldous and James Allen Fill. Reversible Markov Chains and Random Walks on Graphs. Book draft, `http://www.stat.berkeley.edu/~aldous/RWG/book.html`.

[5] Colin Cooper, Alan Frieze, and Tomasz Radszik. Multiple random walks in random regular graphs. *SIAM Journal of Discrete Mathematics*, 23(4):1738–1761, 2009/10.

[6] J. Theodore Cox. Coalescing random walks and voter model consensus times on the torus in Z^d. *Annals of Probability*, 17:1333–1366, 1989.

[7] Rick Durrett. *Random Graph Dynamics*. Cambridge University Press, 2006.

[8] Rick Durrett. Some features of the spread of epidemics and information on a random graph. *Proceedings of the National Academy of Sciences*, 107(10):4491–4498, 2010.

[9] Jure Leskovec, Kevin Lang, Anirban Dasgupta, and Michael Mahoney. Community structure in large networks: Natural cluster sizes and the absence of large well-defined clusters. *Internet Mathematics*, 6(1):29–123, 2009.

[10] Thomas M. Liggett. *Interacting Particle Systems*, volume 276 of *Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]*. Springer-Verlag, 1985.

[11] Roberto Imbuzeiro Oliveira. On the coalescence time of reversible random walks. To appear in Transactions of the American Mathematical Society, 2010.