

Markov Chain Robustness

David Zuckerman   

University of Texas at Austin, TX, USA

Abstract

When a Markov chain models nature or social interactions, it is likely not followed exactly, but only approximately. We therefore introduce several notions of robustness for a Markov chain P . Our standard adversary can dynamically change transition probabilities of P by $1 \pm \varepsilon$, and our strong adversary can completely control each transition independently with probability ε , as in a model by Azar, Broder, Karlin, Linial, and Philips [4]. These adversaries are equivalent up to constant factors if the degrees are constant. Our adversarial chains need not converge.

We define and prove various robustness properties of a reversible chain P , i.e., a random walk on a connected undirected graph G . Let d be the maximum degree, Δ the diameter, π the stationary distribution, and t_{mix} the mixing time.

1. We define a natural analogue $\pi^+(S)$ that upper bounds limiting frequencies in a set S in the adversarial chain. We show that if $\varepsilon = O(1/\sqrt{dt_{\text{up}}})$, where t_{up} is a variant of the mixing time, then $\pi^+(S) = O(\pi(S)^{1-\alpha})$ for any $\alpha > 0$.
2. We define the mixing time robustness as the largest ε such that the approximate mixing time increases by only a constant factor, and prove that it is $\Omega(1/\sqrt{dt_{\text{mix}}})$.
3. We define the hitting time robustness as the largest ε such that the maximum hitting time increases by only a constant factor, and show that it is $\Omega(1/t_{\text{mix}})$. For trees, we show it is $\Omega(1/\Delta)$.
4. We define the cover time robustness as the largest ε such that the cover time increases by only a constant factor. We show that in most graphs it's at least the hitting time robustness.
5. We characterize the mixing, hitting, and cover time robustnesses for constant-degree regular expander graphs up to constant factors. They are $\Theta(1)$, $\Theta(1/\log n)$, and $\Theta(1/\log n)$, respectively.

2012 ACM Subject Classification Theory of computation \rightarrow Random walks and Markov chains

Keywords and phrases Markov chain, random walk, mixing time, hitting time, cover time, robustness, expander graph

Digital Object Identifier 10.4230/LIPIcs.ITCS.2026.118

Funding Supported in part by NSF Grant CCF-2312573 and a Simons Investigator Award (#409864).

Acknowledgements We thank Kuikui Liu and Salil Vadhan for helpful discussions, the anonymous referees for very useful suggestions, and Zhiyang Xun and Dean Doron for helpful comments.

1 Introduction

Markov chains are used to model many scenarios, and there is a huge theory about them. In many situations, such as a Markov chain modeling nature or social behavior, it is likely that the Markov chain is not followed exactly, but only approximately. We therefore study how behavior in Markov chains is affected by small perturbations.

For example, Cannon, Daymude, Randall, and Richa [6] showed that a collection of computational elements can compress – gather tightly as in a sphere – by asynchronously each executing a simple program. (See also [17].) They analyzed this by modeling it as a Markov chain. Such compression and other self-organizing behavior also occurs in nature. In nature, one may expect that the algorithm is not followed exactly but only approximately. Would the same compressive behavior arise in this case? For a Markov chain modeling social behavior, it's even more compelling that the chain wouldn't be followed exactly but could have adversarial influences.



© David Zuckerman;

licensed under Creative Commons License CC-BY 4.0

17th Innovations in Theoretical Computer Science Conference (ITCS 2026).

Editor: Shubhangi Saraf; Article No. 118; pp. 118:1–118:22

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

Another situation where errors arise is in simulating huge physical systems. Such a system may exactly follow a Markov chain; however, researchers simulate such a system by coarsening the exact chain into a discrete approximation. This coarsening necessarily leads to error. How badly does this error affect the simulation? In this case, Kania, Aristoff, and Zuckerman [15] recently developed a randomized algorithm to remove the error, but analyzing the error remains interesting.

There has been some work exploring the power of an adversary on certain Markov chains. Chin, Moitra, Mossel, and Sandon [7] studied the power of an adversary in Glauber dynamics. Glauber dynamics is a type of random walk on Σ^V , where V is a set of nodes in an underlying graph and Σ is a finite alphabet. They proposed an adversary that controls any changes to a subset $A \subseteq V$ with $|A| \leq \varepsilon|V|$. They studied how such an adversary can affect global features of the system, such as the approximate mixing time.

Doron, Moshkovitz, Oh, and Zuckerman [9, 10] studied the power of an extremely strong adversary for random walks on lossless expanders. They showed that even in the presence of such a strong adversary, a random walk will approximately mix in some sense. However, their methods don't appear to extend beyond lossless expanders.

We study a model that can be applied to any random walk. Related questions have been studied before, and we adopt some of those frameworks, but we study different questions.

To study errors in Markov chain behavior, Hartfiel [11, 12] proposed a natural model of *Markov set-chains* defined as follows. Assume the states are $[n] = \{1, 2, \dots, n\}$, and let \mathcal{M}_n denote the set of $n \times n$ stochastic matrices. The Markov set-chain is a stochastic process $\{X_t\}$ defined by a set $\mathcal{P} \subseteq \mathcal{M}_n$.¹ At each time step t , a transition matrix $P_t \in \mathcal{P}$ is chosen arbitrarily. Then P_t determines the transitions from X_t to X_{t+1} : $\Pr[X_{t+1} = j \mid X_t = i] = P_t(i, j)$. Thus a Markov chain is a special case of a Markov set-chain where the set \mathcal{P} consists of one matrix, the transition matrix of the Markov chain. In general, the transition matrices in \mathcal{P} may have different stationary distributions.

Azar, Broder, Karlin, Linial, and Phillips [4] studied a model similar in spirit. More recently, researchers have studied a model very similar to Markov set-chains under the name of dynamic graphs or evolving graphs, e.g., [3, 22].

1.1 The adversary

There appear to be three reasonable choices for the power of an adversary choosing P_t .²

1. Oblivious: P_t is chosen independently of $\{X_s\}$;
2. Markovian: P_t may depend on X_t but not X_s for $s < t$;
3. Non-Markovian: P_t may depend on X_s for all $s \leq t$.

Note that an oblivious adversary corresponds to a non-homogeneous (or time-dependent) Markov chain. Oblivious and Markovian adversaries are equivalent if \mathcal{P} is closed under “row replacement”: replacing a row in any $P \in \mathcal{P}$ by the corresponding row in another $P' \in \mathcal{P}$. The equivalence follows because the adversary can choose the i th row assuming the Markov chain is in state i . The Markovian and non-Markovian adversaries are equivalent if an adversary is trying to affect the hitting time, but not equivalent for cover time. Hartfiel, [4], and recent work on dynamic graphs implicitly use an oblivious adversary (and don't discuss the others).

¹ Hartfiel's definition has \mathcal{P} being compact, but let's ignore that for now.

² Note that we give the adversary more power by allowing P_t to depend on t ; otherwise we could get three more choices for the adversary's power.

We focus on the most powerful, non-Markovian adversary, although we make some observations about the other adversaries.

Can we understand a Markov set chains (MSC) in terms of a related Markov chain? Although reversibility is not required for all of our theorems, we usually assume the related chain is reversible, i.e., a random walk on an undirected graph. We further assume the graph is connected; otherwise we could restrict attention to the appropriate connected component. The most important probability distribution related to a Markov chain is the stationary distribution π . The most important quantities of interest in Markov chains are the mixing time t_{mix} , the time it takes to approach π ; the max hitting time t_{hit} , the maximum expected time to go from one vertex to another; and the cover time t_{cov} , the expected time to visit all nodes from a worst-case starting node.

First, we discuss the stationary distribution. Suppose all chains in an MSC have the same stationary distribution π . Will this MSC converge to π ? It depends on the adversary. For an oblivious adversary, the MSC will converge to π – the usual proof works. For a Markovian adversary, the MSC need not converge, even if all underlying chains are reversible; see Section 3.1.

There are two natural special cases of MSCs that we consider. First, an *interval chain* is an interval of matrices $\mathcal{P} = [P^-, P^+] \cap \mathcal{M}_n$, where $[P^-, P^+]$ denotes the set of matrices P such that for all states i, j , we have $P^-(i, j) \leq P(i, j) \leq P^+(i, j)$. For interval chains, oblivious and Markovian adversaries are equivalent, since \mathcal{P} is closed under row replacement.

For our standard adversary, we consider interval chains of the following form. For P a Markov chain³ and $\varepsilon > 0$, we consider the interval chain $[P^- = (1-\varepsilon)P, P^+ = (1+\varepsilon)P]$. When convenient, we consider the essentially equivalent interval chain $[P^- = e^{-\varepsilon}P, P^+ = e^{\varepsilon}P]$.

Second, [4] introduced a model of biased random walks where a controller biases the random walk. That is, given a Markov chain P and a parameter $\varepsilon > 0$, a controller chooses an arbitrary stochastic matrix B with $\text{supp}(B) \subseteq \text{supp}(P)$ and the biased random walk follows $(1-\varepsilon)P + \varepsilon B$. That is, with probability ε the controller biases the walk, but has to respect the graph structure. Haslegrave, Sauerwald, and Sylvester [13] studied this in the time dependent setting. The focus of these papers is on a friendly controller who tries to improve the situation from the simple random walk. Our focus is on an adversarial controller who seeks to make matters worse. This will be our strong adversary.

In other words, in both cases, we have that $(1-\varepsilon)P$ is left alone. A strong adversary may distribute the remaining ε probability arbitrarily on its support. A standard adversary is limited to change each probability by a bounded relative amount. However, a strong adversary with parameter ε is at most as powerful as a $(d_{\text{max}}\varepsilon)$ -standard adversary, where d_{max} denotes the maximum degree. Since we are thinking of d_{max} as small or constant, we don't dwell on the difference between the adversaries.

One might suggest an even stronger adversary that can take an arbitrary ε probability and distribute it arbitrarily. In other words, it can choose P_t such that every row of P and P_t differs by at most ε in total variation distance. This adversary generalizes the adversary of [7]. However, if $\varepsilon \geq 1/\text{deg}(v)$ for some node v , then such an adversary can force the MSC to avoid hitting v . This renders some of our questions impossible, so we don't consider this adversary.

Main question. For how large an ε does the adversarial chain have similar behavior as the original chain?

³ We identify a Markov chain with its transition matrix.

In order to make sense of this question, we need to define suitable notions for MSCs. We study several notions: the robustness of the stationary distribution, and the robustness of three times: the mixing time, the maximum hitting time, and the cover time. We will use strong adversaries to define stationary and mixing time robustness, and standard adversaries for hitting time and cover time robustness. This is because our proofs work for strong adversaries for the first two robustnesses, but strong adversaries are too powerful for the last two robustnesses, as explained below.

1.2 Stationary robustness

We begin with the robustness of the stationary distribution π . Although MSCs may not converge to a distribution (see Section 3.1), we can still study the maximum and minimum limiting probabilities of sets. We can define this as follows. For $S \subseteq \Omega$, let $N_t(S)$ denote the number of visits to S up to time t .

► **Definition 1.1.** For $\varepsilon > 0$, the upper stationary probability

$$\pi^+(S) = \pi_\varepsilon^+(S) = \sup \limsup_{t \rightarrow \infty} N_t(S)/t,$$

where the outer sup is over the choice of a strong ε -adversary.

We could define a lower limit π^- analogously, but we focus on π^+ .

For our first theorem, we use a variant of the mixing time that we call t_{up} (see Section 2.3). Like mixing time, it satisfies $t_{\text{up}} = O(\frac{\log(1/\pi_{\min})}{\gamma})$, where γ is the spectral gap. We prove the following.

► **Theorem 1.2.** Let $G = (V, E)$ be a graph with all degrees at most d . Let $G = (V, E)$ be a graph with all degrees at most d . For any ε -strong adversary, any starting distribution, and any $S \subseteq V$ with $\pi(S) \leq 1 - \alpha$, we have

$$\pi^+(S) < (1 + \alpha) \exp(2\varepsilon \sqrt{2dt_{\text{up}} \ln(1/\pi(S))}) \pi(S).$$

In particular, for $\varepsilon = O(1/\sqrt{dt_{\text{up}}})$, we have $\pi^+(S) = \pi(S) \cdot \exp(O(\sqrt{\log(1/\pi(S))})) = O(\pi(S)^{1-o(1)})$.

Here the $o(1)$ term goes to 0 as $\pi(S) \rightarrow 0$, so in particular we have $\pi^+(S) = O(\pi(S)^{1-\alpha})$ for any $\alpha > 0$.

The dependence of ε on t_{up} is essentially best possible, as the path graph shows.

To interpret this theorem, consider the case when the state space V is exponentially large. Then any event with negligible stationary probability in the original chain has negligible stationary probability in this MSC with a nontrivial adversary. For example, [6] show that compression in their model happens with negligible probability; we deduce that this conclusion holds even if their chain is not followed exactly.

We could define stationary robustness as the largest ε such that the stationary probability analogues increase from p to at most $O(p^\beta)$ for some $\beta < 1$. The choice of β or quantification over β could a priori change the definition. However, Theorem 1.2 implies a robust stationarity of $\Omega(\sqrt{dt_{\text{up}}})$, and at least for the path the choice of β doesn't matter.

We compare Theorem 1.2 to the work of [4]. They only consider a time-independent adversary, although it wouldn't matter for this theorem. The main difference between their work and ours is that they show a lower bound on the adversarial power, whereas we show an upper bound. Specifically, they showed that for any bounded-degree graph and any

node v , some strong ε -adversary can achieve $\pi^+(v) \geq \pi(v)^{1-\Omega(\varepsilon)}$, and that this is tight for an expander. Haslegrave, Sauerwald, and Sylvester [13] generalized this to more settings. Our Theorem 1.2 upper bounds $\pi^+(S)$ for every graph and every strong ε -adversary.

1.3 Mixing time robustness

We now turn to the robustness of the fundamental times associated with a random walk: mixing time, max hitting time, and cover time. We begin with mixing time, which requires some set up. Let $|\cdot|$ denote variation distance. We will generalize the following for a standard random walk. For a node i , let \Pr_i denote the probability for a random walk starting at i . Let

$$\delta(t) = \max_i \left| \Pr_i[X_t \in \cdot] - \pi(\cdot) \right|$$

► **Definition 1.3.** Fix a small constant α , say $\alpha = 1/4$. The mixing time $t_{\text{mix}} = t_{\text{mix}}(\alpha) = \min\{t : \delta(t) \leq \alpha\}$.

To generalize this to MSCs, we define

$$\delta^+(t) = \sup \left(\max_i \left| \Pr_i[X_t \in \cdot] - \pi(\cdot) \right| \right).$$

Here π is the stationary distribution of the original Markov chain, and the sup is over strong ε -adversaries. Thus $\delta(t) \leq \delta^+(t)$. We can now define the robust mixing time, which will be a workable definition despite the MSC possibly not converging.

► **Definition 1.4.** Fix a small constant α , say $\alpha = 1/4$. Define the robust mixing time as

$$t_{\text{mix}}^+(\varepsilon) = \min\{t : \delta^+(t) \leq \alpha\},$$

where ε is the adversary parameter.

Note that $t_{\text{mix}} \leq t_{\text{mix}}^+(\varepsilon)$, and that $t_{\text{mix}}^+(\varepsilon)$ may be infinite. This leads to the following definition.

► **Definition 1.5.** The mixing time robustness is

$$\rho_{\text{mix}}(C) = \sup\{\varepsilon : t_{\text{mix}}^+(\varepsilon) \leq Ct_{\text{mix}}\}.$$

When C is omitted, we allow C to be an arbitrarily large constant.⁴

It is not hard to show that the mixing time robustness is $\Omega(1/(t_{\text{mix}}d))$, where d is the maximum degree; see Proposition 2.12. We achieve the square root of this quantity for constant-degree graphs.

► **Theorem 1.6.** In a graph with all degrees at most d , the mixing time robustness $\rho_{\text{mix}} = \Omega(1/\sqrt{t_{\text{mix}}d})$.

This is tight for the path and the cycle. However, for expander graphs, this gives $\Omega(1/\sqrt{\log n})$. We prove a constant lower bound for constant-degree expanders.

► **Theorem 1.7.** A d -regular γ -spectral expander has mixing time robustness $\Omega(\gamma/\sqrt{d})$.

⁴ This only makes sense when we have a family of graphs, but this applies to big Oh notation generally.

1.4 Hitting time robustness

We now turn to some expected stopping times. Let \mathbb{E}_u denote the expectation for a Markov (set) chain starting at $u \in V$. Let H_v denote the hitting time of v : the time to first visit v after time 0. The standard hitting times in an unbiased random walk are $h(u, v) = \mathbb{E}_u H_v$, and the max hitting time is $t_{\text{hit}} = \max_{u \neq v} h(u, v)$.

Before defining the robust version, we briefly discuss the adversary. Note that if we have a strong ε -adversary, the robustness for hitting and covering is less than $1/d_{\min}$, where d_{\min} denotes the minimum degree. That's because a strong $1/d_{\min}$ -adversary can always avoid a particular node. Due to this and our proofs working better with standard adversaries, we define the robust versions for standard adversaries.

► **Definition 1.8.** Fix an $\varepsilon > 0$. The robust hitting time $h^+(u, v) = \sup(\mathbb{E}_u H_v)$, where the sup is over strategies for standard ε -adversaries. The robust max hitting time is $t_{\text{hit}}^+ = \max_{u, v} h^+(u, v)$.

Generalizing an observation by [4], the quantities h^+ are the same for oblivious, Markovian, and non-Markovian adversaries. This is because an adversary's best strategy at a given node does not depend on the previous history. This basically follows from Markov decision theory.

We focus on t_{hit}^+ so we have one quantity rather than n^2 quantities. We can now define:

► **Definition 1.9.** The hitting time robustness is

$$\rho_{\text{hit}}(C) = \sup\{\varepsilon : t_{\text{hit}}^+(\varepsilon) \leq Ct_{\text{hit}}\}.$$

When C is omitted, we allow C to be an arbitrarily large constant.

It is not hard to show that the hitting time robustness is $\Omega(1/t_{\text{hit}})$; see Proposition 2.12. We improve this bound dramatically.

► **Theorem 1.10.** The hitting time robustness $\rho_{\text{hit}} = \Omega(1/t_{\text{mix}})$.

We can improve this bound for trees. Let Δ denote the diameter.

► **Theorem 1.11.** The hitting time robustness for any tree is $\Omega(1/\Delta)$.

Note that $\Delta = O(t_{\text{mix}})$, and is often significantly less. For example, for the line and cycle on n nodes, $\Delta = \Theta(n)$ but $t_{\text{mix}} = \Theta(n^2)$. Even more dramatically, for the balanced binary tree $t_{\text{mix}} = \Theta(n \log n)$ but $\Delta = \Theta(\log n)$.

For constant degree expanders, Theorem 1.10 gives a bound of $\Omega(1/\log n)$. We show that this is tight for expanders by giving a matching upper bound, which is a negative result.

► **Theorem 1.12.** The hitting time robustness of any constant degree expander is $\rho_{\text{hit}} = \Theta(1/\log n)$.

1.5 Cover time robustness

We now turn to the cover time. Let the random variable T denote the first time the random walk has visited all nodes. We define the cover time $t_{\text{cov}} = \max_u \mathbb{E}_u T$.

► **Definition 1.13.** Fix an $\varepsilon > 0$. The robust cover time $t_{\text{cov}}^+ = \sup(\mathbb{E}_u T)$, where the sup is over strategies for standard ε -adversaries.

Note that Markovian and non-Markovian adversaries do not appear to be equivalent, since the prior history is very relevant for the cover time. We now define:

► **Definition 1.14.** *The cover time robustness is*

$$\rho_{\text{cov}}(C) = \sup\{\varepsilon : t_{\text{cov}}^+(\varepsilon) \leq Ct_{\text{cov}}\}.$$

When C is omitted, we allow C to be an arbitrarily large constant.

It is not hard to show that the cover time robustness is $\Omega(1/t_{\text{cov}})$; see Proposition 2.12. We improve this in many cases. We always have $t_{\text{cov}} = O(t_{\text{hit}} \log n)$; we can improve the cover time robustness whenever this is tight up to a constant factor.

► **Theorem 1.15.** *In any graph with $t_{\text{cov}} = \Theta(t_{\text{hit}} \log n)$, we have $\rho_{\text{cov}} = \Omega(\rho_{\text{hit}})$.*

There are many graphs where $t_{\text{cov}} = \Theta(t_{\text{hit}} \log n)$, including expanders (and hence most graphs) [5], the complete graph (coupon collecting), two and higher dimensional grids and tori [2, 25], and the balanced binary tree [24].

1.6 Examples

It's instructive to see how our bounds apply to natural examples.

1.6.1 Path/cycle

The path is the graph on $[n]$ with $n - 1$ edges $\{i, i + 1\}$; the cycle contains the additional edge $\{1, n\}$. In the non-adversarial setting, the main quantities are similar for the two graphs. The mixing time, max hitting time, and cover time are all $\Theta(n^2)$.

To understand the adversarial setting, we focus on the path. The worst a strong adversary can do is set all $P_t(i, i + 1) = (1 + \varepsilon)/2$ and all $P_t(i, i - 1) = (1 - \varepsilon)/2$. In this case, the hitting times will increase by a factor of about $((1 + \varepsilon)/(1 - \varepsilon))^n$. Therefore, the hitting time robustness is $\Theta(1/n) = \Theta(1/\sqrt{t_{\text{mix}}})$. All other robustnesses have the same order of magnitude. Thus, Theorem 1.6 is tight for mixing time robustness, but Theorem 1.10 is not tight for hitting time robustness. See Section 4.4 for details.

1.6.2 Expanders

Here we consider d -regular expander graphs, where we think of $d = O(1)$. In the non-adversarial setting, the mixing time $t_{\text{mix}} = \Theta(\log n)$, the max hitting time $t_{\text{hit}} = \Theta(n)$ [21], and the cover time $t_{\text{cov}} = \Theta(n \log n)$ [5, 21]. We actually prove these last two theorems in what we believe is a simpler way.

Our theorems imply tight bounds on the robustness notions for constant-degree regular expanders. Specifically, Theorem 1.7 show that the mixing time robustness is a constant, while Theorem 1.12 and Theorem 1.15 show that the hitting and cover time robustnesses are $\Theta(1/\log n)$.

We remark that previous work shows that lossless expanders have robustness properties beyond what we consider here. A d -regular graph is a lossless expander if sets that are not too large expand by a factor of $(1 - \eta)d$ where $\eta < 1/2$. Doron, Moshkovitz, Oh, and Zuckerman [9, 10] studied random walks where a non-Markovian adversary chooses steps according to a Chor-Goldreich (CG) source, and even generalizations of this. In a CG-source, in a d -regular graph, an adversary can choose any distribution for the next step such that no edge is chosen with probability more than $d^{-\delta}$ for some parameter $\delta > 0$. They show that even this biased walk mixes in the sense that the final node has entropy $\log n - O(1)$. However, their methods do not seem to extend to standard expander graphs.

1.7 Weak randomness

As another perspective, there has been a lot of research investigating the use of weak randomness in computing. In those situations, it is natural to allow computation on the weak randomness, for example to purify it into high-quality randomness. This work disallows such preprocessing, which is particularly natural when the random process models nature or social interactions.

1.8 Techniques

We prove Theorem 1.2 and Theorem 1.6 about stationarity and mixing time robustness by generalizing a potential function argument by Haslegrave, Sauerwald, and Sylvester [13]. Their potential function uses squares to prove a lower bound. We use a potential function based on smaller powers to prove an upper bound. In particular, we need an inequality going in the opposite direction. Specifically, for $x = (x_1, \dots, x_d) \in [0, \infty)^d$ define the ℓ -power mean $M_\ell(x) = ((\sum_i x_i^\ell)/d)^{1/\ell}$, $M_\infty = \max_i x_i$, and the ε -max-avg operator $\text{MA}_\varepsilon = \varepsilon M_\infty + (1 - \varepsilon)M_1$. We use Holder's inequality and other ideas to bound $\text{MA}_\varepsilon(x)$ in terms of $M_\ell(x)$ for some ℓ just slightly larger than 1. We also simplify the [13] method somewhat, avoiding the use of trajectory trees.

We prove Theorem 1.7 about mixing time robustness of expanders by adapting the proof that the second largest eigenvalue determines convergence.

We prove Theorem 1.10 about hitting time robustness by showing that in a standard random walk, there's the "right" probability of hitting a node within twice the separation time $t_{\text{sep}} = O(t_{\text{mix}})$. Specifically, for $h_v = \mathbb{E}_\pi[H_v]$, we show that for all $u, v \in V$, we have $\Pr_u[H_v \leq 2t_{\text{sep}}] = \Omega(t_{\text{sep}}/h_v)$. We then show that this probability doesn't decline much even in the robust setting. Dividing the random walk into epochs of length $2t_{\text{sep}}$, we conclude that the hitting time doesn't increase much. We further use this technique to give a simpler proof of the max hitting time for expanders, which will be in the final version of this paper. Given its utility for both of these theorems, we believe that this technique could be useful elsewhere.

We prove Theorem 1.11 by first observing that we can take the adversary to be oblivious, so the corresponding adversarial walk has a stationary distribution. We then compare the stationary distributions, and apply the "essential edge lemma" from Aldous-Fill about hitting times across cut edges.

We prove Theorem 1.12, the negative result giving a tight bound for expanders, by giving a target distribution and modifying the Metropolis algorithm. Much of this argument is similar to a lower bound on $\pi^+(S)$ by [4].

Theorem 1.15 about cover time robustness follows from Matthews' technique for bounding the cover time in terms of the hitting time.

1.9 Related work

We have not found any previous work giving good bounds on our notions of robustness, except for lossless expanders. Doron, Moshkovitz, Oh, and Zuckerman [9, 10] analyzed extremely strong non-Markovian adversaries for lossless expanders, but their methods don't apply more generally. They also didn't analyze hitting or cover times.

In the model introduced by [4] and in follow-up works, the focus is on a controller seeking to minimize hitting and cover times, and they show lower bounds on increasing the probability of being in sets. For example, Haslegrave, Sauerwald, and Sylvester [13] improved a result of [4] and proved that an ε -controller can increase the probability of being in a set from p

to approximately $p^{1-\varepsilon}$. They used this to minimize hitting and cover times, for example showing how to reduce the cover time of an expander to $O(n \log \log n)$. This is different than our work where we view the controller as an adversary, so hitting times and cover times increase rather than decrease. Moreover, the quantification is different: they show that a suitable adversary exists, whereas we show that all adversaries are not too bad.

Chin, Moitra, Mossel, and Sandon [7] studied the power of an adversary in Glauber dynamics. Glauber dynamics is a type of random walk on Σ^V , where V is a set of nodes in an underlying graph and Σ is a finite alphabet. They proposed an adversary that controls any changes to a subset $A \subseteq V$ with $|A| \leq \varepsilon|V|$. They studied how such an adversary can affect global features of the system, such as the approximate mixing time. Their methods seem tailored to Glauber dynamics and not generalizable.

There is also research analyzing how the stationary distribution of a Markov chain is affected by perturbations. For example, Liu ([18], Theorem 2.2) showed that if P is perturbed to \tilde{P} , which has stationary distribution $\tilde{\pi}$, then

$$|\pi - \tilde{\pi}| \leq t_{\text{hit}}^2 |P - \tilde{P}|,$$

where $|M|$ for a matrix M is the maximum ℓ_1 -norm of any row. Thus, the best bound on robustness that this could give is $\Omega(1/t_{\text{hit}}^2)$, which is much weaker than what we obtain. Moreover, our robustness is for Markov set chains, which are much more general than Markov chains.

Hartfiel [12] studies when Markov set-chains converge, and movement between transient states and absorbing states and the like. He doesn't address the types of questions we ask.

Avin, Koucky, and Lotker [3] study random walks on evolving graphs. Here an adversary is allowed to completely change the graph, i.e., transition matrix, at every time step. However, the adversary doesn't know the location of the random walker; in other words, the adversary decides the graphs in advance. They show that the cover time and mixing time could be exponentially larger than the cover times and mixing times of the individual graphs, but are only polynomially larger if the random walk is made lazy. This model is quite different from ours.

Hunter [14] studies a related model, what in our terminology is a strong oblivious adversary. He defines η to be the "random target time" – the time for a random walk to hit a random node chosen according to π – but confusingly calls this the mixing time. Really, η is more closely related to t_{hit} ; see [16] for more on the random target time. He then shows that the adversarial walk has a stationary distribution at most $\eta\varepsilon$ -far from π in variation distance. This is much weaker than our results.

There are other adversarial variants of Markov chains that don't appear relevant for us. For example, the PageRank algorithm [20] and variants (e.g., [23]) allow perturbations where the adversary need not respect the graph structure, which leads to very different results. For another example, [8] define a model that they call "adversarial Markov chains" on an infinite metric space. In their model, transitions are followed exactly except when the chain is in some prespecified bounded subset, in which case an adversary can make arbitrary bounded jumps. They study whether such an adversarial chain remains bounded in probability.

1.10 Conjectures and Open Problems

We conjecture the following.

► **Conjecture 1.16.** *The hitting time robustness is $\Omega(1/\Delta)$, where Δ denotes the diameter.*

118:10 Markov Chain Robustness

Since $\Delta = O(t_{\text{mix}})$, this would improve Theorem 1.10. Theorem 1.11 establishes this conjecture for trees, which is tight for the path and cycle.

► **Conjecture 1.17.** *For all graphs, the cover time robustness is the hitting time robustness, up to a constant factor.*

► **Conjecture 1.18.** *We can replace t_{up} in Theorem 1.2 with t_{mix} .*

What do we need to know about a graph to determine its robustnesses? Are the different robustnesses efficiently computable, perhaps up to constant factors?

Another class of interesting open problems is to determine the various robustnesses for natural graphs, such as different dimensional grids, hypercubes, and balanced trees.

2 Preliminaries

2.1 Random walk basics

We will assume throughout that any undirected graph $G = (V, E)$ is connected. We denote $n = |V|$, $m = |E|$, $\Gamma(v)$ the set of neighbors of node $v \in V$, and $d_v = |\Gamma(v)|$ the degree of v . A random walk on G is a sequence of random variables (X_0, X_1, \dots) where each $X_t \in V$ and X_{t+1} is uniform on $\Gamma(X_t)$. We use P to denote the transition matrix of the random walk. Letting p_t denote the probability distribution of X_t , viewed as a row vector, so we have $p_{t+1} = p_t P$. If G is not bipartite, then p_t converges to the stationary distribution π given by $\pi(v) = d_v / (2m)$. Even if G is bipartite, $(p_t + p_{t+1})/2$ converges to π .

Let \mathbb{E}_u and \Pr_u denote the expectation and probability, respectively, for a random walk starting at $u \in V$. Let H_v denote the hitting time of v : the time to first visit v after time 0. The standard hitting times in an unbiased random walk are $h(u, v) = \mathbb{E}_u H_v$, and the max hitting time is $t_{\text{hit}} = \max_{u \neq v} h(u, v)$.

The following lemma is well known.

► **Lemma 2.1.** *The expected return time $h(v, v) = 1/\pi(v)$.*

2.2 Expander graphs

Expanders are graphs where every subset of nodes has many neighbors. We will use the standard spectral characterization, which is equivalent for constant expansion. For simplicity, assume that G is regular. Denote the eigenvalues of P by $1 = \lambda_1 > \lambda_2 \geq \dots \geq \lambda_n \geq -1$. The last equality holds only when G is bipartite. Let us assume that G is not bipartite, e.g., by adding self loops. Let

$$\lambda = \max(\lambda_2, -\lambda_n) = \max_{i \neq 1} |\lambda_i|$$

denote the second largest eigenvalue in absolute value. Let $\gamma \stackrel{\text{def}}{=} 1 - \lambda$ denote the spectral gap. We call G a γ -spectral expander.

2.3 Mixing time variants

Recall the definition of mixing time, Definition 1.3. The choice of α only affects the mixing time up to a constant factor. Specifically:

► **Lemma 2.2.** *In any graph, $t_{\text{mix}}(\alpha) \leq \lceil \log_2 1/\alpha \rceil t_{\text{mix}}(1/4)$.*

For this and much more on the mixing time see, for example, the book by Levin, Peres, and Wilmer [16].

Up to a logarithmic factor, the reciprocal of the spectral gap γ characterizes the mixing time. Specifically, it is known (e.g., [16] Theorems 12.3 and 12.4) that

$$\begin{aligned} t_{\text{mix}}(\alpha) &\leq (1/\gamma) \log(1/(\alpha\pi_{\min})) \\ t_{\text{mix}}(\alpha) &\geq (1/\gamma - 1) \log(1/(2\alpha)) \end{aligned}$$

We will use a variant of the mixing time called the separation time t_{sep} .

► **Definition 2.3.** *The separation distance is:*

$$s(t) = \max_{v \in V} (1 - p_t(v)/\pi(v)).$$

In other words, it's the minimum β such that all nodes v have $p_t(v) \geq (1 - \beta)\pi(v)$.

We can now define separation time in the natural way.

► **Definition 2.4.** *Fix a small constant α , say $\alpha = 1/4$. Define the separation time*

$$t_{\text{sep}} = t_{\text{sep}}(\alpha) = \min\{t : s(t) \leq \alpha\}.$$

► **Lemma 2.5.** *In any graph, $t_{\text{sep}} = \Theta(t_{\text{mix}})$.*

For the proof, see Aldous-Fill [1]. Briefly, it's immediate that $\delta(t) \leq s(t)$. Aldous-Fill (Lemma 4.7 in [1]) show that $s(2t) \leq 1 - (1 - \delta(t))^2$.

We further define a variant of the separation distance briefly discussed in Aldous-Fill in Equation 4.14. They left it unnamed, so we name it.

► **Definition 2.6.** *The upper separation distance is:*

$$s_{\text{up}}(t) = \max_{v \in V} (p_t(v)/\pi(v) - 1).$$

In other words, it's the minimum β such that all nodes v have $p_t(v) \leq (1 + \beta)\pi(v)$.

We can now define upper separation time analogously.

► **Definition 2.7.** *Fix a small constant α , say $\alpha = 1/4$. Define the upper separation time*

$$t_{\text{up}} = t_{\text{up}}(\alpha) = \min\{t : s_{\text{up}}(t) \leq \alpha\}.$$

While we have $\delta(t) \leq s_{\text{up}}(t)$, we can't bound $s_{\text{up}}(2t)$ in terms of $\delta(t)$; see Aldous-Fill's Example 4.9. Nevertheless, the upper bounds of t_{mix} in terms of the spectral gap γ also hold for t_{sep} and t_{up} , just by examining the proof.

► **Lemma 2.8.** *We have $t_{\text{sep}}(\alpha), t_{\text{up}}(\alpha) \leq \frac{-\log(\alpha\pi_{\min})}{\gamma}$.*

► **Remark 2.9.** Since standard random walks on bipartite graphs don't converge, the above definitions aren't interesting for bipartite graphs. The usual way of circumventing this is to study lazy random walks on bipartite graphs. Alternatively, we can get meaningful statements for bipartite graphs by replacing p_t with $(p_t + p_{t+1})/2$ in the above definitions, understanding that even though the random walks don't converge, the average of pairs of steps converge.

2.4 Easy bounds on robustness

We will use the following notation.

► **Definition 2.10.** Let $p_t(v, S)$ denote the probability that an unbiased walk, starting at node v , is in the set S at time t . For a given adversary A , let $q_t(v, S)$ denote the probability that an A -biased walk, starting at node v , is in the set S at time t . When the start node v is understood, we often omit it and write $p_t(S)$ and $q_t(S)$.

We have the following simple lemma.

► **Lemma 2.11.** For any standard ε -adversary, any time t , and any set S , we have $q_t(S) \leq e^{t\varepsilon} p_t(S)$.

Note that this does not hold for a strong ε -adversary, since such an adversary could move ε probability to a set S with tiny $p_t(S)$.

It's not hard to show that each robustness is at least the reciprocal of the corresponding quantity.

► **Proposition 2.12.** For any graph G , the following bounds hold:

1. The mixing time robustness, if defined by a standard adversary, is $\Omega(1/t_{\text{mix}})$. Since we defined mixing time robustness ρ_{mix} using a strong adversary, this implies $\rho_{\text{mix}} = \Omega(1/(dt_{\text{mix}}))$ for graphs with maximum degree d .
2. The hitting time robustness $\rho_{\text{hit}} = \Omega(1/t_{\text{hit}})$.
3. The cover time robustness $\rho_{\text{cov}} = \Omega(1/t_{\text{cov}})$.

Proof. First we show the mixing time robustness. Using Lemma 2.2, choose $t = O(t_{\text{mix}})$ such that $|p_t - \pi| \leq \alpha/2$. Set $\varepsilon = \alpha/(2t) = \Omega(1/t_{\text{mix}})$. Then a standard ε -adversary can't make the biased walk probabilities q_t much smaller than the unbiased probabilities p_t . Specifically, for any set S , we have

$$q_t(S) \geq (1 - \varepsilon)^t p_t(S) \geq (1 - t\varepsilon) p_t(S) = (1 - \alpha/2) p_t(S).$$

Therefore, $|q_t - p_t| \leq \alpha/2$, and the triangle inequality gives $|q_t - \pi| \leq \alpha$.

We prove the hitting time and cover time robustness results together. Let R be a stopping time such as a hitting time or cover time. Let $t = 4 \max_u \mathbb{E}_u[R]$. Then by Markov's inequality, for any node v , we have $\Pr_v[R \geq t] \leq 1/4$, so $\Pr_v[R \leq t] \geq 3/4$. Set $\varepsilon = 1/(4t)$. Fix a strong ε -adversary, and let R^+ denote R with respect to this adversary. Then for any node v , we have

$$\Pr_v[R^+ \leq t] \geq (1 - t\varepsilon) 3/4 \geq 1/2.$$

But now dividing time into epochs of length t , the chance of R^+ in each epoch is at least $1/2$, so by Wald's identity we have $\mathbb{E}_u[R^+] \leq 2t \leq 8 \max_u \mathbb{E}_u[R]$. ◀

3 Stationary and Mixing Time Robustness

In this section, we first show that for a Markovian adversary, a Markov set-chain will not necessarily converge to π , even if all underlying chains are reversible and have π as their stationary distribution. We only consider a Markovian adversary in Section 3.1; in the rest of this section we consider standard and strong adversaries. Next we show how to save a square root from the easy mixing time robustness of $\Omega(1/t_{\text{mix}})$, proving Theorem 1.6. We add one more ingredient to prove Theorem 1.2. Finally, we establish that constant-degree expanders have constant mixing time robustness, namely $\Omega(1/\gamma)$.

3.1 MSCs non-convergence

First we show that for a Markovian adversary, the MSC will not necessarily converge to π , even if all underlying chains are reversible. To see this, we define two weighted undirected graphs based on the path graph. Denote $\text{Odd} = \{\{i, i+1\} \mid i \text{ is odd}\}$, and $\text{Even} = \{\{i, i+1\} \mid i \text{ is even}\}$. Graph G_{odd} assigns edges in Odd weight 2, and edges in Even weight 1, and adds weighted self loops at the two endpoints to make all nodes have weighted degree 3. Graph G_{even} is the same but swaps even and odd: it assigns edges in Even weight 2, and edges in Odd weight 1. Let P_{odd} denote a random walk on G_{odd} , and similarly for P_{even} . Since G_{odd} and G_{even} are both regular, both have uniform stationary distributions.

However, consider a Markovian MSC that chooses P_{odd} if X_t is odd and P_{even} otherwise. If X_t is not an endpoint, then $\Pr[X_{t+1} = X_t + 1] = 2/3$. We thus have a biased random walk which converges to a very biased distribution very far from uniform.

3.2 Bounding the increase in probability

We prove that if a set S has probability p , then a strong ε -adversary cannot increase the probability to more than $O(p^{1-o(1)})$ as long as $t\varepsilon^2 = O(1)$. This quadratic dependence on ε improves on the easy linear dependence given by Lemma 2.11. We use a potential function argument that generalizes one by Haslegrave, Sauerwald, and Sylvester [13]. They used a potential function based on squares to prove a lower bound. We use a potential function based on smaller powers to prove an upper bound. In particular, we need an inequality going in the opposite direction as [13].

For $x = (x_1, \dots, x_d) \in [0, \infty)^d$ define the ℓ -power mean $M_\ell(x) = ((\sum_i x_i^\ell)/d)^{1/\ell}$, $M_\infty = \max_i x_i$, and the ε -max-avg operator $\text{MA}_\varepsilon = \varepsilon M_\infty + (1 - \varepsilon)M_1$. We now develop the inequality we need.

► **Lemma 3.1.** *For $k, r \geq 1$ and $kr\varepsilon \leq 1$, we have $(1 + r\varepsilon)^k + r(1 - \varepsilon)^k \leq r + 1 + 2k^2r^2\varepsilon^2$.*

Proof.

$$\begin{aligned}
(1 + r\varepsilon)^k + r(1 - \varepsilon)^k &= 1 + kr\varepsilon + \binom{k}{2}(r\varepsilon)^2 + \binom{k}{3}(r\varepsilon)^3 + \dots \\
&\quad + r - kr\varepsilon + \binom{k}{2}r\varepsilon^2 - \binom{k}{3}r\varepsilon^3 + \dots \\
&\leq r + 1 + 2 \left(\binom{k}{2}(r\varepsilon)^2 + \binom{k}{3}(r\varepsilon)^3 + \dots \right) \\
&\leq r + 1 + 2 \binom{k}{2}(r\varepsilon)^2 (1 + (kr\varepsilon/2) + (kr\varepsilon/2)^2 + \dots) \\
&\leq r + 1 + 4 \binom{k}{2}(r\varepsilon)^2 \leq r + 1 + 2k^2r^2\varepsilon^2. \quad \blacktriangleleft
\end{aligned}$$

► **Lemma 3.2.** *For $x \in [0, \infty)^d$ and $k, \ell > 1$ satisfying $1/k + 1/\ell = 1$ and $k(d-1)\varepsilon \leq 1$, we have $\text{MA}_\varepsilon(x) \leq e^{2k(d-1)\varepsilon^2} M_\ell(x)$.*

Proof. Let $x = (x_1, \dots, x_d)$, and assume without loss of generality that $x_1 \geq x_2 \geq \dots \geq x_d \geq 0$. Using Holder's inequality and Lemma 3.1 with $r = d - 1$,

118:14 Markov Chain Robustness

$$\begin{aligned}
\text{MA}_\varepsilon(x) &= \left(\varepsilon + \frac{1-\varepsilon}{d} \right) x_1 + \frac{1-\varepsilon}{d} x_2 + \cdots + \frac{1-\varepsilon}{d} x_d \\
&\leq \frac{1}{d} \left((1+(d-1)\varepsilon)^k + (d-1)(1-\varepsilon)^k \right)^{1/k} (x_1^\ell + \cdots + x_d^\ell)^{1/\ell} \\
&\leq \frac{1}{d} (d + 2k^2(d-1)^2\varepsilon^2)^{1/k} (x_1^\ell + \cdots + x_d^\ell)^{1/\ell} \\
&\leq (1 + 2k^2(d-1)\varepsilon^2)^{1/k} M_\ell(x) \\
&\leq e^{2k(d-1)\varepsilon^2} M_\ell(x). \quad \blacktriangleleft
\end{aligned}$$

Recall Definition 2.10 for the definitions of p_t and q_t .

► **Theorem 3.3.** *Let $G = (V, E)$ be a graph with all degrees at most d , and fix a starting node. Let $k, \ell > 1$ satisfy $1/k + 1/\ell = 1$ and $kd\varepsilon \leq 1$. For any ε -strong adversary, any start node, and any $S \subseteq V$, we have $q_t(S) \leq e^{2tkd\varepsilon^2} p_t(S)^{1/\ell}$.*

Proof. We follow the structure of [13] but both simplify it somewhat, avoiding the use of the trajectory tree, and generalize it. Fix the transition matrix P , a start node u , target set S , and stopping time t . Define the vectors

$$\begin{aligned}
p_t &= (p_t(v))_{v \in V} \\
s_t &= (q_t(v, S)^\ell)_{v \in V}.
\end{aligned}$$

When intermixed with matrices, we will view p_t as a row vector and s_t as a column vector, for reasons that will become apparent. Note that $p_{t+1} = p_t P$. Define the potential function

$$\Phi_i = \langle p_{t-i}, s_i \rangle = \sum_{v \in V} p_{t-i}(v) q_i(v, S)^\ell.$$

Since $q_0(v, S) = 1$ if $v \in S$ and 0 otherwise, we have that

$$\Phi_0 = \sum_{v \in V} p_t(v) q_0(v, S)^\ell = p_t(S).$$

Since $p_0(v) = 1$ if $v = u$ and 0 otherwise, we have that

$$\Phi_t = \sum_{v \in V} p_0(v) q_t(v, S)^\ell = q_t(u, S)^\ell = q_t(S)^\ell.$$

If the random walk is at node v with $i + 1$ steps remaining, the adversary's strategy to maximize $q_{i+1}(v, S)$ is to use the ε probability it controls to move to a neighbor w maximizing $q_i(w, S)$. Letting d_v denote the degree of v , this strategy gives

$$q_{i+1}(v, S) = \text{MA}_\varepsilon((q_i(w, S))_{w \in \Gamma(v)}) \leq e^{2kd_v\varepsilon^2} M_\ell((q_i(w, S))_{w \in \Gamma(v)}),$$

using Lemma 3.2. Raising both sides to the ℓ th power gives

$$q_{i+1}(v, S)^\ell \leq e^{2k\ell d_v \varepsilon^2} \frac{1}{d_v} \sum_{w \in \Gamma(v)} q_i(w, S)^\ell.$$

Using the fact that all degrees are at most d , we can deduce an inequality with vectors and matrices:

$$s_{i+1} \leq e^{2k\ell d \varepsilon^2} P s_i.$$

Now we can bound the potential function by

$$\Phi_{i+1} = \langle p_{t-i-1}, s_{i+1} \rangle \leq e^{2k\ell d\varepsilon^2} p_{t-i-1} P s_i = e^{2k\ell d\varepsilon^2} \langle p_{t-i}, s_i \rangle = e^{2k\ell d\varepsilon^2} \Phi_i.$$

We conclude that $\Phi_t \leq \exp(2tk\ell d\varepsilon^2)\Phi_0$, and

$$q_t(S) = \Phi_t^{1/\ell} \leq e^{2tkd\varepsilon^2} \Phi_0^{1/\ell} = e^{2tkd\varepsilon^2} p_t(S)^{1/\ell}. \quad \blacktriangleleft$$

► **Corollary 3.4.** *The conclusion of Theorem 3.3 holds for any starting distribution.*

Proof. This follows from the concavity of the function x^ℓ . ◀

► **Corollary 3.5.** *For any ε -strong adversary, any starting distribution, and any $S \subseteq V$, we have*

$$q_t(S) \leq \exp(2\varepsilon\sqrt{2td\ln(1/p)}) \cdot p,$$

where $p = p_t(S)$. In particular, for $\varepsilon = O(1/\sqrt{td})$, we have $q_t(S) = O(p^{1-o(1)})$.

Proof. We choose k to optimize the upper bound for $q_t(S)$. Let $c = 2td\varepsilon^2$, $p = p_t(S)$, and $r = \ln(1/p)$. Since $1/\ell = 1 - 1/k$, the upper bound is

$$e^{ck} p^{1-1/k} = e^{ck+r/k} p.$$

By the arithmetic-geometric mean inequality, this is minimized when $ck = r/k$ and achieves the bound $e^{2\sqrt{cr}} p = \exp(2\varepsilon\sqrt{2tdr}) \cdot p$. ◀

We now restate and prove Theorem 1.6, which is tight for the path and the cycle.

► **Theorem 1.6.** *In a graph with all degrees at most d , the mixing time robustness $\rho_{\text{mix}} = \Omega(1/\sqrt{t_{\text{mix}}d})$.*

Proof. While we can deduce this from Corollary 3.5, it's more straightforward to apply Theorem 3.3 directly. We will choose suitable constants k and ℓ that satisfy the assumptions of Theorem 3.3 and use the conclusion. Letting $c \stackrel{\text{def}}{=} \exp(2tkd\varepsilon^2)$, we can bound the variation distance

$$|q_t - p_t| \leq \sup_{p \in [0,1]} (cp^{1/\ell} - p).$$

The derivative of $cp^{1/\ell} - p$ is $(c/\ell)p^{1/\ell-1} - 1 = (c/\ell)p^{-1/k} - 1$. There are two cases. If $c \geq \ell$, then the derivative is nonnegative on $[0, 1]$, so the sup is achieved at $p = 1$ and the sup is $c - 1$. If $c < \ell$, then the sup is achieved when $cp^{-1/k} = \ell$. This value of p evaluates to:

$$cp^{1/\ell} - p = p(cp^{-1/k} - 1) = p(\ell - 1) \leq 2p/k \leq 2/k,$$

using that $\ell = 1/(1 - 1/k) \leq 1 + 2/k$ for $k \geq 2$.

We conclude that $|q_t - p_t| \leq \max(c - 1, 2/k)$. By choosing $k = 4/\alpha$ and $\varepsilon = \sqrt{\alpha/(8tkd)}$, and using that $e^x \leq 1 + 2x$ for $x \leq 1/2$, we get that this maximum is at most $\alpha/2$. By Lemma 2.2, after $t = O(t_{\text{mix}})$ steps we have $|p_t - \pi| \leq \alpha/2$. Therefore, by the triangle inequality, $|q_t - \pi| \leq \alpha$. Thus the robustness is bounded below by

$$\varepsilon = \sqrt{\alpha/(8tkd)} = \Omega(1/\sqrt{t_{\text{mix}}d}). \quad \blacktriangleleft$$

We now use Corollary 3.5 to upper bound $\pi^+(S)$ and prove Theorem 1.2.

► **Theorem 1.2.** *Let $G = (V, E)$ be a graph with all degrees at most d . Let $G = (V, E)$ be a graph with all degrees at most d . For any ε -strong adversary, any starting distribution, and any $S \subseteq V$ with $\pi(S) \leq 1 - \alpha$, we have*

$$\pi^+(S) < (1 + \alpha) \exp(2\varepsilon \sqrt{2dt_{\text{up}} \ln(1/\pi(S))}) \pi(S).$$

In particular, for $\varepsilon = O(1/\sqrt{dt_{\text{up}}})$, we have $\pi^+(S) = \pi(S) \cdot \exp(O(\sqrt{\log(1/\pi(S))})) = O(\pi(S)^{1-o(1)})$.

Proof. As usual, let $p_t(\cdot)$ and $q_t(\cdot)$ denote unbiased and adversarially-biased probabilities, respectively. Let $S \subseteq V$ be arbitrary, and fix a starting node v for the random walk. First assume that G is bipartite. Note that to show $\pi^+(S) \leq B$, it suffices to show that for all large enough t we have $q_t(S) \leq B$. Let $c = 2\varepsilon \sqrt{2dt_{\text{up}}}$ and $f(x) = xe^{c\sqrt{\ln(1/x)}}$. To bound $q_t(S)$ for $t \geq t_{\text{up}}$, set $t' = t - t_{\text{up}}$, and condition on $X_{t'}$. Recall that by the definition of t_{up} , we have $p_{t_{\text{up}}}(S) \leq (1 + \alpha)\pi(S)$. Using $f(rx) \leq rf(x)$ for $rx < 1$ and $r > 1$, and applying Corollary 3.5, we have

$$q_t(v, S) = \mathbb{E}_{X_{t'}} [q_{t_{\text{up}}}(X_{t'}, S)] \leq \mathbb{E}_{X_{t'}} [f(p_{t_{\text{up}}}(X_{t'}, S))] \leq (1 + \alpha)f(\pi(S)),$$

as required for the nonbipartite case.

The bipartite case follows similarly by comparing $(q_t + q_{t+1})/2$ with $(p_t + p_{t+1})/2$. ◀

3.3 Mixing time robustness for expanders

We now establish robustness in terms of spectral gap. See Section 2.2 for the definition of a spectral expander. We first restate our theorem

► **Theorem 1.7.** *A d -regular γ -spectral expander has mixing time robustness $\Omega(\gamma/\sqrt{d})$.*

Proof. Our strong adversary chooses a stochastic matrix B_t with $\text{supp}(B_t) \subseteq \text{supp}(P)$, and the transition matrix $P_t = (1 - \varepsilon)P + \varepsilon B_t$. Let q_t denote the vector of probabilities at time t . We can write $q_t = \pi + r_t$, where $r_t \perp \pi$ can be viewed as an error term. Therefore,

$$\begin{aligned} q_{t+1} &= (1 - \varepsilon)Pq_t + \varepsilon B_t q_t \\ &= (1 - \varepsilon)\pi + (1 - \varepsilon)Pr_t + \varepsilon B_t \pi + \varepsilon B_t r_t \end{aligned}$$

Subtracting π gives

$$r_{t+1} = (1 - \varepsilon)Pr_t + \varepsilon B_t r_t + \varepsilon(B_t \pi - \pi)$$

Using that for stochastic B with row sums at most d , we have $\|Bv\|_2 \leq \sqrt{d}\|v\|_2$, $\|(B-I)v\|_2 \leq \sqrt{d}\|v\|_2$, and $\|Pw\|_2 \leq \lambda\|w\|_2$ for $w \perp \pi$ yields:

$$\|r_{t+1}\|_2 \leq ((1 - \varepsilon)\lambda + \varepsilon\sqrt{d})\|r_t\|_2 + \varepsilon\sqrt{d/n} \tag{1}$$

We will write (1) as $\|r_{t+1}\|_2 \leq b\|r_t\|_2 + \delta$ to simplify our calculation. Our calculation will show that we can choose ε small enough to make r_t converge quickly. For $\varepsilon_0 \leq 1$ to be chosen later, set

$$\varepsilon = \frac{(1 - \lambda)\varepsilon_0}{2(\sqrt{d} - \lambda)}.$$

Then $b \stackrel{\text{def}}{=} (1 - \varepsilon)\lambda + \varepsilon\sqrt{d} = \lambda + \varepsilon(\sqrt{d} - \lambda) \leq \lambda + (1 - \lambda)/2 = (1 + \lambda)/2 < 1$.

Let $\delta \stackrel{\text{def}}{=} \varepsilon\sqrt{d/n}$. Note that for $x' \leq bx + \delta$ with $b < 1$, and for $x \geq 2\delta/(1-b)$, we have $x' \leq bx + (1-b)x/2 = ((1+b)/2)x$.

Plugging into (1), we deduce that for $\|r_t\|_2 \geq 4\varepsilon\sqrt{d/n}/(1-\lambda)$, we have

$$\|r_{t+1}\|_2 \leq \frac{1+\lambda}{2}\|r_t\|_2.$$

Consequently, after $t = O(\log_{1/b} n)$ steps, we have

$$\|r_t\|_1 \leq \sqrt{n}\|r_t\|_2 \leq \frac{4\varepsilon\sqrt{d}}{1-\lambda} = \frac{2\varepsilon_0\sqrt{d}}{\sqrt{d}-\lambda} \leq \frac{2\varepsilon_0\sqrt{d}}{\sqrt{d}/4} = 8\varepsilon_0.$$

Here we used that for $d \geq 2$, we have $\sqrt{d}-1 \geq \sqrt{d}/4$.

Thus, by choosing ε_0 a small enough constant, we can make the variation distance to π , which is $\|r_t\|_1/2$, an arbitrarily small constant. \blacktriangleleft

4 Hitting Time Robustness

4.1 Inverse of mixing time

We now show that the hitting time robustness is $\Omega(1/t_{\text{mix}})$. First recall the separation time t_{sep} and that $t_{\text{sep}} = \Theta(t_{\text{mix}})$ (Lemma 2.5). We first prove that short unbiased walks hit nodes with at least the “right” probability. Let $h_v = \mathbb{E}_\pi[H_v]$.

► **Lemma 4.1.** *For all $u, v \in V$, we have $\Pr_u[H_v \leq 2t_{\text{sep}}] \geq t_{\text{sep}}/(16h_v)$.*

Proof. Let

$$p \stackrel{\text{def}}{=} \Pr_\pi[H_v \leq t_{\text{sep}}] = \sum_{w \in V} \pi(w) \Pr_w[H_v \leq t_{\text{sep}}].$$

Since for any starting u , after t_{sep} steps every node w has probability at least $\pi(w)/2$, we have

$$\Pr_u[H_v \leq 2t_{\text{sep}}] \geq \sum_{w \in V} \frac{\pi(w)}{2} \Pr_w[H_v \leq t_{\text{sep}}] = \frac{p}{2}.$$

We now show that $p \geq t_{\text{sep}}/(8h_v)$, which will prove the lemma. Suppose that $p < t_{\text{sep}}/(8h_v)$. We will show that $\Pr_\pi[H_v \leq 2h_v] < 1/2$, which contradicts Markov’s inequality, proving the lemma.

To show this, let $k = 2h_v/t_{\text{sep}}$. We prove by induction that for all $i \leq k$, we have

$$\Pr_\pi[H_v \leq it_{\text{sep}}] < 2ip \leq 1/2.$$

This is true for $i = 1$. Suppose it’s true for a given i . Then

$$\Pr_\pi[H_v \leq (i+1)t_{\text{sep}}] = \Pr_\pi[H_v \leq it_{\text{sep}}] + \Pr_\pi[H_v \leq (i+1)t_{\text{sep}} \mid H_v > it_{\text{sep}}].$$

To analyze the conditional probability, we bound the conditional distribution of X_t given that $H_v > t$, where $t = it_{\text{sep}}$. Specifically, for any node w ,

$$\Pr_\pi[X_t = w \mid H_v > t] = \frac{\Pr_\pi[X_t = w \wedge H_v > t]}{\Pr_\pi[H_v > t]} \leq \frac{\Pr_\pi[X_t = w]}{\Pr_\pi[H_v > t]} = \frac{\pi(w)}{\Pr_\pi[H_v > t]}$$

118:18 Markov Chain Robustness

Thus, conditioning on $H_v > t$ can only increase probabilities by at most $1/\Pr_\pi[H_v > t] \leq 2$. Therefore,

$$\Pr_\pi[H_v \leq (i+1)t_{\text{sep}} \mid H_v > it_{\text{sep}}] \leq 2p,$$

and we are done. \blacktriangleleft

We are now ready to prove our theorem about hitting time robustness.

► **Theorem 4.2.** *For all ε and all nodes u, v , we have $h_{u,v}^+ = O(h_v \exp(O(\varepsilon t_{\text{mix}})))$.*

Proof. We prove the theorem with t_{sep} in place of t_{mix} ; this suffices since $t_{\text{sep}} = O(t_{\text{mix}})$. By Lemma 4.1, even with an ε -adversary, we have

$$q \stackrel{\text{def}}{=} \Pr_u[H_v^+ \leq 2t_{\text{sep}}] \geq \exp(-2t_{\text{sep}}\varepsilon)t_{\text{sep}}/(16h_v).$$

By Wald's identity, we have $h_{u,v}^+ \leq 2t_{\text{sep}}/q = O(h_v \exp(O(\varepsilon t_{\text{mix}})))$. \blacktriangleleft

Theorem 1.10, stating that the hitting time robustness $\rho_{\text{hit}} = \Omega(1/t_{\text{mix}})$, follows.

4.2 Hitting time robustness for trees

We now show that the hitting time robustness is $\Omega(1/\Delta)$ for trees, confirming Conjecture 1.16 in this case. We first restate the theorem.

► **Theorem 1.11.** *The hitting time robustness for any tree is $\Omega(1/\Delta)$.*

Fix a tree G with at least two edges (otherwise it's trivial), so $\Delta \geq 2$. Let $\varepsilon = 1/(2\Delta)$, and fix a standard ε -adversary. For the hitting time, we can assume without loss of generality that the adversary is oblivious, and hence the adversarial random walk P' has a stationary distribution π' . We first show that π' is not too different than the stationary distribution π of the uniform random walk P .

► **Lemma 4.3.** *For any node v , we have $\pi(v)/4 \leq \pi'(v) \leq 4\pi(v)$.*

Proof. First consider two neighbors v and w . Let $\pi'(v, w) = \pi'(v)P'(v, w)$ denote the stationary probability of the directed edge (v, w) . Since $\{v, w\}$ is a cut edge, we must have $\pi'(v, w) = \pi'(w, v)$. Therefore,

$$\frac{\pi'(v)}{\pi'(w)} = \frac{P'(w, v)}{P'(v, w)} \geq \frac{(1-\varepsilon)d_v}{(1+\varepsilon)d_w} \geq (1-2\varepsilon)\frac{d_v}{d_w}.$$

For v, w that may not be neighbors, an inductive argument gives a bound in terms of the distance $\Delta(v, w)$:

$$\frac{\pi'(v)}{\pi'(w)} \geq (1-2\varepsilon)^{\Delta(v, w)} \frac{d_v}{d_w} \geq \frac{d_v}{4d_w},$$

because $(1-1/x)^x \geq 1/4$ for $x \geq 2$. The lemma will follow because $\pi(v) = d_v/(2m)$ is exactly proportional to the degree. To elaborate, define the ratio $r(v) = \pi'(v)/\pi(v)$. Then

$$\frac{r(v)}{r(w)} = \frac{\pi'(v)}{\pi'(w)} \cdot \frac{\pi(w)}{\pi(v)} = \frac{\pi'(v)}{\pi'(w)} \frac{d_w}{d_v} \geq \frac{1}{4}.$$

Since some $r(u) \geq 1$ and some $r(u') \leq 1$, we have $1/4 \leq r(v) \leq 4$ for all v , as claimed. \blacktriangleleft

Proof of Theorem 1.11. As in the above lemmas, we set $\varepsilon = 1/(2\Delta)$ and fix a standard ε -adversary. We now show that $h^+(v, w) < 8h(v, w)$, which will prove the theorem. We consider the oblivious ε -adversary that maximizes the hitting time from v to w . It suffices to show $h^+(v, w) \leq 8h(v, w)$ for neighbors v and w , since for arbitrary v and w in a tree the hitting time is the sum of the hitting times between neighbors.

We now use the idea behind the “essential edge lemma” from Aldous-Fill (Lemma 5.1 in [1]). For neighbors v and w , consider the graph H obtained from G by removing all nodes and edges that are strictly closer to w than to v . Thus, in H the node w has degree 1, and $h(v, w)$ and $h^+(v, w)$ are the same in H as in G . Let π and π' denote the corresponding stationary distributions in H , and let $h(\cdot, \cdot)$ and $h^+(\cdot, \cdot)$ denote hitting times in H . Then since a random walk in H starting at w is forced to go to v , we have the following:

$$\begin{aligned} h(v, w) &= h(w, w) - 1 = 1/\pi(w) - 1 \geq 1/(2\pi(w)) && \text{(since } \pi(w) \leq 1/2) \\ h^+(v, w) &= h^+(w, w) - 1 = 1/\pi'(w) - 1 < 4/\pi(w) \leq 8h(v, w) && \text{(using Lemma 4.3)} \end{aligned}$$

The theorem follows. ◀

4.3 Negative result

► **Definition 4.4.** Fix a simple undirected graph $G = (V, E)$. Let $\Delta(v, w)$ denote the distance between v and w in G . Define the average distance $\bar{\Delta} = \mathbb{E}_{v,w}[\Delta(v, w)]$.

► **Theorem 4.5.** For every regular graph G and $\varepsilon \in [0, 1]$, there is a strong oblivious ε -adversary with $t_{\text{hit}}^+ = \Omega(ne^{\varepsilon\bar{\Delta}})$.

Proof. Much of this proof is similar to a lower bound on $\pi^+(S)$ by [4].

The idea is to construct a “lazy” adversary – allowing self loops – where some node has stationary probability at most $e^{-\varepsilon\bar{\Delta}}/n$, and then apply Lemma 2.1. The lazy adversary won’t be a valid adversary because the original graph doesn’t have self loops, so we then convert it into a valid adversary.

Let G be d -regular. Fix a node u such that $\mathbb{E}_w[\Delta(u, w)] \geq \bar{\Delta}$. We first define a lazy adversary using a Metropolis algorithm. The target stationary distribution ρ will be proportional to the weight defined as $\text{wt}(v) = (1 - \varepsilon)^{-\Delta(u, v)}$. That is, $\rho_v = \text{wt}(v)/z$, where $z = \sum_v \text{wt}(v)$. Since for an edge $\{v, w\}$ we have $|\Delta(u, v) - \Delta(u, w)| \leq 1$, the lazy adversary corresponds to the Metropolis Markov chain Q where

$$Q_{v,w} = \begin{cases} 1/d & \text{if } \{v, w\} \in E \text{ and } \Delta(u, w) \geq \Delta(u, v) \\ (1 - \varepsilon)/d & \text{if } \{v, w\} \in E \text{ and } \Delta(u, w) = \Delta(u, v) - 1 \\ 1 - \sum_{x \in \Gamma(v)} Q_{v,x} & \text{if } w = v \end{cases}$$

Since Q corresponds to a Metropolis algorithm for ρ , the stationary distribution of Q is ρ . Observe that

$$z = \sum_v (1 - \varepsilon)^{-\Delta(u, v)} = n \mathbb{E}_v[(1 - \varepsilon)^{-\Delta(u, v)}] \geq n(1 - \varepsilon)^{-\mathbb{E}_v[\Delta(u, v)]} \geq n(1 - \varepsilon)^{-\bar{\Delta}}.$$

Let $h^*(\cdot)$ denote the hitting times for Q . Since some neighbor v of u has $h^*(v, u) \geq h^*(u, u) - 1$, by Lemma 2.1 we have

$$t_{\text{hit}}^* \geq 1/\rho_u - 1 = z - 1 \geq n(1 - \varepsilon)^{-\bar{\Delta}} - 1.$$

However, Q is not a valid adversary, since a valid adversary must only place weight on edges of G , but G has no self loops. To circumvent this issue, distribute all the self loop probabilities proportional to the other edges. That is, let $Q_v = \sum_{x \in \Gamma(v)} Q_{v,x}$, and define

$$R_{v,w} = \begin{cases} (1 + Q_{v,v})/(dQ_v) & \text{if } \{v,w\} \in E \text{ and } \Delta(u,w) \geq \Delta(v,w) \\ (1 - \varepsilon)/(dQ_v) & \text{if } \{v,w\} \in E \text{ and } \Delta(u,w) = \Delta(v,w) - 1 \end{cases}$$

In other words R is a sped-up version of Q . A random trajectory of Q can be converted into a random trajectory of R by simply deleting consecutive repetitions. Thus, expected hitting times of R are at least $\min_v Q_v \geq 1 - \varepsilon$ times those of Q . Since R is a strong oblivious ε -adversary, we have

$$t_{\text{hit}}^+ \geq (1 - \varepsilon)t_{\text{hit}}^* \geq n(1 - \varepsilon)^{1 - \bar{\Delta}} - 1 = \Omega(ne^{\varepsilon\bar{\Delta}}). \quad \blacktriangleleft$$

► **Corollary 4.6.** *For every regular graph G and $\varepsilon \in [0, 1]$, there is a strong oblivious ε -adversary with $t_{\text{hit}}^+ = \Omega(ne^{(\varepsilon \log_d n)/2})$.*

Proof. Every d -regular graph has $\bar{\Delta} \geq (\log_d n)/2$. ◀

This enables us to conclude:

► **Theorem 1.12.** *The hitting time robustness of any constant degree expander is $\rho_{\text{hit}} = \Theta(1/\log n)$.*

Proof. The unbiased max hit time for any expander is $O(n)$ [21]. In order for t_{hit}^+ to be $O(n)$, we need $\varepsilon = \Omega(1/\log n)$. Also, $\varepsilon = \Theta(\log n)$ suffices because of Theorem 1.10 and the fact that constant degree expanders have $t_{\text{mix}} = \Theta(\log n)$. ◀

4.4 Tight analysis of path/cycle

Notice that Theorem 1.11 is mainly useful when the unbiased $t_{\text{hit}} = \Theta(n)$. We now give a tight analysis of the robustnesses of the path and cycle, which have $t_{\text{hit}} = \Theta(n^2)$. We focus on the hitting time robustness, as the others will be the same. We now show:

► **Theorem 4.7.** *The hitting time robustness of the path and cycle are $\Theta(1/n)$.*

Proof. From Theorem 1.11, we have $\rho_{\text{hit}} = \Omega(1/n)$. It therefore suffices to show the upper bound and negative result that $\rho_{\text{hit}} = O(1/n)$.

The idea is to use the essential edge lemma and its weighted version. It says that if $\{v, w\}$ is a cut edge, then $h(v, w) = 2e(G_v) + 1$, where G_v is the connected component of v if $\{v, w\}$ is removed. The weighted version simply counts the weights of the edges, and normalizes the edge $\{v, w\}$ to have weight 1 (or divide by $\text{wt}(v, w)$).

For simplicity we analyze the path on $n + 1$ nodes labeled 0 to n and edges $e_i = \{i, i + 1\}$. Consider the hitting time $h(n, 0) = h(n, n - 1) + h(n - 1, n - 2) + \dots + h(1, 0)$. Let's see how often edge e_i appears in $2e(G_v) + 1$, where the $+1$ gets charged to $\{v, w\}$. The contribution of edge e_i to $h(n, 0)$ is $1 + 2i$.

Now compare it when we bias the random walk by ε . We can do this by assigning $\text{wt}(e_i) = (1 + \varepsilon)^i$. The contribution of edge e_i to the biased hitting time $h^+(n, 0)$ is

$$1 + 2((1 + \varepsilon) + (1 + \varepsilon)^2 + \dots + (1 + \varepsilon)^{i-1}) \approx 2(1 + \varepsilon)^i/\varepsilon > 2\binom{i}{2}\varepsilon.$$

Thus the ratio of this to the contribution in the unbiased random walk is at least $\Omega(i\varepsilon)$. Now, over half the contribution comes from $i \geq n/2$, and for such i we can lower bound this ratio by $\Omega(n\varepsilon)$. Therefore, if $\varepsilon = \omega(1/n)$, then the ratio $h^+(n, 0)/h(n, 0) = \omega(1)$, as we wanted. ◀

5 Cover Time Robustness

In the non-adversarial setting, Matthews [19] upper bounded the cover time in terms of the max hitting time:

► **Theorem 5.1** ([19]). $t_{\text{cov}} \leq t_{\text{hit}}(1 + 1/2 + \dots + 1/n) \leq t_{\text{hit}}(1 + \ln n)$.

There are many examples where this is tight up to a constant factor: expanders (and hence most graphs) [5], the complete graph (coupon collecting), two and higher dimensional grids and tori [2, 25], and the balanced binary tree [24].

We observe that in any graph with $t_{\text{cov}} = \Theta(t_{\text{hit}} \log n)$, we have $\rho_{\text{cov}} = \Theta(\rho_{\text{hit}})$. This is because Matthews' proof works the same in the adversarial setting.

► **Theorem 5.2.** *For any adversary, $t_{\text{cov}}^+ \leq t_{\text{hit}}^+(1 + \ln n)$.*

Proof. Matthews [19] proof works in the adversarial setting, so

$$t_{\text{cov}}^+ \leq t_{\text{hit}}^+(1 + 1/2 + \dots + 1/n) \leq t_{\text{hit}}^+(1 + \ln n). \quad \blacktriangleleft$$

► **Corollary 5.3.** *In any graph with $t_{\text{cov}} = \Theta(t_{\text{hit}} \log n)$, we have $\rho_{\text{cov}} = \Omega(\rho_{\text{hit}})$.*

We conjecture that more generally, in any graph, $\rho_{\text{cov}} = \Theta(\rho_{\text{hit}})$.

References

- 1 D. Aldous and J. A. Fill. Reversible Markov chains and random walks on graphs, 2002. Unfinished monograph, recompiled 2014, available at <http://www.stat.berkeley.edu/~aldous/RWG/book.html>.
- 2 D.J. Aldous. On the time taken by random walks on finite groups to visit every state. *Z. Wahrscheinlichkeitstheorie verw Gebiete*, 62:361–374, 1983.
- 3 C. Avin, M. Koucky, and Z. Lotker. Cover time and mixing time of random walks on dynamic graphs. *Random Structures and Algorithms*, 52:576–596, 2018. doi:10.1002/RSA.20752.
- 4 Y. Azar, A. Z. Broder, A. R. Karlin, N. Linial, and S. J. Phillips. Biased random walks. *Combinatorica*, 16:1–18, 1996. doi:10.1007/BF01300124.
- 5 A.Z. Broder and A.R. Karlin. Bounds on the cover time. *Journal of Theoretical Probability*, 2:101–120, 1989.
- 6 S. Cannon, J. J. Daymude, D. Randall, and A.W. Richa. A Markov chain algorithm for compression in self-organizing particle systems. In *Proceedings of the 2016 Symposium on Principles of Distributed Computing*, pages 279–288, 2016.
- 7 B. Chin, A. Moitra, E. Mossel, and C. Sandon. The power of an adversary in Glauber dynamics. In *The Thirty Seventh Annual Conference on Learning Theory*, volume 247 of *Proceedings of Machine Learning Research*, pages 1102–1124, 2024. URL: <https://proceedings.mlr.press/v247/chin24a.html>.
- 8 R.V. Craiu, L. Gray, K. Latuszynski, N. Madras, G.O. Roberts, and J.S. Rosenthal. Stability of adversarial Markov chains, with an application to adaptive MCMC algorithms. *Annals of Applied Probability*, 25:3592–3623, 2015.
- 9 D. Doron, D. Moshkovitz, J. Oh, and D. Zuckerman. Almost Chor-Goldreich sources and adversarial random walks. In *Proceedings of the 55th Annual ACM Symposium on Theory of Computing*, pages 1–9, 2023.
- 10 D. Doron, D. Moshkovitz, J. Oh, and D. Zuckerman. Online condensing of unpredictable sources via random walks. In *Proceedings of the 40th Computational Complexity Conference*, volume 339 of *LIPICs*, pages 30:1–30:17, 2025. doi:10.4230/LIPICs.CCC.2025.30.
- 11 D.J. Hartfiel. On the limiting set of stochastic products $xA_1 \dots A_k$. *Proceedings of the American Mathematical Society*, 81:201–206, 1981.

118:22 Markov Chain Robustness

- 12 D.J. Hartfiel. *Markov Set-Chains*, volume 1695 of *Lecture Notes in Mathematics*. Springer-Verlag, 1998.
- 13 J. Haslegrave, T. Sauerwald, and J. Sylvester. Time dependent biased random walks. *ACM Transactions on Algorithms*, 18, 2022.
- 14 J.J. Hunter. Mixing times with applications to perturbed markov chains. *Linear Algebra and Its Applications*, 417:108–123, 2006.
- 15 S. Kania, D. Aristoff, and D.M. Zuckerman. Riteweight: Randomized iterative trajectory reweighting for steady-state distributions without discretization error. Technical report, Arxiv, 2024. arxiv:2401.05597.
- 16 D.A. Levin and Y. Peres. *Markov Chains and Mixing Times*. American Mathematical Society, 2nd edition, 2017.
- 17 S. Li, B. Dutta, S. Cannon, J. J. Daymude, R. Avinery, E. Aydin, A. W. Richa, D. I. Goldman, and D. Randall. Programming active cohesive granular matter with mechanically induced phase changes. *Science Advances*, 7:eabe8494, 2021.
- 18 Y. Liu. Perturbation bounds for the stationary distributions of Markov chains. *SIAM Journal on Matrix Analysis and Applications*, 33:1057–1074, 2012. doi:10.1137/110838753.
- 19 P. Matthews. Covering problems for Markov chains. *Annals of Probability*, 16:1215–1228, 1988.
- 20 L. Page, S. Brin, R. Motwani, and T. Winograd. The pagerank citation ranking: Bringing order to the web. Technical Report Technical Report 1999-66, Stanford InfoLab, 1999.
- 21 R. Rubinfeld. The cover time of a regular expander is $O(n \log n)$. *Information Processing Letters*, 35:49–51, 1990. doi:10.1016/0020-0190(90)90173-U.
- 22 T. Sauerwald and L. Zanetti. Random walks on dynamic graphs: Mixing times, hitting times, and return probabilities. In *Proceedings of the 46th International Colloquium on Automata, Languages, and Programming*, volume 132 of *LIPICs*, pages 93:1–93:15, 2019. doi:10.4230/LIPICs.ICALP.2019.93.
- 23 D. Vial and V.G. Subramanian. Restart perturbations for reversible markov chains: Trichotomy and pre-cutoff equivalence. *Random Structures and Algorithms*, 66, 2025.
- 24 D. Zuckerman. Covering times of random walks on bounded degree trees and other graphs. *Journal of Theoretical Probability*, 2:147–157, 1989.
- 25 D. Zuckerman. A technique for lower bounding the cover time. *SIAM Journal on Discrete Mathematics*, 5:254–259, 1992.