

On coalescence time in graphs

When is coalescing as fast as meeting?

Varun Kanade*
University of Oxford
varunk@cs.ox.ac.uk

Frederik Mallmann-Trenn†
MIT
mallmann@mit.edu

Thomas Sauerwald‡
University of Cambridge
thomas.sauerwald@cl.cam.ac.uk

Abstract

Coalescing random walks is a fundamental stochastic process, where a set of particles perform independent discrete-time random walks on an undirected graph. Whenever two or more particles meet at a given node, they merge and continue as a single random walk. The *coalescence time* is defined as the expected time until only one particle remains, starting from one particle at every node. Despite recent progress such as by Cooper, Elsässer, Ono, Radzik [CEOR13] and Cooper, Frieze and Radzik [CFR09a], the coalescence time for graphs such as binary trees, d -dimensional tori, hypercubes and more generally, vertex-transitive graphs, remains unresolved.

We provide a powerful toolkit that results in tight bounds for various topologies including the aforementioned ones. The meeting time is defined as the worst-case expected time required for two random walks to arrive at the same node at the same time. As a general result, we establish that for graphs whose meeting time is only marginally larger than the mixing time (a factor of $\log^2 n$), the coalescence time of n random walks equals the meeting time up to constant factors. This upper bound is complemented by the construction of a graph family demonstrating that this result is the best possible up to constant factors. For almost-regular graphs, we bound the coalescence time by the hitting time, resolving the discrete-time variant of a conjecture by Aldous for this class of graphs. Finally, we prove that for any graph the coalescence time is bounded by $O(n^3)$ (which is tight for the Barbell graph); surprisingly even such a basic question about the coalescing time was not answered before this work. By duality, our results give bounds on the voter model and therefore give bounds on the consensus time in arbitrary undirected graphs.

We also establish a new bound on the hitting time and cover time of regular graphs, improving and tightening previous results by Broder and Karlin [BK89], as well as those by Aldous and Fill [AF02].

keywords— coalescing time, meeting time, random walks, voter model

*This work was supported in part by The Alan Turing Institute under the EPSRC grant EP/N510129/1.

†This work was carried out while a student at the École normale supérieure and Simon Fraser University. This work was supported in part by NSF Award Numbers BIO-1455983, CCF-1461559, and CCF-0939370.

‡This work was supported by the ERC Starting Grant (DYNAMIC MARCH).

Contents

1	Introduction	1
1.1	Contributions	3
1.2	Proof Ideas and Technical Contributions	6
1.3	Discussion and Future Work	10
2	Notation and Preliminaries	10
3	Bounding t_{coal} for large $t_{\text{meet}}/t_{\text{mix}}$	11
3.1	Stochastic Process	11
3.2	Meeting Time Distribution Prior to t_{meet}	16
3.3	Upper Bound - Proof of Theorem 1.1	18
3.4	Lower Bound - Proof of Theorem 1.2	23
4	Bounding $t_{\text{coal}} = O(t_{\text{hit}})$ for Almost-Regular Graphs	33
4.1	Concentration Inequalities for Random Walks	34
4.2	Reducing the Walks from n to $O(\log^3 n)$ in $O(t_{\text{hit}})$	36
4.3	Reducing the Walks from $\log^4 n$ to $(\Delta/d)^{O(1)}$ in $O(t_{\text{hit}})$	39
4.3.1	Proof Overview	39
4.3.2	Definitions and Lemmas required to prove Theorem 4.4	40
4.3.3	Completing the Proof of Theorem 4.4	49
4.4	Bounding t_{coal} in terms of t_{hit}	49
4.5	Proof of Theorem 1.3	50
4.6	Conjecture and a Possible Improvement for Non-Regular Graphs	50
5	Bounding $t_{\text{coal}} \in [\Omega(\log n), O(n^3)]$	51
5.1	General Upper Bound $t_{\text{coal}} = O(n^3)$	51
5.2	General Lower Bound $t_{\text{coal}} = \Omega(\log n)$	52
5.3	Proof of Theorem 1.5	52
A	Basic Results about Markov Chains	55
B	Bounding t_{meet} and Implications for $t_{\text{coal}}, t_{\text{hit}}$ and t_{cov}	60
B.1	Relating Meeting Time to t_{mix} and $\frac{1}{1-\lambda_2}$	60
B.2	Relating Meeting Time to t_{hit}	68
B.3	Proof of Theorem 1.4	69
C	Applications to Concrete Topologies	70
C.1	2-Dimensional Grids/Tori and Paths/Cycles	70
C.2	d -Dimensional Grids and Tori, $d \geq 3$	70
C.3	Hypercubes	70
C.4	(Regular) Expanders	71
C.5	Real World Graph Models	71
C.6	Binary Trees	72
C.6.1	Bounds on the t -step probabilities	72
C.6.2	Establishing the Lower Bound on the Meeting Time	74
C.7	Star	76

1 Introduction

Coalescing random walks is a fundamental stochastic process on *connected* and *undirected* graphs. The process begins with particles on some subset of the nodes in the graph. At discrete time-steps, every particle performs one step of an independent random walk.¹ Whenever two or more particles arrive at the same node at the same time-step, they merge into a single particle and continue as a single random walk. The *coalescence time* is defined as the first time-step when only one particle remains. The coalescence time depends on the number and starting positions of the particles.

Studying the coalescence time is of substantial importance in distributed computing: At the heart of many distributed computing applications lie consensus protocols and leader election *e.g.*, data consistency, consolidation of replicated states, synchronization of processes and devices [Pel02; DGM+11] and communication networks [PVV09]). Other applications of the coalescence process appear in robotics [GORN17]; here, robots perform random walks to gather samples from their environment and need to communicate these samples to all other robots. Studying the coalescence time also implies results for other interaction types of random walks including predator and prey particles as well as annihilating particles [CFR09b].

Relationship to consensus protocols Arguably the simplest consensus protocol achieving consensus on any undirected graph is the voter model. Initially, every node has a distinct opinion. At every round, each node chooses synchronously one of its neighbors at random and adopts that node’s opinion. The *consensus time* is defined as the time it takes until only one opinion remains. The voting process viewed backwards is exactly the same as the coalescence process starting with a random walk on every node; thus, the coalescence time and consensus time have the same distribution. Despite recent progress by Cooper *et al.* [CEOR13; CFR09a] and Berenbrink *et al.* [BGKM16], the coalescence time and consensus time are far from being well-understood—even for certain fundamental graphs as we describe below. Recently, there have been several studies on variants of the voter model, most notably 2-Choices and 3-Majority which received ample attention [CER14; BCN+15; CER+15; BCN+16; CRRS16; EFK+16; BCE+17; GL17]. However, the behavior of these processes is fundamentally different and despite their efficiency in reaching consensus on expanders cliques, they are unsuitable on more general undirected graphs as the consensus time is exponential in some graphs.

In this paper, we follow the approach of Cooper *et al.* [CEOR13] and Hassin and Peleg [HP01] and study the consensus time through the more tangible analysis of the coalescence time. When starting with two particles, the coalescence time is referred to as the *meeting time*. Let t_{meet} denote the worst-case expected meeting time over all pairs of starting nodes and let t_{coal} denote the expected coalescence time starting from one particle on every node. It is clear that $t_{\text{meet}} \leq t_{\text{coal}}$; as for an upper bound, it can be shown that $t_{\text{coal}} = O(t_{\text{meet}} \log n)$, where n is the number of nodes in the graph. The main idea used to obtain the bound is that the number of surviving random walks halves roughly every t_{meet} steps. A proof of the result appears implicitly in the work of Hassin and Peleg [HP01].

Aldous [Ald91] showed in *continuous-time* that the meeting time is bounded by the maximum hitting time, $t_{\text{hit}} := \max_{u,v} t_{\text{hit}}(u, v)$, where $t_{\text{hit}}(u, v)$ denotes the expected time required to hit v starting from vertex u . We observe that the result of Aldous also holds in discrete time. Thus, this gives a bound of $O(t_{\text{hit}} \log n)$ for the coalescing time; however, in general $O(t_{\text{hit}})$ may be a loose upper bound on t_{meet} . In recent work, Cooper *et al.* [CEOR13] provide results that are better

¹Throughout this paper, we use random walk and particle interchangeably, assuming that every random walk has an identifier.

than $O(t_{\text{meet}} \log n)$ for several interesting graph classes, notably expanders and power-law graphs. They show that $t_{\text{coal}} = O((\log^4 n + \|\pi\|_2^{-2}) \cdot (1 - \lambda_2)^{-1})$, where λ_2 is the second largest eigenvalue of the transition matrix of the random walk and π is the stationary distribution. Berenbrink *et al.* [BGKM16] show that $t_{\text{coal}} = O(m/(d_{\min} \cdot \Phi))$, where m is the number of edges, d_{\min} is the minimum degree and Φ is the conductance. Their result improves on that of Cooper *et al.* for certain graph classes, *e.g.*, cycles.

As mentioned before, despite the recent progress due to Cooper *et al.* [CEOR13] and Berenbrink *et al.* [BGKM16], for many fundamental graphs such as the binary tree, hypercube and the (d -dimensional) torus, the coalescing time in the discrete setting remains unsettled. We provide a rich toolkit allowing us to derive tight bounds for many graphs including all of the aforementioned ones. One of our main results establishes a relationship between the ratios $t_{\text{coal}}/t_{\text{meet}}$ and $t_{\text{meet}}/t_{\text{mix}}$, where $t_{\text{mix}} = t_{\text{mix}}(1/e)$ denotes the mixing time.² In particular, the result shows that if $t_{\text{meet}}/t_{\text{mix}} = \Omega(\log^2 n)$, then $t_{\text{coal}} = O(t_{\text{meet}})$; however, we also provide a more fine-grained tradeoff. For almost-regular graphs,³ we bound the coalescence time by the hitting time. For vertex-transitive graphs we show that the coalescence time, the meeting time, and the hitting time are equal up to constant factors. Finally, we prove that for any graph the coalescence time is bounded by $O(n^3)$; it can be easily verified that this is tight by considering the barbell graph. Surprisingly, the right bound on this fundamental quantity was not known prior to this work. Unlike in the analogous case of the cover time [AF02] where such a bound can be easily derived, the argument in the case of coalescence time appears significantly involved.⁴ Prior to this work, Hassin and Peleg [HP01] had shown a worst-case upper bound of $O(n^3 \log n)$. We also give worst-case upper and lower bounds on the meeting time and coalescence time that are tight for general graphs and regular (or nearly-regular) graphs.

In the process of establishing bounds on the coalescence time, we develop techniques to give tight bounds on the meeting time. We apply these to various topologies such as the binary tree, torus and hypercube. We believe that these techniques might be of more general interest.

The process of coalescing random walks was first studied in *continuous time*; in this case, particles jump to a random neighboring node when activated according to a Poisson clock with mean 1. As Cooper and Rivera [CR16] recently pointed out “*It is however, not clear whether the continuous-time results apply to the discrete-time setting*”, and to the best of our knowledge, there is no general way in which results in *continuous time* can be transferred to *discrete time* or vice versa, even when the random walks in discrete-time are *lazy*. In the continuous time setting, Cox [Cox89] show that the coalescence time is bounded by $\Theta(t_{\text{hit}})$ for tori. Oliveira [Oli12] showed that the coalescence time is $O(t_{\text{hit}})$ in general. In a different work, Oliveira [Oli13] derived so-called mean field conditions, which are sufficient conditions for the coalescing process on a graph to behave similarly to that on the complete graph up to scaling by the expected meeting time. His main result (for non vertex-transitive graphs) in [Oli13, Theorem 1.2], implies that $t_{\text{coal}} = O(t_{\text{meet}})$ whenever $t_{\text{mix}} \cdot \pi_{\max} = O(1/\log^4 n)$. One of our main results, Theorem 1.1, implies $t_{\text{coal}} = O(t_{\text{meet}})$ whenever $t_{\text{mix}}/t_{\text{meet}} = O(1/\log^2 n)$. Notice that since $t_{\text{meet}} \geq 1/(\|\pi\|_2^2) \geq 1/\pi_{\max}$, our condition is considerably more general—however, the results in [Oli13] also establish mean-field behavior (that is, when suitably scaled, the distribution of the coalescence time is similar to that on a complete graph), while ours are only concerned with the expected coalescence time, t_{coal} . On the other hand, our result also applies to graphs where $t_{\text{coal}} \gg t_{\text{meet}}$ such as the star graph, and together

²The *mixing time* is the first time-step at which the distribution of a random walk starting from an arbitrary node is close to the stationary distribution.

³We call a graph *almost-regular* if $\deg(u) = \Theta(\deg(v))$ for all $u, v \in V$.

⁴Cooper *et al.* [CEOR13] mistakenly stated, as a side remark, that this last result was a simple consequence of their main result.

Graph	t_{mix}	t_{meet}	t_{coal}	t_{hit}
Binary tree	$\Theta(n)$	$\Theta(n \log n)$	<i>Thm. 1.3 & Thm. C.6</i>	$\Theta(n \log n)$
Clique	$\Theta(1)$	$\Theta(n)$	[CEOR13; BGKM16] & <i>Thm. 1.1</i>	$\Theta(n)$
Cycle	$\Theta(n^2)$	$\Theta(n^2)$	[BGKM16] & <i>Thm. 1.3</i>	$\Theta(n^2)$
Rand. r -reg.	$\Theta(\log n)$	$\Theta(n)$	[CFR09a; CEOR13; BGKM16] & <i>Thm. 1.1</i>	$\Theta(n)$
Hypercube	$\Theta(\log n \log \log n)$	$\Theta(n)$	<i>Thm. 1.3</i>	$\Theta(n)$
Path	$\Theta(n^2)$	$\Theta(n^2)$	[BGKM16] & <i>Thm. 1.3</i>	$\Theta(n^2)$
Star	$\Theta(1)$	$\Theta(1)$	folklore	$\Theta(n)$
Torus ($d = 2$)	$\Theta(n)$	$\Theta(n \log n)$	<i>Thm. 1.3</i>	$\Theta(n \log n)$
Torus ($d > 2$)	$\Theta(n^{2/d})$	$\Theta(n)$	<i>Thm. 1.1</i>	$\Theta(n)$

Table 1: A summary of bounds on the mixing, meeting, coalescence and hitting times for fundamental topologies for discrete-time random walks. All bounds on the mixing and hitting times appear directly or implicitly in [AF02].

with Theorem 1.2, demonstrate that the trade-off between meeting and mixing time is the best possible.

1.1 Contributions

In this work, we provide several results relating the coalescence and meeting times to each other and to other fundamental quantities of random walks on undirected graphs. In particular, our focus is on understanding for which graphs the coalescence time is the same as the meeting time, as we know that t_{coal} is always in the rather narrow interval of $[t_{\text{meet}}, O(t_{\text{meet}} \cdot \log n)]$. As a consequence of our results, we derive new and re-derive existing bounds on the meeting and coalescence times for several graph families of interest. These results are summarized in Table 1 and discussed in greater detail in Appendix C. Formal definitions of all quantities used below appear in Section 2. Throughout this paper, we assume that random walks are *lazy* meaning that w.p. 1/2 the walk stays put.

Our first main result relates t_{coal} to t_{meet} and t_{mix} . As already mentioned in the introduction, the crude bound $t_{\text{coal}} = O(t_{\text{meet}} \log n)$ is well-known. However, this bound is not in general tight, as demonstrated by our result below.

Theorem 1.1. *For any graph G , we have*

$$t_{\text{coal}} = O\left(t_{\text{meet}} \left(1 + \sqrt{\frac{t_{\text{mix}}}{t_{\text{meet}}}} \cdot \log n\right)\right),$$

Consequently, when $t_{\text{meet}} \geq t_{\text{mix}} \log^2 n$, $t_{\text{coal}} = O(t_{\text{meet}})$.

The proof of Theorem 1.1 appears in Section 3. One interesting aspect about this bound is that it can be used to establish $t_{\text{coal}} = \Theta(t_{\text{meet}})$ even without having to know the quantities t_{meet} or t_{mix} . This flexibility turns out to be particularly useful when dealing with random graph models for “real world” networks, where we establish (nearly-)tight and sublinear bounds (w.r.t. to the number of vertices) in Section C.5.

Another interesting feature of our theorem is that the main result of Cooper *et al.* [CEOR13, Theorem 1] can be reproven by combining [CEOR13, Theorem 2] with Theorem 1.1 (see Proposition B.2).

Our next main result shows that the bound in [Theorem 1.1](#) is tight up to a constant factor, which we establish by constructing an explicit family of graphs. Interestingly, for this family of almost-regular graphs we also have $t_{\text{hit}} \gg t_{\text{meet}}$, thus showing that t_{hit} may be a rather loose upper bound for t_{coal} in some cases.⁵

Theorem 1.2. *For any sequence $(\alpha_n)_{n \geq 0}$, $\alpha_n \in [1, \log^2 n]$ there exists a family of almost-regular graphs (G_n) , with G_n having $\Theta(n)$ nodes and satisfying $\frac{t_{\text{meet}}}{t_{\text{mix}}} = \Theta(\alpha_n)$ such that*

$$t_{\text{coal}} = \Omega \left(t_{\text{meet}} \cdot \left(1 + \sqrt{\frac{t_{\text{mix}}}{t_{\text{meet}}}} \cdot \log n \right) \right).$$

The above two results show that that $t_{\text{meet}}/t_{\text{mix}}$ should be $\Omega(\log^2 n)$ to guarantee that $t_{\text{coal}} = O(t_{\text{meet}})$.

A natural question is therefore whether in the case of structured sub-classes such as regular graphs, or vertex-transitive graphs, or special graphs such as grids, tori, binary trees, cycles, real-world (power-law) graphs, *etc.*, better bounds can be obtained through other methods. We provide results that are tight or nearly tight in several of these cases; some of these results were previously known using other methods, some are novel to the best of our knowledge.

Theorem 1.3. *The following hold for graphs of the stated kind*

(i) *For any graph G ,*

$$t_{\text{coal}} = O(t_{\text{hit}} \cdot \log \log n).$$

(ii) *For any graph G with maximum degree Δ and average degree d ,*

$$t_{\text{coal}} = O(t_{\text{hit}} + t_{\text{meet}} \cdot \log(\Delta/d)).$$

Hence for any almost-regular graph G , $t_{\text{coal}} = O(t_{\text{hit}})$.

(iii) *For any vertex-transitive G ,*

$$t_{\text{coal}} = \Theta(t_{\text{meet}}) = \Theta(t_{\text{hit}}).$$

(iv) *In the case of binary trees, d -dimensional tori/grids, paths/cycles, expanders, hypercubes, random power law graphs,⁶ we have $t_{\text{coal}} = \Theta(t_{\text{meet}})$.*

The proof of the first three statements of [Theorem 1.3](#) appear in [Section 4](#) and the last statement follows from the results in [Appendix C](#). We point out that since $t_{\text{meet}} = O(t_{\text{hit}})$ for any graph,⁷ [Theorem 1.3](#) implies the bound $t_{\text{coal}} = O(t_{\text{hit}})$ not only for almost-regular graphs, but also for dense graphs where $|E| = \Theta(n^2)$. This settles the discrete-time analogue of a conjecture by Aldous [[AF02](#), Open Problem 14.13] for these graph classes. It remains open whether $t_{\text{coal}} = O(t_{\text{hit}})$ holds for all undirected graphs.

Another natural question is to express t_{meet} or t_{coal} solely in terms of t_{mix} , the spectral gap $1 - \lambda_2$ or other connectivity properties of G . We derive several such bounds on t_{meet} , t_{hit} and t_{coal} .

As a by-product of our techniques, we also derive new bounds on t_{hit} and t_{cov} , the cover-time. The detailed results are given in [Appendix B](#), but we highlight the results for regular graphs here:

⁵Note that the star also exhibits $t_{\text{hit}} \gg t_{\text{meet}}$. However, the star is not almost-regular.

⁶The exact model is specified in [Section C.5](#).

⁷In [Proposition B.9](#), we prove this formally by following the proof for the continuous setting [[AF02](#), Proposition 14.5].

Theorem 1.4. *Let G be any graph with $\Gamma = \Delta/\delta$, where Δ is the maximum degree and δ the minimum degree. It holds that*

$$t_{\text{hit}} = O(\Gamma n / \sqrt{1 - \lambda_2}) = O(\Gamma n / \Phi),$$

where Φ is the conductance of the graph and λ_2 is the second largest eigenvalue of the transition matrix P of a lazy random walk. Consequently, $t_{\text{meet}} \leq t_{\text{coal}} = O(\Gamma n \log(\Gamma) / \sqrt{1 - \lambda_2}) = O(\Gamma n \log(\Gamma) / \Phi)$ and $t_{\text{cov}} = O(\Gamma n \log n / \Phi)$.

We point out that so far the best possible bound on t_{coal} for regular graphs has been $t_{\text{coal}} = O(n/(1 - \lambda_2))$ from [CEOR13].⁸ The best possible bound on t_{hit} (and t_{cov}) in terms of $1 - \lambda_2$, was $t_{\text{hit}} = O(n/(1 - \lambda_2))$ and $t_{\text{cov}} = O(n \log n / (1 - \lambda_2))$ due to Broder and Karlin [BK89] from 1989. In all four cases, t_{meet} , t_{coal} , t_{hit} , and t_{cov} , Theorem 1.4 improves the dependency on $1/(1 - \lambda_2)$ (or, equivalently t_{mix}), by almost a square-root (we refer the reader to Theorem B.6 and Theorem B.8 for further details). As a result of this improvement, we get a bound of $O(n/\Phi)$ on the hitting time which is the best known bound on the hitting time (and cover time) in terms of the conductance and improves the bound of [AF02, Corollary 6.2.1] by a factor of $1/\Phi$.

We also derive a general lower bound on t_{meet} that combines the trivial bound, $1/\|\pi\|_2^2$, with the minimum number of collisions (see Theorem B.1.(iii)). Although this bound does not directly yield the correct lower bound for binary trees, it forms the basis of a later analysis in Theorem C.6.

Finally, we also provide asymptotically tight worst-case bounds on t_{meet} and t_{coal} . We show that on any graph the coalescence time must be at least $\Omega(\log n)$ and is no more than $O(n^3)$. For regular (and in particular vertex-transitive) graphs these bounds become $\Omega(n)$ and $O(n^2)$ (See also Table 2 on page 51, which also contains an explanation why these bounds are asymptotically tight.) These two new upper bounds for general and regular graphs complete the picture of worst-case bounds:

Theorem 1.5. *The following hold for graphs of the stated kind.*

- (i) *For any graph G we have $t_{\text{meet}} \in [\Omega(1), O(n^3)]$ and $t_{\text{coal}} \in [\Omega(\log n), O(n^3)]$.*
- (ii) *For any regular (including vertex-transitive) graph G we have $t_{\text{meet}}, t_{\text{coal}} \in [\Omega(n), O(n^2)]$.*

The proof of Theorem 1.5 appears in Section 5.

Summary of Technical Contributions

Our work also makes several technical contributions, which might be of interest for future research on coalescing walks and other stochastic processes; these are explained in greater detail in Section 1.2. Below we give a very brief summary.

- **Conditional Expectation Approach.** Most of our results make use of the conditional expectation approach given in (1), a very simple yet extremely powerful tool, which to the best of our knowledge has not been used in the context of meeting and coalescing times before.
- **Division of Particles into two Groups.** One basic ingredient in our proof is a domination result that allows us to divide random walks into a group of “destroyers” (\mathcal{G}_1), which are particles that cannot be eliminated, and a group of remaining particles (\mathcal{G}_2), which can be eliminated by any other random walk. This domination result might be helpful to analyse other stochastic processes involving different types of particles, e.g. [CFR09a].

⁸Alternatively, the same bound as the known bound can also be derived from the bound on the conductance in [BGKM16] together with Cheeger’s inequality.

- **New Concentration Inequalities.** We derive a new concentration inequality for random walks on graphs in [Section 4.1](#). Unlike previous approaches which are based on the mixing time (or the closely related spectral gap), our new inequality depends only on the hitting time and improves on the existing bounds when the mixing time is close to the hitting time. These tighter inequalities are required to derive worst-case upper bounds on the coalescence time.

1.2 Proof Ideas and Technical Contributions

When dealing with processes involving concurrent random walks, a significant challenge is to understand the behavior of “short” random walks. This challenge appears in several settings, *e.g.*, in the context of cover time of multiple random walks [[AAK+11](#); [ER09](#)], where Efremenko and Reingold [[ER09](#), Section 6] highlight the difficulty in analyzing the hitting time distribution before its expectation. In the context of concentration inequalities for Markov chains, Lezaud [[Lez89](#), p. 863] points out the requirement to spend at least mixing time steps before taking any samples. Related to that, in property testing, dealing with graphs that are far from expanders has been mentioned as one of the major challenges to test the expansion of the graph by Czumaj and Sohler [[CS10](#)].

In our setting, we also face these generic problems and devise different methods to get a handle on the meeting time distribution before its expectation. Despite our focus being on coalescing and meeting times, several of our approaches can be leveraged to derive new bounds on other random walk quantities such as hitting times or cover times (see [Appendix B](#)).

Bounds on t_{coal} in terms of t_{mix} and t_{meet}

The key ingredient in the proof of [Theorem 1.1](#), where we express t_{coal} as a tradeoff between t_{meet} and t_{mix} is a better understanding of meeting events prior to the meeting time. More precisely, we derive a tight bound on the probability p_ℓ that two random walks meet before ℓ time-steps, for ℓ in the range $[t_{\text{mix}}, t_{\text{meet}}]$. Arguing about meeting probabilities of walks that are much shorter than t_{meet} allows us to understand the rate at which the number of *alive* random walks is decreasing.

Optimistically, one may hope that starting with k random walks, as there are $\binom{k}{2}$ possible meeting events, roughly $\binom{k}{2} \cdot p_\ell$ meetings may have occurred after ℓ time-steps. However, the non-independence of these events turns out to be a serious issue and we require a significantly more sophisticated approach to account for the dependencies. We divide the k random walks into disjoint groups \mathcal{G}_1 and \mathcal{G}_2 (with $|\mathcal{G}_1|$ usually being much smaller than $|\mathcal{G}_2|$) and walks of \mathcal{G}_1 can’t be eliminated. The domination of the real process by the group-restricted one is established by introducing a formal concept called immortal process at the beginning of [Section 3.1](#). In this stochastic process, we can expose the random walks of \mathcal{G}_1 first and consider meetings with random walks in \mathcal{G}_2 (for an illustration, see [Figure 2](#) on page 12). Conditioning on a specific exposed walk in \mathcal{G}_1 , the events of the different walks in \mathcal{G}_2 meeting this exposed walk are indeed independent. In fact, we will also use the symmetric case where the roles of \mathcal{G}_1 and \mathcal{G}_2 are switched. Thus, the problem then reduces to calculating the probability of a random walk in \mathcal{G}_2 having a ‘good trajectory’, *i.e.*, one which many random walks in \mathcal{G}_1 would meet with large enough probability.

Surprisingly, it suffices to divide trajectories into only two categories ([Lemma 3.3](#)). Although, one may expect that a more fine-grained classification of trajectories would result in better bounds, this turns out not to be the case. In fact, the bound that we derive on the coalescing time in [Theorem 1.1](#) is tight, and this is precisely due to the tightness of [Lemma 3.3](#). The tightness is established by the following construction (cf. [Figure 1](#)). The graph is designed such that the vast majority of meetings (between any two random walks) occur in a relatively small part of the graph




Figure 1: The graph described in [Section 3.4](#) with $t_{\text{coal}} = \Omega(t_{\text{meet}} + \sqrt{t_{\text{meet}}/t_{\text{mix}}} \cdot \log n \cdot t_{\text{mix}})$.

(G_2 in [Figure 1](#)). On average, it takes a considerable number of time-steps before random walks actually get to this part of the graph. What this implies is that for relatively short trajectories (of length significantly smaller than t_{meet}), it is quite likely that other random walks will not meet them (cf. [Lemma 3.3](#)). There is a bit of a dichotomy here, once a walk reaches G_2 it is likely that many random walks will meet it; however, a random walk not reaching G_2 is unlikely to be met by any other random walk.

Equipped with [Theorem 1.1](#), we can bound $t_{\text{coal}} = \Theta(t_{\text{meet}})$ for all graphs satisfying $t_{\text{meet}}/t_{\text{mix}} \geq \log^2 n$. Therefore, the problem of bounding t_{coal} reduces to bounding t_{meet} .

For some of the other results including [Theorem 1.2](#) and [Theorem 1.3](#), we will need a more fine-grained approach to derive lower (or upper bounds) on the probability that two walks meet during a certain number of steps, which may or may not be smaller than the mixing time or meeting time. The starting point is the following simple observation. If we have two random walks $(X_t)_{t \geq 0}$ and $(Y_t)_{t \geq 0}$, and count the number of collisions $Z := \sum_{t=0}^{\tau-1} \mathbf{1}_{X_t=Y_t}$ before time-step $\tau \in \mathbb{N}$, then

$$\mathbb{P}[Z \geq 1] = \frac{\mathbb{E}[Z]}{\mathbb{E}[Z \mid Z \geq 1]}. \quad (1)$$

If we further assume that both walks start from the stationary distribution, then we have

$$\mathbb{P}[Z \geq 1] = \frac{\tau \cdot \|\pi\|_2^2}{\mathbb{E}[Z \mid Z \geq 1]}.$$

To the best of our knowledge, this is the first application of this formula to meeting (and coalescence) times. However, we should mention that variants of this formula have been used by Cooper and Frieze in several works (*e.g.*, [\[CF05\]](#)) to derive accurate bounds on the hitting (and cover time) on various classes of random graphs, and in Barlow et al. [\[BPS12\]](#) to bound the collisions of random walks on infinite graphs. Using [\(1\)](#), we are able to obtain several improvements to existing bounds on the meeting time, and as a consequence for coalescing time. We believe that our work further highlights the power of this basic identity.

The crux of (1) is that in order to lower (or upper) bound the probability that the two walks meet, we need to derive a corresponding bound on $\mathbb{E}[Z \mid Z \geq 1]$, *i.e.*, the number of collisions conditioning on the occurrence of at least one collision. Our results employ various tools to get a handle on this quantity, but here we mention one that is quite intuitive:

$$\mathbb{E}[Z \mid Z \geq 1] \leq \max_{u \in V} \sum_{t=0}^{\tau-1} \sum_{v \in V} (p_{u,v}^t)^2. \quad (2)$$

The inner summand $\sum_{v \in V} (p_{u,v}^t)^2$ is the probability that two walks starting from the same vertex u will meet after a further t steps. Thus, summing over t and conditioning on the first meeting happening (*i.e.*, the condition $Z \geq 1$) at some vertex u before time-step τ yields the bound in (2). Despite the seemingly crude nature of this bound, it can be used to derive new results for t_{hit} , t_{meet} and t_{coal} that significantly improve over the state-of-the-art for regular graphs (see [Appendix B](#), or the last paragraph in this section for a summary).

Bounds on t_{coal} in terms of t_{hit}

The derivation of our bounds on t_{coal} in terms of t_{hit} ([Theorem 1.3](#)) are based on two general reduction results, that might be useful in other applications:

Theorem 1.6 (Reduction Results). *The following results hold for any graph G :*

1. *The coalescence process reduces the number of walks from n to $O(\log^3 n)$ in $O(t_{\text{hit}})$ steps with probability at least $1 - n^{-1}$. (see [Theorem 4.3](#))*
2. *The coalescence process reduces the number walks from $\log^4 n$ to $(\Delta/d)^{O(1)}$ in $O(t_{\text{hit}})$ steps in expectation, where Δ is the maximum degree and d is the average degree (see [Theorem 4.4](#))*

A basic ingredient are new concentration inequalities, which are derived in [Section 4.1](#). Our concentration inequalities yield sufficiently strong bounds for upper tails of returns (or other, possibly more complex random variables) by a random walk of length t_{hit} , while most of the existing bounds (*e.g.*, [[CLLM12](#); [Lez89](#)]) require that the expectation of the random variable is at least as large as t_{mix} . While $t_{\text{mix}} \leq t_{\text{hit}}$ in general, the challenging case in our analysis is when $t_{\text{mix}} \approx t_{\text{hit}}$ and in this cases our concentration inequalities provide stronger upper tails than the existing ones.

Equipped with these concentration results, the proof of [Theorem 4.3](#) is surprisingly simple and rests again on (1). First, by a straightforward bucketing argument on the degree distribution, we show that with high probability, we can find for each random walk $(X_t)_{t \geq 0}$ with label i a set S (depending on the trajectory of X_t), so that with high probability, (i) each vertex in S is visited frequently during $O(t_{\text{hit}})$ steps, and (ii) each vertex in S has the same degree up to constant factors. Conditioning on this, it follows that a second random walk $(Y_t)_{t \geq 0}$ will have sufficient collisions with $(X_t)_{t \geq 0}$ in expectation, *i.e.*, $\mathbb{E}[Z]$ is large enough. To bound $\mathbb{E}[Z \mid Z \geq 1]$, we use the concentration inequalities to establish that with high probability, the trajectory $(X_t)_{t \geq 0}$ will be good in the sense that $\mathbb{E}[Z \mid (x_0, x_1, \dots), Z \geq 1]$ is not too large. Combining these bounds yields $\mathbb{P}[Z \geq 1] = \Omega(1/\log^3 n)$, and a straightforward division into groups \mathcal{G}_1 and \mathcal{G}_2 of sizes $\Theta(\log^3 n)$ and $n - |\mathcal{G}_1|$ shows that all random walks in \mathcal{G}_2 can be eliminated in $O(t_{\text{hit}})$ steps.

The proof of the second reduction result ([Theorem 4.4](#)) is more involved, although it again revolves around (1). The issue is that we can no longer repeat the simple bucketing argument from [Theorem 4.3](#) about the degree distribution, since the number of buckets may vastly exceed the number of walks. Furthermore, we may no longer obtain “w.h.p.”-bounds on the probability for certain good events. For all these reasons, a refined approach is needed.

Our analysis allocates small phases of length $O(t_{\text{hit}}/\kappa)$ in order to halve the number of random walks, where $k = \kappa^c$ is the number of walks at the beginning of the phase, for some suitably large constant c . The first step is to show that starting from any vertex, there exists a large set of vertices, so that each vertex is visited the “right” amount of time, but also that it was not too unexpected to visit that vertex. The latter condition is quite subtle, but it allows us to arrange a proper scheduling of the walks to show that, regardless of which vertices the random walk i decides to visit in that set, there are enough walks that are able to reach these vertices by then. In other words, it rules out the possibility that, despite two random walks visiting the same set of vertices, they never collide (for an illustration, see [Figure 5](#) on page 45). Using our concentration bounds with a careful choice of the slackness parameters in terms of κ , the above approach can eventually be shown to reduce the number of random walks k by a constant fraction within $O(t_{\text{hit}}/\kappa)$ steps. Repeating this iteratively yields the bound $O(t_{\text{hit}})$.

Bounds on t_{hit} and Worst-Case Bounds

With the two reduction results, [Theorem 1.3](#) follows immediately. Furthermore, the aforementioned results can be also used to derive worst-case upper and lower bounds on meeting and coalescing time on general and regular graphs that are tight up to constant factors. Some of these were known, or follow directly from existing results, the others are novel to the best of our knowledge.

We proceed by establishing that $t_{\text{coal}} = O(n^3)$ on all graphs. The proof of $t_{\text{coal}} = O(n^3)$ ([Theorem 1.5](#)) follows by first applying both reductions ([Theorem 4.3](#) and [Theorem 4.4](#)) to reduce the number of walks from n to $(\Delta/d)^{O(1)} \leq (n^2/|E|)^{O(1)}$ in $O(t_{\text{hit}})$. We have, by [Proposition B.9](#), $t_{\text{meet}} \leq 4t_{\text{hit}} = O(n \cdot |E|)$, where this last bound follows from [\[AKL+79\]](#).

Finally, combining the bound $t_{\text{meet}} = O(n \cdot |E|)$ together with $t_{\text{coal}}(S_0) = O(t_{\text{meet}} \cdot \log(|S_0|))$ ([Proposition 3.4](#)) for any set of start vertices S_0 , yields that after additional

$$O(t_{\text{meet}} \cdot \log(|S_0|)) = O(n \cdot |E| \cdot \log(n^2/|E|)) = O(n^3)$$

steps the coalescing terminates. The fact that this is tight can be easily verified by considering the Barbell graph.⁹

For regular graphs, the same argument as before shows that $t_{\text{coal}} = O(n^2)$, and this is matched by the cycle, for instance. The proofs of the other results are straightforward, and we refer the reader to [Section 5.3](#).

Bounds on t_{meet} and Other Results

In [Appendix B](#), we derive several bounds on t_{meet} . These bounds are derived more directly by (1) and/or (2), and involve other quantities such as $\|\pi\|_2^2$ or the eigenvalue gap $1 - \lambda_2$. One important technical contribution is to combine routine spectral methods involving the spectral representation and fundamental matrices that have been used in previous works, *e.g.*, Cooper et al. [\[CEOR13\]](#) with some short-time bounds on the t -th step probabilities. This allows us to improve several bounds, not only on t_{meet} and t_{coal} but also t_{hit} and t_{cov} , by significantly reducing the dependency on the spectral gap or mixing time—by almost a square root factor. As a corollary, we also derive a new bound on the cover time for regular graphs that considerably improves over the best known bound by Broder and Karlin [\[BK89\]](#) from 1989.

⁹This n -vertex graph is constructed by taking two cliques of size $n/4$ each, and connecting them through a path of length $n/2$.

Concrete Topologies

Finally, in [Appendix C](#), we apply the derived upper and lower bounds on t_{meet} and t_{coal} on various fundamental topologies including grids, expanders and hypercubes. In most cases, these results follow immediately from the general bounds by plugging in corresponding values for $\|\pi\|_2^2$, t_{hit} or t_{mix} . One exception is the binary tree, for which it seems surprisingly non-trivial to derive a lower bound of $t_{\text{meet}} = \Omega(n \log n)$. Here again we use a refinement of (1) that restricts the vertices to leaf-nodes u , for which $\sum_{t=1}^{t_{\text{mix}}} (\sum_{v \in V} p_{u,v}^t)^2 = \Omega(\log n)$. The matching upper bound $t_{\text{meet}} = O(n \log n)$ follows from $t_{\text{coal}} = O(t_{\text{hit}})$ for almost-regular graphs ([Theorem 1.3](#)).

Of particular interest might be the analysis of “real-world” graph models given in [Section C.5](#). There we show how to utilize our bounds from earlier sections to establish $t_{\text{coal}} = \Theta(t_{\text{meet}})$ on two random graph models, leading to bounds on t_{coal} that are sublinear in the number of vertices.

1.3 Discussion and Future Work

In this work we derived several novel bounds on t_{coal} . Our first main result implies that a gap of just $\Omega(\log^2 n)$ between t_{mix} and t_{meet} is sufficient to have $t_{\text{coal}} = \Theta(t_{\text{meet}})$. We also proved that this result is essentially tight. Further, we derived several new bounds on t_{coal} based on t_{hit} . For almost-regular-graphs, our new result implies the following hierarchy for the discrete-time setting,

$$t_{\text{meet}} \leq t_{\text{coal}} = O(t_{\text{hit}}),$$

which refines the already known result $t_{\text{meet}} = O(t_{\text{hit}})$. Finally, we also determined tight worst-case lower and upper bound for t_{coal} .

For future work, an obvious problem is to extend the $t_{\text{coal}} = O(t_{\text{hit}})$ result to all graphs (so far, we only know $t_{\text{coal}} = O(t_{\text{hit}} \cdot \log \log n)$). Even more ambitious would be to try to prove that the continuous-time variant and the discrete-time process are (asymptotically) equivalent, as this would immediately resolve the $t_{\text{coal}} = O(t_{\text{hit}})$ problem. A different direction may be to further explore lower bounds on t_{meet} ; in this work we only derived one lower bound on t_{meet} in [Theorem B.1](#).

2 Notation and Preliminaries

Throughout the paper, let $G = (V, E)$ denote an undirected, connected graph with $|V| = n$ and $|E| = m$. For a node $u \in V$, $\deg(u)$ denotes the degree of u and $N(u) = \{v : (u, v) \in E\}$ the *neighborhood* of u . By Δ , δ and $d = \frac{1}{n} \sum_{u \in V} \deg(u)$, we denote the maximum, minimum and average degree, respectively. We say G is Γ -approximative regular if $\Delta/\delta = \Gamma$.

Unless stated otherwise, all random walks are assumed to be discrete-time (indexed by natural numbers) and lazy, *i.e.*, if P denotes the $n \times n$ transition matrix of the random walk, $p_{u,u} = 1/2$, $p_{u,v} = 1/(2 \deg(u))$ for any edge $(u, v) \in E$ and $p_{u,v} = 0$ otherwise. We define $p_{u,v}^t$ to be the probability that a random walk starting at $u \in V$ is at node $v \in V$ at time $t \in \mathbb{N}$. Furthermore, let $p_{u,\cdot}^t$ be the probability distribution of the random walk after t time steps starting at u . By π we denote the *stationary distribution*, which satisfies $\pi(u) = \deg(u)/(2m)$ for all $u \in V$.

Let $d(t) = \max_u \|p_{u,\cdot}^t - \pi\|_{\text{TV}}$ and $\bar{d}(t) = \max_{u,v} \|p_{u,\cdot}^t - p_{v,\cdot}^t\|_{\text{TV}}$, where $\|\cdot\|_{\text{TV}}$ denotes the total variation distance. Following Aldous and Fill [[AF02](#)], we define the *mixing time* to be $t_{\text{mix}}(\varepsilon) = \min\{t \geq 0 : \bar{d}(t) \leq \varepsilon\}$ and for convenience we will write $t_{\text{mix}} = t_{\text{mix}}(1/e)$. We define separation from stationarity to be $s(t) = \min\{\varepsilon : p_{u,v}^t \geq (1 - \varepsilon)\pi(v) \text{ for all } u, v \in V\}$. Then $s(\cdot)$ is submultiplicative, so in particular, non-increasing [[AF02](#)], and we can define the *separation threshold time* $t_{\text{sep}} = \min\{t \geq 0 : s(t) \leq e^{-1}\}$ and, by [[AF02](#), Lemma 4.11], $t_{\text{sep}} \leq 4t_{\text{mix}}$. We write $T_{\text{hit}}(u, v)$ to denote the first time-step $t \geq 0$ at which a random walk starting at u hits v . In particular, $T_{\text{hit}}(u, u) = 0$.

The *hitting time* $t_{\text{hit}}(u, v) = \mathbb{E}[T_{\text{hit}}(u, v)]$ of any pair of nodes $u, v \in V$ is the expected time required for a random walk starting at u to hit v . Thus, $t_{\text{hit}}(u, v)$ is the expectation of $T_{\text{hit}}(u, v)$. The hitting time of a graph $t_{\text{hit}} = \max_{u, v} t_{\text{hit}}(u, v)$ is the maximum over all such pairs.

For $A \subseteq V$, we use $t_{\text{hit}}(u, A)$, to denote the expected time required for a random walk starting to u to hit some node in the set A . Furthermore, we define $t_{\text{hit}}(\pi, u) = \sum_{v \in V} t_{\text{hit}}(v, u) \cdot \pi(v)$. Furthermore, we define $t_{\text{avg-hit}} = \sum_{u, v \in V} \pi(u) \cdot \pi(v) \cdot t_{\text{hit}}(u, v)$.

Let $t_{\text{meet}}(u, v)$ denote the expected time when two random walks starting at u and v first arrive at the same node at the same time, and we write t_{meet}^π for the expected meeting time of two random walks starting at two independent samples from the stationary distribution. Finally, let $t_{\text{meet}} = \max_{u, v} t_{\text{meet}}(u, v)$ denoted the worst-case expected meeting time.

We define the coalescence process as a stochastic process as follows: Let $S_0 \subseteq V$ be the set of nodes for which there is initially one random walk on it, and for all $v \in S_t$ let

$$Y_v(t) = \begin{cases} u \in N(v) & \text{w.p. } \frac{1}{2|N(v)|} \\ v & \text{w.p. } \frac{1}{2} \end{cases}$$

The set of *active* nodes in step $t + 1$ is given by $S_{t+1} = \{Y_v(t) \mid v \in S_t\}$. The process satisfies the Markov property, *i.e.*,

$$\mathbb{P}[S_{t+1} \mid \mathcal{F}_t] = \mathbb{P}[S_{t+1} \mid S_t], \quad (3)$$

where \mathcal{F}_t is the filtration up to time t , which, informally speaking, is the history of all random decisions up to time t . Finally, we define the *time of coalescence* as $T_{\text{coal}}(S_0) = \min\{t \geq 0 \mid |S_t| = 1\}$. Throughout this paper, the expression w.h.p. (*with high probability*) means with probability at least $1 - n^{-\Omega(1)}$ and the expression w.c.p. (*with constant probability*) means with probability $c > 0$ for some constant c . We use $\log n$ for the natural logarithm. [Appendix A](#) contains some known results about Markov Chains that we frequently use in our proofs.

3 Bounding t_{coal} for large $t_{\text{meet}}/t_{\text{mix}}$

In this section we prove [Theorem 1.1](#), one of our main results. We refer the reader to [Section 1.2](#) for a high-level description of the proof ideas.

3.1 Stochastic Process

In order to prove our first main result, it is helpful to consider a more general stochastic process, P_{imm} , called the immortal process, involving multiple independent random walks. In the immortal process, whenever several random walks arrive at the same node at the same time a subset of them (rather than just one) may survive, while the remaining are merged with one of the surviving walks. To identify the random walks, we assume that each walk has a natural number (in \mathbb{N}) as an identifier. In order to define this process formally, we introduce some additional notation and definitions; then we state and prove some auxiliary lemmas. A related concept was introduced in [\[Oli12, Section 3.4\]](#) under the name of “allowed killings”.

As mentioned before, we assume that every random walk r has a unique identifier $\text{id}(r) \in \mathbb{N}$. We divide the ids into two groups \mathcal{G}_1 , the group of immortal walks and \mathcal{G}_2 the group of the remaining (mortal) walks. Whenever two or more walks collide at a node and at least one of these walks is in \mathcal{G}_1 , then all walks with ids in \mathcal{G}_1 survive, while all walks with ids in \mathcal{G}_2 are killed (merged with some walk with id in \mathcal{G}_1). Furthermore, if all walks have ids in \mathcal{G}_2 , *i.e.*, there are no walks with id in \mathcal{G}_1 , then the walk with the minimum id among these walks survives. The ids along with the




Figure 2: Illustration of the process P_{imm} .

assignment of ids to groups determine which of the random walks that arrive at a given node at the same time survive.

Formally, let P_{imm} denote the following process:

1. At time 0, $S_0 = \{(u_r, \text{id}(r))\}$, where u_r is the starting node of random walk r and $\text{id}(r)$ is its identifier.
2. At time t , several random walks may arrive at the same node. The process P_{imm} allows some subset of them to survive, while the rest ‘coalesce’ with one of the surviving walks. Formally, S_{t+1} is defined using S_t as follows. Define the (random) next-step position of the random walk with $\text{id } i \in \mathbb{N}$ which is on node $v \in V$ to be

$$Y_{v,i}(t) := \begin{cases} u & \text{where } u \in N(v) & \text{w.p. } \frac{1}{2|N(v)|} \\ v & & \text{w.p. } \frac{1}{2}, \end{cases}$$

Let $R_v(t) := \{(Y_{v,i}(t), i) \mid (v, i) \in S_t\}$, $v \in V$ be the set of next-step positions (before merging happens) for random walks that were at node v at time t . Let

$$\hat{R}_v(t) := \{(v, i) \mid \exists u \in V, (v, i) \in R_u(t)\}$$

be the random walks that have arrived at node v at time-step $t + 1$, just before merging happens. Then, merging happens w.r.t. the ids as follows:

- (a) If there exists $i \in \mathcal{G}_1$ such that $(v, i) \in \hat{R}_v(t)$ (at least one walk with id in \mathcal{G}_1 arrives at v), then

$$S_v(t+1) := \{(v, j) \mid (v, j) \in \hat{R}_v(t), j \in \mathcal{G}_1\}$$

- (b) If there is no $i \in \mathcal{G}_1$, such that $(v, i) \in \hat{R}_v(t)$ and $\hat{R}_v(t) \neq \emptyset$ (no walk with id in \mathcal{G}_1 arrives at v , but at least one walk arrives at v), then

$$S_v(t+1) := \{(v, j)\},$$

where $j = \min\{i \mid (v, i) \in \hat{R}_v(t)\}$.

- (c) Otherwise, $S_v(t+1) := \emptyset$, *i.e.*, no walk arrived at v .

Finally, let

$$S_{t+1} := \bigcup_{v \in V} S_v(t+1).$$

We now relate this more general process, P_{imm} , to the coalescing process defined in [Section 2](#). Let P be regarded as a special instance of P_{imm} with $\mathcal{G}_1 = \{1\}$. In process P , only one of several walks arriving at the same node survives and by convention the one having the smallest id is chosen. Let $(S_t)_{t=0}^\infty$ denote the stochastic process P . If we define $\bar{S}_t := \{v \mid (v, i) \in S_t\}$, then $(\bar{S}_t)_{t=0}^\infty$ is a coalescence process as defined in [Section 2](#). Moreover, P represented by $(S_t)_{t=0}^\infty$ is the coalescence process which additionally keeps track of the ids. Throughout this paper, we assume that every random walk of S_0 is on a distinct node.

In the following we show that the time it takes to reduce to k random walks in the original process P is majorized by the time it takes in P_{imm} to reduce to k random walks. While this might be intuitive, one needs to be very careful about the dependencies between the meetings of different random walks: For instance a random walk which is immortal in P_{imm} might eliminate many other random walks whereas the corresponding coupled random walk in P might be eliminated early and therefore cannot eliminate said random walks.

Proposition 3.1. *Consider the following two processes:*

1. *Process P is the standard process of coalescing random walks, viewed as a special case of P_{imm} with $\mathcal{G}_1 = \{1\}$ as described above.*
2. *Process P_{imm} is the process defined above using groups \mathcal{G}_1 and \mathcal{G}_2 , where $1 \in \mathcal{G}_1$.*

*Let T^k, T_{imm}^k be the stopping times given by the condition that fewer than k random walks remain for the two processes respectively. Assume both processes start with the same initial configuration, *i.e.*, the vertices occupied by walks in both processes are identical and there is only one walk per vertex in either process. Then, there exists a coupling such that*

$$T^k \leq T_{\text{imm}}^k.$$

Proof. We will give a coupling between the moves of walks in P_{imm} and P_{int} , a new process that is essentially intermediate between P and P_{imm} ; furthermore, we will show that the original process P is essentially a restricted view of the process P_{int} . The process P_{int} will label the walks *dead*, *alive*, and *phantom*. We emphasize that a phantom walk is not considered alive. Note that the processes P and P_{imm} can be viewed as processes which assign labels to each random walk of the type alive and dead.

Let S_t^Q denote the set of tuples of alive walks in process $Q \in \{P, P_{\text{int}}, P_{\text{imm}}\}$ at time t . Let $\bar{S}_t^Q = \{v \mid (v, i) \in S_t^Q\}$ for $Q \in \{P, P_{\text{int}}, P_{\text{imm}}\}$ be the set of nodes which are occupied by at least one alive walk (there might be several in P_{imm} at $t \geq 1$). In order to prove the proposition, we show that there exists a coupling, such that for any $t \in \mathbb{N}$

$$\bar{S}_t^P \subseteq \bar{S}_t^{P_{\text{int}}} \tag{4}$$

$$\bar{S}_t^{P_{\text{int}}} \subseteq \bar{S}_t^{P_{\text{imm}}} \quad (5)$$

implying that $|\bar{S}_t^P| \leq |\bar{S}_t^{P_{\text{imm}}}|$ which yields the claim since

$$T^k = \min\{t \geq 0: |\bar{S}_t^P| \leq k\} \leq \min\{t \geq 0: |\bar{S}_t^{P_{\text{imm}}}| \leq k\} = T_{\text{imm}}^k.$$

We now define P_{int} . As mentioned above, the walks in P_{int} will be given three kinds of labels alive, dead, or phantom; the dead walks do not continue ahead in time; alive and phantom walks do.

Formally, P_{int} using the groups \mathcal{G}_1 and \mathcal{G}_2 is defined as follows. We say that walk r is of type \mathcal{G}_i , if $\text{id}(r) \in \mathcal{G}_i$ for $i \in \{1, 2\}$. Whenever at least one walk arrives¹⁰ on a node, then the following happens.

1. At least one of the walks is of type \mathcal{G}_1
 - (a) At least one walk of type \mathcal{G}_1 is alive
 - i. the walk of \mathcal{G}_1 with the smallest id is labeled as alive (even if it was labeled phantom before)
 - ii. all other walks of type \mathcal{G}_1 (if there are any) are labeled as phantom
 - iii. alive walks of type \mathcal{G}_2 are labeled dead (if present).
 - (b) All walks of type \mathcal{G}_1 are phantom walks
 - i. There is no walk of type \mathcal{G}_2
 - A. No label is changed
 - ii. There is at least one walk of type \mathcal{G}_2
 - A. the walk of type \mathcal{G}_1 with the smallest id is labeled as alive
 - B. all other walks of type \mathcal{G}_1 (if there are any) are labeled as phantom
 - C. alive walks of type \mathcal{G}_2 are labeled dead.
2. All walks are of type \mathcal{G}_2
 - (a) the walk of \mathcal{G}_2 with the smallest id is labeled as alive
 - (b) all other walks are labeled as dead.

Note, that walks of \mathcal{G}_1 are either alive or phantom and walks of \mathcal{G}_2 are either alive or dead. Also, note that in the process P_{int} , there is at most one *alive* walk at any given node. Throughout the proof we regard the processes in two stages: First, each random walk selects a destination (possibly the same node it was on) and moves there. In the second phase the walks are merged according to the process. See [Figure 3](#) for an illustration.

We prove (4) by induction on t starting from the same initial configuration: if $v \in \bar{S}_t^P$, then $v \in \bar{S}_t^{P_{\text{int}}}$. Consider the inductive step from t to $t+1$ and assume that the claim holds at the end of round t (after merging happened). For the (unique) random walk at $v \in \bar{S}_t^P$ under process P , we couple its transition to node $Y_v(t+1)$ (where we possibly have $Y_v(t+1) = v$) with the corresponding alive walk of $\bar{S}_t^{P_{\text{int}}}$ (there might be several walks of P_{int} , however only one is alive and we couple with this alive walk). Let S be the set of nodes to which a random walk in P moved, *i.e.*, $S = \{Y_v(t+1): v \in \bar{S}_t^P\}$. Observe, that before the merging takes place in round $t+1$ (but moves have been made), there is, by induction hypothesis and the coupling, at least one alive walk

¹⁰Throughout, by arrive we take into account that walks may arrive at a node from the same node through laziness.

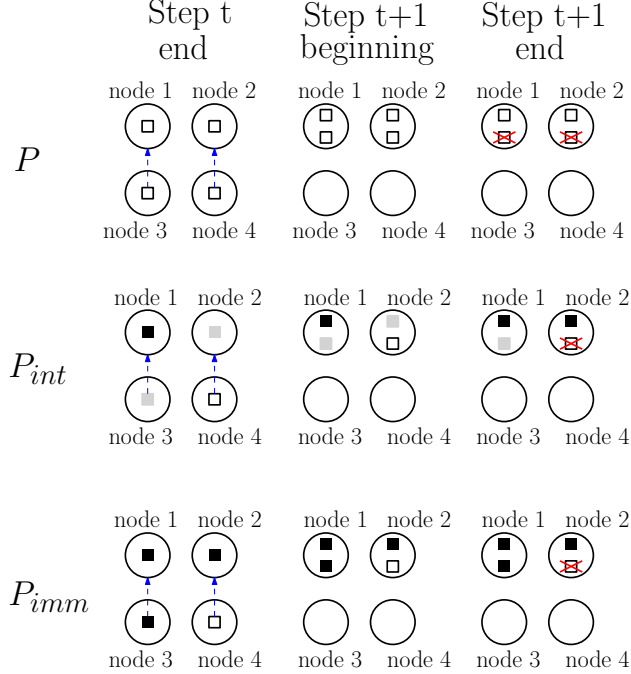


Figure 3: An illustration of couplings between the processes. The squares depict the random walks. Walks of \mathcal{G}_1 are coloured black and grey (phantom) and the nodes of \mathcal{G}_2 are white. The blue arrows denote the moving decisions. Observe that in P_{int} a phantom becomes alive (and a walk of \mathcal{G}_2 is labeled dead).

of P_{int} on each node of S . Furthermore, the definition of P_{int} ensures that whenever an alive random walk moves to a node, then after merging takes place, at least¹¹ one alive walk remains. Thus, our coupling ensures that if $v \in \tilde{S}_{t+1}^P$, then $v \in \tilde{S}_{t+1}^{P_{\text{int}}}$. In words, if one looks at the subsets where there is an alive walk of P_{int} , this is essentially the standard coalescence process. This finishes the proof of (4) and we turn to proving (5).

When starting from the same initial configuration, we will provide a coupling that satisfies the following invariants.

1. There is a bijective map from the alive and phantom walks of P_{int} to the alive walks of P_{imm} , such that the following holds. All walks of P_{int} of type \mathcal{G}_i are mapped to walks of P_{imm} of type \mathcal{G}_i , for $i \in \{1, 2\}$.
2. Whenever a walk of type \mathcal{G}_2 is labeled dead in P_{imm} , then it is also labeled dead in P_{int} and vice versa.

At the beginning there are no dead or phantom walks in P_{int} , there are no dead walks in P_{imm} , all walks are alive and as the starting positions in P_{imm} and P_{int} are the same, an arbitrary bijective mapping may be chosen, so long as it respects node positions and walk types.

Assume the invariant holds at time t . We take one random walk step for each alive or phantom random walk in P_{int} . These are coupled with the corresponding walks in P_{imm} , under the chosen map. Walks that are already dead are neither simulated in P_{int} nor in P_{imm} . Hence, we can ensure the bijection between the walks of \mathcal{G}_1 in both processes holds at time $t + 1$.

¹¹By definition, there is actually exactly one alive walk.

We now prove the second invariant. Note that whenever a walk r of type \mathcal{G}_2 in $P_{\text{imm}}(P_{\text{int}})$ is labeled dead, this implies there must have been another walk r' on the same node at the same time. Since there is a bijective map, r' must be on the same node in $P_{\text{int}}(P_{\text{imm}})$. We have that either r' is of type \mathcal{G}_1 or r' is of type \mathcal{G}_2 and that $\text{id}(r') < \text{id}(r)$. In either case, r is also killed (labeled dead) in $P_{\text{int}}(P_{\text{imm}})$. Hence, we can ensure the bijection between the walks of \mathcal{G}_2 in both processes holds at time $t + 1$. Thus, the invariant holds at time $t + 1$. By induction, and since the alive walks of P_{int} are a subset of the alive walks of P_{imm} the invariant holds throughout the process and yielding (5). This finishes the proof. \square

3.2 Meeting Time Distribution Prior to t_{meet}

Let $(X_t)_{t \geq 0}$ and $(Y_t)_{t \geq 0}$ be independent random walks starting at arbitrary positions. For τ a multiple of t_{mix} , the following lemma gives a lower bound on the probability of intersection of the two random walks in τ steps.

Lemma 3.2. *Let $(X_t)_{t \geq 0}$ and $(Y_t)_{t \geq 0}$ be two independent random walks starting at arbitrary positions. Let $\text{intersect}(X_t, Y_t, \tau)$ be the event that there exists $0 \leq s \leq \tau$, such that $X_s = Y_s$. Then*

$$\mathbb{P}[\text{intersect}(X_t, Y_t, 5t_{\text{mix}})] \geq \frac{1}{32\alpha},$$

where $\alpha = t_{\text{meet}}/t_{\text{mix}}$. Furthermore, there exists a constant $c > 0$, such that for any $1 \leq b \leq \frac{e-1}{e} \cdot \alpha$, we have

$$\mathbb{P}[\text{intersect}(X_t, Y_t, cbt_{\text{mix}})] \geq \frac{b}{\alpha},$$

Proof. First, let $(\tilde{X}_t)_{t \geq 0}$ and $(\tilde{Y}_t)_{t \geq 0}$ be two random walks that start from two independent samples drawn from the stationary distribution and are run for $\ell := 2\lceil\alpha\rceil\lceil t_{\text{mix}}\rceil$ steps. Notice that $\ell \geq 2t_{\text{meet}}$, and hence, by Markov's inequality,

$$\mathbb{P}[\text{intersect}(\tilde{X}_t, \tilde{Y}_t, \ell)] \geq \frac{1}{2}. \tag{6}$$

Furthermore, if we divide the interval $[1, \ell]$ into $2\lceil\alpha\rceil$ consecutive sections of length $\lceil t_{\text{mix}}\rceil$ each, the probability for a collision in each of these section is identical and therefore the union bound implies

$$\mathbb{P}[\text{intersect}(\tilde{X}_t, \tilde{Y}_t, \ell)] \leq 2\lceil\alpha\rceil \cdot \mathbb{P}[\text{intersect}(\tilde{X}_t, \tilde{Y}_t, t_{\text{mix}})], \tag{7}$$

and hence combining equation (6) and (7) yields

$$\mathbb{P}[\text{intersect}(\tilde{X}_t, \tilde{Y}_t, t_{\text{mix}})] \geq \frac{1}{4\lceil\alpha\rceil}.$$

Consider now two independent random walks $(X_t)_{t \geq 0}$ and $(Y_t)_{t \geq 0}$ starting at arbitrary positions. By applying Lemma A.5 to both walks, with probability at least $(1 - e^{-1})^2$ both $X_{4t_{\text{mix}}}$ and $Y_{4t_{\text{mix}}}$ are drawn independently from the stationary distribution since $4t_{\text{mix}} \geq t_{\text{sep}}$. Therefore,

$$\mathbb{P}[\text{intersect}(X_t, Y_t, 5t_{\text{mix}})] \geq (1 - e^{-1})^2 \cdot \mathbb{P}[\text{intersect}(\tilde{X}_t, \tilde{Y}_t, t_{\text{mix}})] \geq (1 - e^{-1})^2 \cdot \frac{1}{4\lceil\alpha\rceil}.$$

Observing that for any $\alpha \geq 1$, the RHS above expression is greater than $1/(32\alpha)$ completes the proof of the first part. For the second part, we consider k blocks of length $5t_{\text{mix}}$. Due to independence

of different blocks, the probability of that the two walks meet in at least one of the k blocks is at least $1 - (1 - \frac{1}{32\alpha})^k$. We set $k := \lceil 32b/(1 - e^{-1}) \rceil$, $x := 1/(32\alpha)$. We distinguish between two cases.

Case $k \cdot x < 1$: We use the fact that $(1 - x)^k \leq e^{-xk} \leq 1 - (1 - e^{-1})xk$ for $0 \leq x < 1$, $k \geq 0$ and $xk \leq 1$. We derive that the probability of intersecting after k blocks is at least $1 - (1 - \frac{1}{32\alpha})^k \geq (1 - e^{-1})k/(32\alpha) = b/\alpha$.

Case $k \cdot x \geq 1$: We have $1 - (1 - \frac{1}{32\alpha})^k \geq 1 - (1 - \frac{1}{32\alpha})^{32\alpha} \geq 1 - 1/e \geq b/\alpha$. In both cases the second part follows. \square

At the heart of the proof of [Theorem 1.1](#) lies the following lemma that analyses the marginal distribution of the meeting time distribution. That is, we only expose the first random walk $(X_t)_{t=0}^\tau$, and look at how this affects the probability of meeting. In essence, we show that at least one of the two ‘‘orthogonal’’ cases hold. In Case 1 (corresponding to set C_1), there is at least a modest probability that after exposing (X_t) , (Y_t) will intersect with significant probability. Otherwise, in Case 2 (corresponding to set C_2), there is a significant probability that after exposing (X_t) , (Y_t) will intersect with at least a modest probability.

Lemma 3.3. *Fix $\tau \in \mathbb{N}$ and a graph G . Let $(X_t)_{t=0}^\tau$ and $(Y_t)_{t=0}^\tau$ be independent random walks, where the starting nodes X_0 and Y_0 are drawn independently from the stationary distribution π (w.r.t. to G), and the walks are run for τ steps. Let $p = \mathbb{P}[\text{intersect}(X_t, Y_t, \tau)]$ and let \mathcal{T}_τ denote the set all possible trajectories of a walk of length τ in G (including possible self-loops). We define the following two categories C_1 and C_2 with $C_1 \subseteq C_2$*

$$\begin{aligned} C_1 &:= \{(z_0, \dots, z_\tau) \in \mathcal{T}_\tau : \mathbb{P}[\exists 0 \leq s \leq \tau, Y_s = z_s] \geq \sqrt{p}\} \\ C_2 &:= \{(z_0, \dots, z_\tau) \in \mathcal{T}_\tau : \mathbb{P}[\exists 0 \leq s \leq \tau, Y_s = z_s] \geq p/3\}. \end{aligned}$$

Then, $\mathbb{P}[(X_t)_{t=0}^\tau \in C_1] \geq \frac{p}{3}$ or $\mathbb{P}[(X_t)_{t=0}^\tau \in C_2] \geq \frac{\sqrt{p}}{3}$.

While the actual lower bounds on the probabilities appear rather crude, it turns out that the ‘‘significant probability’’ $\sqrt{p}/3$ is best possible, as we demonstrate in our lower bound construction later. Remarkably, the fact that the ‘‘modest probability’’ is only $p/3$ and much smaller than $\sqrt{p}/3$ does not affect the tightness of our bound, since in [Claim 3.5](#), we can make up for this gap in both cases through a simple amplification argument over the unexposed random walks.

Proof. Let us suppose that $\mathbb{P}[(X_t)_{t=0}^\tau \in C_1] < \frac{p}{3}$. We show that this implies $\mathbb{P}[(X_t)_{t=0}^\tau \in C_2] \geq \frac{\sqrt{p}}{3}$. Assume for the sake of contradiction $\mathbb{P}[(X_t)_{t=0}^\tau \in C_2] < \frac{\sqrt{p}}{3}$. We have

$$\begin{aligned} p &= \mathbb{P}[\text{intersect}(X_t, Y_t, \tau)] \\ &\leq \mathbb{P}[(X_t)_{t=0}^\tau \in C_1] \cdot 1 + \mathbb{P}[(X_t)_{t=0}^\tau \in (C_2 \setminus C_1)] \cdot \sqrt{p} + \mathbb{P}[(X_t)_{t=0}^\tau \notin C_2] \cdot \frac{p}{3} \\ &< p/3 + \sqrt{p}/3 \cdot \sqrt{p} + p/3 \leq p, \end{aligned}$$

a contradiction. This completes the proof. \square

It is well-known that starting with k random walks, the coalescence time is bounded by $O(t_{\text{meet}} \log k)$, this can be deduced from the proof presented in [\[HP01\]](#). For the sake of completeness, we give a self-contained proof¹².

¹²One might be tempted to pair random walks in groups of two and run them for $2t_{\text{meet}}$ time steps so that, by Markov inequality, they meet with probability at least $1/2$. Repeating this iteratively would yield the claim. To formalize such an argument one would need to disallow coalescence between different pairs of random walk which differs from the stochastic process we reduce to in [Section 3.1](#).

Proposition 3.4. *We have $t_{\text{coal}}(S_0) = O(t_{\text{meet}} \log |S_0|)$.*

Proof. Let P be the coalescing process (with ids) defined in [Section 3.1](#). Recall that $\mathcal{G}_1 = \{1\}$. Let S_t be set of coalescing random walks at an arbitrary time-step t . In the following we show the slightly stronger claim that the expected time to reduce the number of random walks by a constant factor is $O(t_{\text{meet}})$.

Formally, we fix an arbitrary time-step t_0 . With $T := \min\{t \geq t_0 : |S_t| \leq 99/100 \cdot |S_{t_0}|, |S_{t_0}| \geq 100\}$ denoting the first time-step the number of coalescing random walks reduces by a factor of $99/100$, we will prove that $\mathbb{E}[T] = O(t_{\text{meet}})$. Iterating the argument $O(\log |S_0|)$ times implies that the expected time it takes to reduce to 100 random walks is $O(t_{\text{meet}} \log |S_0|)$. Note that the expected time to reduce from 100 random walks to 1 is bounded by $O(t_{\text{meet}})$. Hence, the claim $t_{\text{coal}}(S_0) = O(t_{\text{meet}} \log |S_0|)$ follows.

It remains to show that the expected number of time steps it takes to reduce the number of random walks by a factor of $99/100$ is indeed $O(t_{\text{meet}})$.

We divide time into blocks of length $\tau := c \frac{e-1}{e} t_{\text{meet}} + 4t_{\text{mix}}$, where c is the constant of [Lemma 3.2](#), *i.e.*, $\mathbb{P}[\text{intersect}(X_t, Y_t, c \frac{e-1}{e} t_{\text{meet}})] \geq \frac{e-1}{e}$. We are primarily interested in what happens at the end of the blocks, *i.e.*, at time steps $t_0, t_0 + \tau, t_0 + 2\tau, \dots$. For simplicity, we will start counting time from 0 at the beginning of each block. Let $(X_t)_{t \geq 0}$ be the random walk with id 1. After $4t_{\text{mix}}$ steps, we can couple the state of the random walk $(X_t)_{t \geq 4t_{\text{mix}}}$ with a node drawn from π with probability at least $(1 - e^{-1})$, since $4t_{\text{mix}} \geq t_{\text{sep}}$ (see [Lemma A.5](#)). Further, note that conditioned on this coupling, the statement of [Lemma 3.3](#) implies that $(X_t)_{t \geq 4t_{\text{mix}}} \in C_2$ w.p. at least $p/3$, where we used $C_2 \subseteq C_1$, and where $p := \mathbb{P}[\text{intersect}(\tilde{X}_t, \tilde{Y}_t, c \cdot \frac{e-1}{e} \cdot t_{\text{meet}})] \geq \frac{e-1}{e}$ for $\tilde{X}_0, \tilde{Y}_0 \sim \pi$.

We condition on the successful coupling of $X_{4t_{\text{mix}}}$ with a node drawn from π and that $(X_t)_{t \geq 4t_{\text{mix}}} \in C_2$, which happens with probability at least $(1 - e^{-1})p/3 = \frac{(e-1)^2}{3e^2}$ (called event \mathcal{E}). Finally, consider any random walk $(Y_t)_{t \geq 0}$ with id other than 1. Again with probability at least $1 - e^{-1}$ we can couple $Y_{4t_{\text{mix}}}$ with a node drawn from π and conditioned on successful coupling, $(Y_t)_{t \geq 4t_{\text{mix}}}$ meets $(X_t)_{t \geq 4t_{\text{mix}}}$ between time-steps $[4t_{\text{mix}}, \tau]$ with probability at least $p/3$, by definition of C_2 . Thus, conditioned on event \mathcal{E} , each walk of \mathcal{G}_2 vanishes w.p. $(1 - e^{-1})p/3 = \frac{(e-1)^2}{3e^2}$ and thus the expected fraction of walks killed in the τ time-steps is at least $\frac{(e-1)^2}{3e}$.

Let $Z_\ell = |S_{t_0 + \ell \cdot \tau}|$ denote the number of random walks alive at the beginning of block ℓ .

$$\mathbb{E}[Z_\ell \mid \mathcal{F}_{t_0 + (\ell-1) \cdot \tau}] \leq Z_{\ell-1} - (Z_{\ell-1} - 1) \cdot \frac{(e-1)^4}{9e^4} \leq Z_{\ell-1} - \frac{Z_{\ell-1}}{100}.$$

The above holds as long as $Z_{\ell-1} \geq 100$. We can therefore apply [Lemma A.9](#) with parameters $g = 99/100 \cdot S_0$ and $\beta = 99/100$ to obtain that $\mathbb{E}[T] = O(\tau) = O(t_{\text{meet}})$, which completes the proof. \square

3.3 Upper Bound - Proof of [Theorem 1.1](#)

We commence by considering the process P_{imm} defined in [Section 3.1](#). This allows us to establish [Claim 3.5](#) providing us with the following tradeoff. For a given period τ of length at least t_{mix} we obtain a bound on the required number of periods to reduce the number of random walks by an arbitrary factor. The proof relies heavily on [Lemma 3.3](#) which divides the walks of \mathcal{G}_1 into two groups allowing us to expose the walks of \mathcal{G}_1 first and then to calculate the probability of the walks of \mathcal{G}_2 to intersect with them. In fact, we will also use the symmetric case where the roles of \mathcal{G}_1 and \mathcal{G}_2 are switched. These probabilities are derived from the time-probability tradeoff presented in [Lemma 3.2](#). We then use [Claim 3.5](#) to derive a bound on the number of time steps it takes

to reduce the number of walks to $\lceil 2\alpha \rceil$, where $\alpha = t_{\text{meet}}/t_{\text{mix}}$ (Corollary 3.6). From there on we employ Claim 3.5 to reduce the number of walks to 1 in $O(t_{\text{meet}})$ time steps. Melding both phases together yields the bound of Theorem 1.1.

We now define a process $P_{\text{imm}}(S_0, k)$ with $k < |S_0|$, which is a parameterized version of the process P_{imm} defined in Section 3.1:

- Let $|S_0| = k'$; there are k' random walks with ids $1, \dots, k'$ and starting nodes $v_1, \dots, v_{k'}$. Thus, $S_0 = \{(v_i, i) \mid 1 \leq i \leq k'\}$.
- Let $\mathcal{G}_1 = \{1, \dots, k\}$ and $\mathcal{G}_2 = \{k+1, \dots, k'\}$. Recall that, by definition of P_{imm} , we have that if some random walks with ids in \mathcal{G}_1 and some with ids in \mathcal{G}_2 are present on the same node at the same time, only the ones with ids in \mathcal{G}_1 survive. If all the random walks have ids only in \mathcal{G}_1 , then all of them survive. If all random walks have ids only in \mathcal{G}_2 , then only the one with the smallest id survives.

We define

$$\text{IDs}(S_t) := \{\text{id}(r) \mid (u_r, \text{id}(r)) \in S_t\}, t \in \mathbb{N}.$$

The following lemma gives the expected time it takes to reduce the number of random walks in \mathcal{G}_2 from $k' - k$ to some arbitrary integer $g \geq k$: given a period of length τ and integer g , assuming that $k = |\mathcal{G}_1|$ is large enough, we derive a bound on the number of periods of length τ until the walks in \mathcal{G}_2 are reduced to g . The required size of k is a function of the probability for two random walks drawn from π intersecting after τ time steps.

Claim 3.5. *Let $\tau \in \mathbb{N}$, let $(X_t)_{t=0}^\tau$ and $(Y_t)_{t=0}^\tau$ be independent random walks run for τ steps, with X_0 and Y_0 drawn independently from π . Let $p_\tau \leq \mathbb{P}[\text{intersect}(X_t, Y_t, \tau)]$ be a lower bound on the probability of the intersection of the two walks during the τ steps. Consider an instantiation of $P_{\text{imm}}(S_0, k)$. Suppose that $k \geq \frac{3}{(1-e^{-1}) \cdot p_\tau}$. For some $1 \leq g \leq |S_0| - k$, define the stopping condition $T_g = \min\{t \geq 0 \mid |\text{IDs}(S_t) \cap \mathcal{G}_2| \leq g\}$. Then the expected stopping time satisfies*

$$\mathbb{E}[T_g] = O\left((4t_{\text{mix}} + \tau) \cdot \sqrt{\frac{1}{p_\tau}} \cdot (\log |\mathcal{G}_2| - \log g)\right).$$

We first describe the high-level proof idea, before delving into the formal proof. We divide time into blocks of size $4t_{\text{mix}} + \tau$. For any random walk $(Z_t)_{t=0}^{4t_{\text{mix}}+\tau}$ we can couple its position after $4t_{\text{mix}} \geq t_{\text{sep}}$ w.c.p. with a node drawn from π . Thus, conditioning on the success of this coupling we have, by Lemma 3.3, $\mathbb{P}\left[(Z_t)_{t=4t_{\text{mix}}}^{4t_{\text{mix}}+\tau} \in C_1\right] \geq \frac{p_\tau}{3}$ or $\mathbb{P}\left[(Z_t)_{t=4t_{\text{mix}}}^{4t_{\text{mix}}+\tau} \in C_2\right] \geq \frac{\sqrt{p_\tau}}{3}$. In the former case we have that w.c.p. there is at least one random walk r in \mathcal{G}_1 which is, due to independence of the walks, in class C_1 . The hypothetical extension of the trajectory of any random walk in $r' \in \mathcal{G}_2$ intersects with r w.p. $c\sqrt{p_\tau}/3$, where the constant arises due to the fact that we also need to couple the state of r' at time $4t_{\text{mix}}$ to a node drawn according to π . (We need to consider the hypothetical extension because the walk r' may get eliminated sooner—this only helps us.) Thus, r' gets eliminated w.p. at least $c\sqrt{p_\tau}$ for a suitable constant c .

In the latter case we have that w.p. at least $c\sqrt{p_\tau}/3$ a random walk of \mathcal{G}_2 is in class C_2 . Every random walk in that class intersects w.c.p. with at least one of the walks of \mathcal{G}_1 . Thus, in both cases, we have that in each block a random walk of \mathcal{G}_2 is eliminated w.p. at least $c\sqrt{p_\tau}$ for some constant c . Thus, the number of random walks in \mathcal{G}_2 decrease in expectation by a factor of $c\sqrt{p_\tau}$.

Proof. We will consider the process in *blocks* each consisting of $4t_{\text{mix}} + \tau$ time-steps. For convenience in the proof, we'll restart counting time-steps from 0 at the beginning of each block; we keep track

of the total number of time-steps by counting the number of blocks. Let C_1 and C_2 be as defined in [Lemma 3.3](#). Then we perform a case analysis by considering the two possible outcomes described in [Lemma 3.3](#) separately. We define $Z_j = |\text{IDs}(S_{j,(4t_{\text{mix}}+\tau)} \cap \mathcal{G}_2)|$, *i.e.*, the number of walks remaining in \mathcal{G}_2 after j blocks of time have passed. For any $j \geq 1$, we will show that there exists a constant $c > 0$ such that,

$$\mathbb{E}[Z_j \mid \mathcal{F}_{j-1}] \leq Z_{j-1} \cdot (1 - c\sqrt{p_\tau}).$$

By using [Lemma A.9](#), we get $\mathbb{E}[T_g] = O\left((4t_{\text{mix}} + \tau) \cdot \frac{1}{\sqrt{p_\tau}} \cdot (\log |\mathcal{G}_2| - \log g)\right)$ (the factor $(4t_{\text{mix}} + \tau)$ appears as the size of the block). Recall that \mathcal{F}_j is the filtration up to end of the j th block. In the remainder we show that we have indeed $\mathbb{E}[Z_j \mid \mathcal{F}_{j-1}] \leq Z_{j-1} \cdot (1 - c\sqrt{p_\tau})$.

Case 1. $\mathbb{P}[(X_t)_{t=0}^\tau \in C_1] \geq \frac{p_\tau}{3}$:

Consider any random walk r in \mathcal{G}_1 at the beginning of a *block*. Using [Lemma A.5](#), after $4t_{\text{mix}}$ steps we can couple the state of the random walk with a node drawn from π with probability at least $(1 - e^{-1})$. Furthermore, conditioned on this coupling, the portion of the random walk between time-steps $4t_{\text{mix}}$ and $4t_{\text{mix}} + \tau$ of the walk is in class C_1 with probability at least $\frac{p_\tau}{3}$. Since $k \geq \frac{3}{p_\tau \cdot (1 - e^{-1})}$, w.p. $c_1 > 0$, in any block, there exists a walk in \mathcal{G}_1 that has the portion between time-steps $4t_{\text{mix}}$ and $4t_{\text{mix}} + \tau$ in C_1 .

Fix a block and condition on the event that there is a walk in \mathcal{G}_1 , denoted by r_1 , whose portion between time-steps $4t_{\text{mix}}$ and $4t_{\text{mix}} + \tau$ is in C_1 . Consider any walk in \mathcal{G}_2 , denoted by r_2 , at the beginning of the block. We want to argue that this walk r_2 has a reasonable probability of intersecting some walk in \mathcal{G}_1 in this block of time-steps. First, consider (the possibly hypothetical continuation of r_2) walk r'_2 for the entire length of the block. The reason for this is that if r_2 and some walk from \mathcal{G}_1 are at the same node at the same time sometime in the block, r_2 will be eliminated in the process $P_{\text{imm}}(S_0, k)$; however, we can consider its hypothetical extension to the entire length of the block. Using [Lemma A.5](#) the state of the walk r'_2 at time-step $4t_{\text{mix}}$ can be coupled with a node drawn from π with probability at least $c_2 := 1 - e^{-1}$. Then conditioned on successful coupling, the probability that r'_2 and r_1 collide during time-steps $4t_{\text{mix}}$ and $4t_{\text{mix}} + \tau$ is at least $\sqrt{p_\tau}$ (by definition of C_1 in [Lemma 3.3](#)). Thus, the probability that r_2 hits at least one walk in \mathcal{G}_1 is at least $c_1 \cdot c_2 \cdot \sqrt{p_\tau}$. Note that it is also possible for r'_2 to be eliminated by another walk from \mathcal{G}_2 . In any case, we have that r_2 is eliminated w.p. at least $c\sqrt{p_\tau}$ and we get

$$\mathbb{E}[Z_j \mid \mathcal{F}_{j-1}] \leq Z_{j-1} \cdot (1 - c_1 \cdot c_2 \sqrt{p_\tau}).$$

Case 2. $\mathbb{P}[(X_t)_{t=0}^\tau \in C_2] \geq \frac{\sqrt{p_\tau}}{3}$:

Consider a walk in \mathcal{G}_2 , denoted by r_2 , at the beginning of a block; as in the previous case, we will consider a possibly hypothetical continuation r'_2 of r_2 . Using [Lemma A.5](#) we can couple the state of r'_2 at time-step $4t_{\text{mix}}$ with a node drawn from π with probability at least $1 - e^{-1}$. Furthermore, conditioned on the successful coupling, with probability at least $\frac{\sqrt{p_\tau}}{3}$ the trajectory of r'_2 between the time-steps $4t_{\text{mix}}$ to $4t_{\text{mix}} + \tau$ is in C_2 . Thus, with probability at least $p := (1 - e^{-1}) \frac{\sqrt{p_\tau}}{3}$, r'_2 has a trajectory between time-steps $4t_{\text{mix}}$ and $4t_{\text{mix}} + \tau$ that lies in C_2 . Now consider any random walk $r_1 \in \mathcal{G}_1$ at the beginning of the block. Again, using [Lemma A.5](#) with probability at least $1 - e^{-1}$, we can couple the state of the random walk at time $4t_{\text{mix}}$ with a node drawn from π . Conditioned on this between time-steps $4t_{\text{mix}}$ to $4t_{\text{mix}} + \tau$, this random walk hits any trajectory whose portion between time-steps $4t_{\text{mix}}$ to $4t_{\text{mix}} + \tau$ lies in C_2 with probability at least $p_\tau/3$ (by definition of C_2 in [Lemma 3.3](#)). Since $k = |\mathcal{G}_1| \geq \frac{3}{(1 - e^{-1}) \cdot p_\tau}$, with at least constant probability $c_1 > 0$ there is some walk in \mathcal{G}_1 that intersects any fixed trajectory whose portion between time-steps $4t_{\text{mix}}$ to $4t_{\text{mix}} + \tau$ lies in C_2 . Since the random walks in \mathcal{G}_1 are independent, by the definition

of the immortal process, we have that any walk in \mathcal{G}_2 is eliminated by the end of the block with probability at least $c_1 \cdot p = c\sqrt{p\tau}$ for some constant $c > 0$. Similarly as before, it is possible that r_2 is eliminated by at least one of the walks of \mathcal{G}_2 , which only increases the probability for r_2 of being eliminated. We get

$$\mathbb{E}[Z_j \mid \mathcal{F}_{j-1}] \leq Z_{j-1} \cdot (1 - c\sqrt{p\tau}).$$

□

In the following we bound the time T required to reduce to $2\lceil\alpha\rceil$ random walks. The claim follows by applying [Claim 3.5](#) to derive a bound on T_{imm} for process P_{imm} , and using the majorization of T by T_{imm} ([Proposition 3.1](#)).

Corollary 3.6. *Consider the coalescence process starting with set S_0 and let $\alpha = t_{\text{meet}}/t_{\text{mix}}$. Let $T_1 = \min\{t \geq 0 \mid |S_t| \leq 2\lceil\alpha\rceil\}$. Then $\mathbb{E}[T_1] = O(t_{\text{mix}} \cdot \sqrt{\alpha} \cdot \log |S_0|)$.*

Proof. We consider the process P (defined in [Section 3.1](#)), which is identical to the coalescence process, but in addition also keeps track of ids of random walks and that allows only the walk with the smallest id to survive. We assume that the ids are from the set $\{1, 2, \dots, |S_0|\}$. Let $S_0 = \{(v_1, 1), \dots, (v_{|S_0|}, |S_0|)\}$ and $\bar{S}_0 = \{i : (v, i) \in S_0\}$. We consider the process $P_{\text{imm}}(S_0, k)$ and $k = \lceil\alpha\rceil$. Let T_1^* be the stopping time defined by $|\text{IDs}(\bar{S}_t) \cap \mathcal{G}_2| \leq \alpha$ for the process $P_{\text{imm}}(S_0, k)$. By definition of P_{imm} and [Proposition 3.1](#), it follows that T_{imm} stochastically dominates T . Thus, it suffices to bound $\mathbb{E}[T_{\text{imm}}]$. W.l.o.g. we assume that $\alpha \geq 6\frac{e-1}{e}$, otherwise the claim follows directly from [Proposition 3.4](#). We apply [Lemma 3.2](#) with $b = 6$ and derive that for some suitable constant c ,

$$p = \mathbb{P}[\text{intersect}(X_{t \geq 0}, Y_{t \geq 0}, 6ct_{\text{mix}})] \geq \frac{6}{\alpha},$$

Thus, we have

$$\frac{3}{(1 - e^{-1}) \cdot p} \leq \frac{3}{\frac{1}{2} \cdot p} \leq \alpha \leq k$$

Applying [Claim 3.5](#) with $g = \alpha$, $\tau = 6ct_{\text{mix}}$ (where c is a constant as given by [Lemma 3.2](#)), $p_\tau = 6/\alpha$, and observing that $k \geq \frac{3}{(1 - e^{-1}) \cdot p_\tau}$, we get the required result. □

In the following we bound the time T required to reduce from $2\lceil\alpha\rceil$ random walks to a single random walk. The proof uses the same ideas as before ([Corollary 3.6](#)) however, this time we consider several phases and in each we reduce the number of random walks by a constant factor. The expected time per phase is geometrically increasing as the number of walks decreases and the overall time is essentially dominated by the time for a constant number of random walks to meet, which is $O(t_{\text{meet}})$.

Lemma 3.7. *Consider the coalescence process starting with set S_0 , satisfying $|S_0| \leq 4\alpha \log \alpha$, where $\alpha = t_{\text{meet}}/t_{\text{mix}}$. Let $T_2 := \min\{t \geq 0 \mid |S_t| \leq 1\}$. Then $\mathbb{E}[T_2] = O(t_{\text{meet}})$.*

Proof. We will consider the coalescence process in phases. Let ℓ be the largest integer such that $|S_0| \geq (\frac{4}{3})^\ell$. For $j \geq 1$, the j^{th} phase ends when $|S_t| < (\frac{4}{3})^{\ell-j+1}$. The $(j+1)^{\text{th}}$ phase begins as soon as the j^{th} phase ends. Note that it may be the case that some phases are empty. Let $T_2(j)$ denote the time for phase j to last. We will only consider phases up to which $\ell - j + 1 \geq 32$.

Now we focus on a particular phase j . Let t_j be the time when the j^{th} phase begins and let S_{t_j} denote the corresponding set at that time. Thus, we have

$$\left(\frac{4}{3}\right)^{\ell-j+1} \leq |S_{t_j}| < \left(\frac{4}{3}\right)^{\ell-j+2} \tag{8}$$

We consider the process P_{imm} defined in [Section 3.3](#) as follows. Define $n_j = |S_{t_j}|$. Fix a phase j and define $S'_0 = \{(v_1, 1), \dots, (v_{n_j}, n_j)\}$ and $\bar{S}'_0 = \{v_1, \dots, v_{n_j}\}$. Then, consider again the set of occupied vertices (ignoring the labels) $\bar{S}_{t_j+t} = \{v \mid \exists i \in \mathbb{N}, (v, i) \in S'_t\}$ with $t \in \mathbb{N}$. Thus, phase j ends when $|S'_t| = |\bar{S}_{t_j+t}| < \left(\frac{4}{3}\right)^{\ell-j+1}$. Let

$$k_j := \left\lfloor \frac{|S'_0|}{2} \right\rfloor$$

be the size of \mathcal{G}_1 and consider the process $P_{\text{imm}}(S'_0, k_j)$ as defined in [Section 3.3](#). Let

$$g_j := \left\lfloor \frac{|S'_0| - k_j}{3} \right\rfloor$$

and

$$T_2^*(j) := \min\{t \mid |\text{IDs}(S'_t) \cap \mathcal{G}_2| \leq g_j\}.$$

We note that as long as $\ell - j + 1 \geq 32$, $g_j \geq 1$ and at time $T_2^*(j)$,

$$|S'_t| \leq g_j + k_j \leq \frac{|S'_0| - k_j}{3} + k_j = \frac{|S'_0|}{3} + \frac{2k_j}{3} \leq \frac{|S'_0|}{3} + \frac{|S'_0|}{3} + \frac{2}{3} < \frac{4}{3} \cdot |S'_0|.$$

By [Proposition 3.1](#), $T_2^*(j)$ stochastically dominates $T_2(j)$ and hence it suffices to bound $\mathbb{E}[T_2^*(j)]$. In order to bound $\mathbb{E}[T_2^*(j)]$, we define

$$b_j := 32\alpha \log(4/3)(\ell - j + 1)(3/4)^{\ell-j+1}.$$

Since we only consider phases with j respecting $\ell - j + 1 \geq 32$ we have $b_j \leq b_{\ell-31} \leq ((e-1)/e)\alpha$. Furthermore, we have $b_j \geq b_0 \geq 4\alpha \log \alpha (3/4)^\ell \geq 1$, where the last inequality follows from $(4/3)^\ell \leq |S_0| \leq 4\alpha \log \alpha$, which in turn follows from definition of ℓ and the assumed bound on $|S_0|$. Applying [Lemma 3.2](#) with this value of b_j , we get that for

$$\tau_j := cb_j t_{\text{mix}},$$

for independent random walks $(X_t)_{t=0}^{\tau_j}, (Y_t)_{t=0}^{\tau_j}$, $\mathbb{P}[\text{intersect}(X_t, Y_t, \tau_j)] \geq p_j$, where

$$p_j := 32 \log(4/3)(\ell - j + 1)(3/4)^{\ell-j+1}.$$

We seek to apply [Claim 3.5](#) to bound $\mathbb{E}[T_2^*(j)]$. We first verify that the conditions of [Claim 3.5](#) are fulfilled. In particular, we verify that $k_j \geq \frac{8}{p_j}$; to see this consider the following:

$$\frac{8}{p_j} = \frac{8}{32 \log(4/3)(\ell - j + 1)} (4/3)^{\ell-j+1} \leq \frac{1}{4} \cdot \left(\frac{4}{3}\right)^{\ell-j+1} \leq \frac{1}{2} \cdot |S'_0| \leq k_j,$$

where we used [\(8\)](#) and $|S'_0| = |S_{t_j}|$ in the second-last inequality. Thus we can apply [Claim 3.5](#) and derive

$$\mathbb{E}[T_2^*(j)] \leq (\tau_j + 4t_{\text{mix}}) \cdot \frac{1}{\sqrt{p_j}} \cdot (\log |\text{IDs}(S'_0) \cap \mathcal{G}_2| - \log g_j)$$

and we continue by dissecting that bound. Since $b_j \geq 1$, there exists a suitably large constant c_1 , so that $\tau_j + 4t_{\text{mix}} \leq c_1 b_j t_{\text{mix}}$. Furthermore,

$$\frac{b_j}{\sqrt{p_j}} = \frac{32\alpha \log(4/3)(\ell - j + 1)(3/4)^{\ell-j+1}}{\sqrt{32 \log(4/3)(\ell - j + 1)(3/4)^{\ell-j+1}}} = O\left(\alpha \sqrt{\ell - j + 1} \cdot \left(\frac{3}{4}\right)^{(\ell-j+1)/2}\right).$$

Observe that, by definition, $|\text{Ds}(S'_0) \cap \mathcal{G}_2|/g_j \leq 3$, hence $\log |\text{Ds}(S'_0) \cap \mathcal{G}_2| - \log g_j \leq \log(3)$. Putting everything together, we get that there is a constant c_2 such that,

$$\mathbb{E}[T_2^*(j)] \leq c_2 \cdot t_{\text{mix}} \cdot \alpha \cdot \sqrt{\ell - j + 1} \left(\frac{3}{4}\right)^{(\ell-j+1)/2} \quad (9)$$

Note that since we stop when $\ell - j + 1 < 32$, there are at most $\ell - 30$ phases considered. Let \tilde{T} be the random variable denoting the time-step when the last phase ends; at this point $|S_{\tilde{T}}| = O(1)$. Therefore, using [Proposition 3.4](#), $\mathbb{E}[T_2 - \tilde{T} | \tilde{T}] = O(t_{\text{meet}})$. But, clearly \tilde{T} is stochastically dominated by $\sum_{j=0}^{\ell-30} T_2^*(j)$. Thus, we have

$$\begin{aligned} \mathbb{E}[T_2] &= \mathbb{E}[\tilde{T}] + \mathbb{E}\left[\mathbb{E}[T_2 - \tilde{T} | \tilde{T}]\right] \\ &\leq c_2 \cdot t_{\text{mix}} \cdot \alpha \sum_{j=0}^{\ell-30} \sqrt{\ell - j + 1} \left(\frac{3}{4}\right)^{(\ell-j+1)/2} + c_3 t_{\text{meet}} \end{aligned} \quad (10)$$

$$\leq c_2 \cdot t_{\text{mix}} \cdot \alpha + c_3 t_{\text{meet}} = O(t_{\text{meet}}) \quad (11)$$

Above, in [\(10\)](#) we used [\(9\)](#) and the fact that $\mathbb{E}[T_2 - \tilde{T} | \tilde{T}] \leq c_3 t_{\text{meet}}$ for some constant $c_3 > 0$ and in step [\(11\)](#), we used the fact that $\sum_{j=32}^{\infty} j c^j < 1$ for $c \leq \sqrt{3/4}$. \square

Thus, the first phase ([Corollary 3.6](#)) and the second phase ([Lemma 3.7](#)) take together $O(\sqrt{\alpha} \cdot \log n \cdot t_{\text{mix}} + t_{\text{meet}})$ time-steps, which yields [Theorem 1.1](#).

3.4 Lower Bound - Proof of [Theorem 1.2](#)

In this section we give a construction of a graph family in order to establish lower bounds on $t_{\text{coal}}(G)$ in terms of $t_{\text{meet}}(G)$ and $t_{\text{mix}}(G)$ demonstrating that [Theorem 1.1](#) is asymptotically tight. Additionally, our construction generalizes a claim of Aldous and Fill [[AF02](#), Chapter 14]: They mention that it is possible to construct regular graphs that mimic the n -star in the sense that the $t_{\text{meet}} = o(t_{\text{avg-hit}})$, without giving further details of the construction. Our construction shows that even the coalescence time can be significantly smaller than the average hitting time for almost-regular graphs. For our family of almost-regular graphs, there is a polynomial gap between t_{meet} and $t_{\text{avg-hit}}$. More importantly, we show that these almost-regular graphs have a gap of $\sqrt{t_{\text{mix}}/t_{\text{meet}}} \cdot \log n$ between coalescing and meeting time. This shows that the bound in [Theorem 1.1](#) is best possible, even if we constrain G to be almost-regular. We refer the reader to [Section 1.2](#) for a high-level description of the proof ideas.

More precisely, in the proof of [Theorem 1.2](#) we shall give an explicit construction of a graph family $G = G_n$ with $t_{\text{coal}} = \Omega(\sqrt{\alpha_n} \cdot \log n \cdot t_{\text{mix}})$, where $\alpha_n = t_{\text{meet}}/t_{\text{mix}}$. For the remainder of this section, we will drop the dependence on n and will simply use G instead of G_n and α instead α_n .

The construction of G (see [Figure 4](#) for an illustration) is based on two building blocks, G_1 and G_2 . First, let $G_1 = (V_1, E_1)$ be a clique of size \sqrt{n} . Let $G_2 = (V_2, E_2)$ be a \sqrt{n} -regular bipartite Ramanujan Graph on $n/\sqrt{\alpha'}$ nodes [[MSS15](#)], where $\alpha' = \max\{\alpha, 2^{20} \cdot C^2\}$, where $C > 1$ is the universal constant of [Corollary A.3](#). The graph G is made of one copy of G_2 , $\kappa = \sqrt{n}$ copies of G_1 (denoted by $G_1^1, G_1^2, \dots, G_1^\kappa$), and a node \hat{z} , which has an edge to $\sqrt{n/\alpha'}$ distinct nodes of G_2 and to each of the designated nodes $z^i \in V_1^i$ in G_1^i for $i \in [1, \kappa]$. It is not difficult to see that this graph is almost-regular, *i.e.*, maximum and minimum degree differ by at most a constant factor.

In [Lemma 3.12](#), [Lemma 3.13](#), [Lemma 3.14](#) and [Lemma 3.15](#) respectively we show that $t_{\text{mix}} = \Theta(n)$, $t_{\text{meet}} = \Theta(\alpha' n)$, $t_{\text{coal}} = \Omega(\sqrt{\alpha'} \cdot n \log n)$, and $t_{\text{avg-hit}} = \Omega(n^{3/2})$. We start with the following

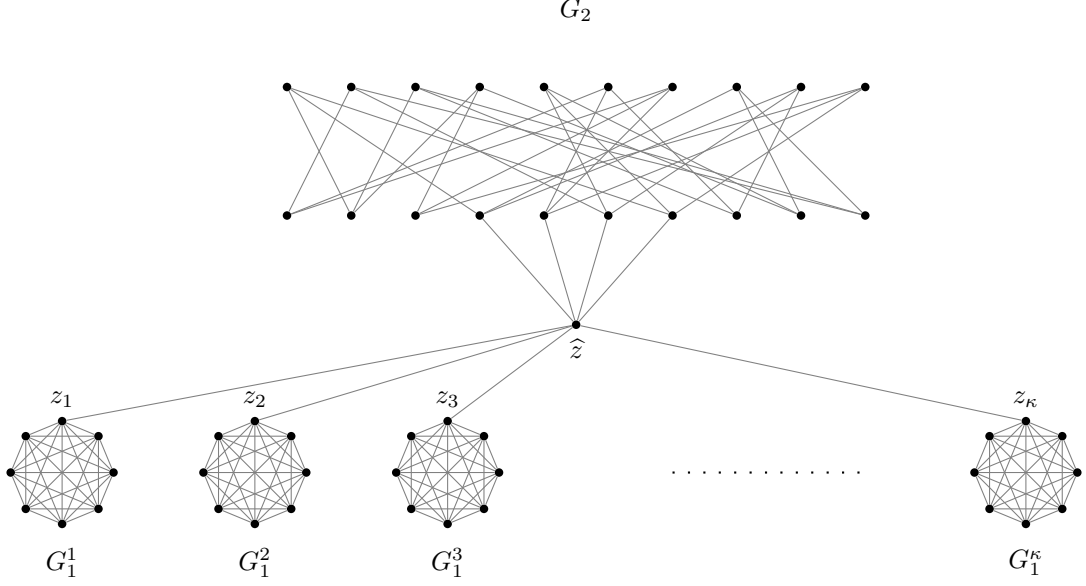


Figure 4: The graph described in [Section 3.4](#) with $t_{\text{coal}} = \Omega(t_{\text{meet}} + \sqrt{t_{\text{meet}}/t_{\text{mix}}} \cdot \log n \cdot t_{\text{mix}})$.

auxiliary lemma which shows that the walk restricted to V_2 behaves similarly to the walk restricted to $V_2 \cup \{\hat{z}\}$, meaning that the walks have very similar t -step probabilities.

Lemma 3.8. *Let P denote the transition matrix of the random walk on G , Q the transition matrix of the random walk on G_2 and \hat{Q} be the transition matrix of the random walk on the subgraph of G induced by $V_2 \cup \{\hat{z}\}$. Let $p_{u,v}^t, q_{u,v}^t, \hat{q}_{u,v}^t$ denote the corresponding transition probabilities for a walk starting at u to end up at node v after t steps. Let $S^* = \{u \in V_2 \cap N(\hat{z})\}$. Then the following statements hold:*

(i) For any $u, v \in V_2$ we have $\|p_{u,\cdot}^t - q_{u,\cdot}^t\|_{\text{TV}} \leq \sum_{i=1}^{t-1} p_{u,S^*}^i / (2\sqrt{n}) \leq t / (2\sqrt{n})$.

(ii) For any $u, v \in V_2$ we have $\|\hat{q}_{u,\cdot}^t - q_{u,\cdot}^t\|_{\text{TV}} \leq \sum_{i=1}^{t-1} p_{u,S^*}^i / (2\sqrt{n}) \leq t / (2\sqrt{n})$.

(iii) For any $u, v \in V_2$ we have that after $t = t_{\text{mix}}(G_2)$ time steps $\|p_{u,\cdot}^t - p_{v,\cdot}^t\|_{\text{TV}} \leq o(1) + 2/e$.

Proof. Let $(X_t)_{t \geq 0}$ be the Markov chain with transition matrix P and let $(Y_t)_{t \geq 0}$ be the Markov chain with transition matrix Q . We will inductively couple these two random walks starting from $X_0 = Y_0 = u$. Given that we coupled both chains up to time $t-1$, we can couple (X_t, Y_t) such that $X_t = Y_t$ with an error probability

$$\begin{aligned} \mathbb{P}[X_t \neq Y_t \mid X_{t-1} = Y_{t-1}] &= \mathbb{P}[X_t \neq Y_t \mid X_{t-1} = Y_{t-1}, X_{t-1} \in S^*] \cdot \mathbb{P}[X_{t-1} \in S^*] \\ &\quad + \mathbb{P}[X_t \neq Y_t \mid X_{t-1} = Y_{t-1}, X_{t-1} \in V_2 \setminus S^*] \cdot \mathbb{P}[X_{t-1} \in V_2 \setminus S^*] \\ &\leq p_{u,S^*}^{t-1} / (2\sqrt{n}) + 0. \end{aligned}$$

We have, by [[LPW06](#), Proposition 4.7], $\|p_{u,\cdot}^t - p_{v,\cdot}^t\|_{\text{TV}} = \inf\{\mathbb{P}[X \neq Y] \mid (X, Y) \text{ is a coupling of } p_{u,\cdot}^t \text{ and } p_{v,\cdot}^t\}$. Hence, by a union bound over t steps,

$$\begin{aligned} \|p_{u,\cdot}^t - p_{v,\cdot}^t\|_{\text{TV}} &= \inf\{\mathbb{P}[X \neq Y] \mid (X, Y) \text{ is a coupling of } p_{u,\cdot}^t \text{ and } p_{v,\cdot}^t\} \leq \mathbb{P}[X_t \neq Y_t] \\ &\leq \sum_{i=1}^{t-1} p_{u,S^*}^i / (2\sqrt{n}) \leq \frac{t}{2\sqrt{n}}. \end{aligned}$$

To prove the second part we redefine $(X_t)_{t \geq 0}$ to be the Markov chain with transition matrix \widehat{Q} and the proof is identical.

We proceed with the last part. For $u, v \in V_2$ we have that after $t = t_{\text{mix}}(G_2)$ time steps, by the triangle inequality and using that $t_{\text{mix}}(G_2) = O(1)$, by [Lemma A.2](#), we get

$$\begin{aligned} \|p_{u,\cdot}^t - \pi^{G_2}(\cdot)\|_{\text{TV}} &\leq \|p_{u,\cdot}^t - q_{u,\cdot}^t\|_{\text{TV}} + \|q_{u,\cdot}^t - \pi^{G_2}(\cdot)\|_{\text{TV}} \\ &\leq \frac{t_{\text{mix}}(G_2)}{2\sqrt{n}} + \|q_{u,\cdot}^t - \pi^{G_2}(\cdot)\|_{\text{TV}} \\ &\leq o(1) + \|q_{u,\cdot}^t - \pi^{G_2}(\cdot)\|_{\text{TV}} \leq o(1) + 1/e, \end{aligned}$$

where the last inequality follows from the definition of mixing time. Again, by the triangle inequality, $\|p_{u,\cdot}^t - p_{v,\cdot}^t\|_{\text{TV}} \leq o(1) + 2/e$. \square

Based on [Lemma 3.8](#), we can now bound the hitting time to reach \widehat{z} , which will later be used to establish the bounds on the mixing and meeting time of the whole graph G . But first, we prove that the mixing time of the graph \widehat{G} induced by $V_2 \cup \{\widehat{z}\}$ is constant and that after mixing on \widehat{G} , the random walk has a probability of $\Omega(1/n)$ to hit \widehat{z} in a constant number of time steps.

Lemma 3.9. *The following three statements hold.*

- (i) *Let \widehat{G} be the induced graph by the vertices $V_2 \cup \{\widehat{z}\}$. Then $t_{\text{mix}}(\widehat{G}) = O(1)$.*
- (ii) *Let $u \in V \setminus \{\widehat{z}\}$. Then there exists a constant $c \geq 1$ such that $\mathbb{P}[T_{\text{hit}}(u, \widehat{z}) \geq n/c] \geq 1/2$.*
- (iii) *Let $u \in V \setminus \{\widehat{z}\}$. Then $t_{\text{hit}}(u, \widehat{z}) = O(n)$.*

Proof. We prove the statements one by one.

- (i) Let Q be the transition matrix of a random walk restricted to G_2 . Let $d^Q(t)$ be the total variation distance w.r.t. the transition matrix Q . Further, let \widehat{Q} be the transition matrix of a random walk restricted to \widehat{G} . Recall that $t_{\text{mix}}(G_2) = O(1)$, by [Lemma A.2](#).

Fix an arbitrary $t \in [2t_{\text{mix}}(G_2), 2t_{\text{mix}}(G_2) + 7]$. In the following we show $\|\widehat{q}_{u,\cdot}^t - \pi^{\widehat{G}}(\cdot)\|_{\text{TV}} \leq 1/e$. We first consider any start vertex $u \in V_2 \setminus \{\widehat{z}\}$ and afterwards the vertex $u = \widehat{z}$. Let \mathcal{D} be the set of distributions over $V(\widehat{G}) = V_2 \cup \{\widehat{z}\}$ assigning no probability mass to \widehat{z} , i.e.,

$$\mathcal{D} = \{D' : \text{for } u \sim D' \text{ we have } \mathbb{P}[u = \widehat{z}] = 0\}. \quad (12)$$

For any such $D' \in \mathcal{D}$, we have, by definition of the total variation distance,

$$\|\widehat{q}_{u \sim D', \cdot}^t - \pi^{\widehat{G}}(\cdot)\|_{\text{TV}} = 0 + \frac{1}{2} \sum_{v \in V_2} \left| \widehat{q}_{u \sim D', v}^t - \pi^{\widehat{G}}(v) \right| + \frac{1}{2} \left| \widehat{q}_{u \sim D', \widehat{z}}^t - \pi^{\widehat{G}}(\widehat{z}) \right|.$$

For $u \in V_2$ observe that $\pi^{\widehat{G}}(u) \in [\pi^{G_2}(u)(1 - \zeta), \pi^{G_2}(u)(1 + \zeta)]$ for some $\zeta = o(1)$. By [[LPW06](#), Exercise 4.1] we have the following identity for $d^Q(t)$. Let \mathcal{D}^* be the set of all distributions over $V(G_2)$, then

$$d^Q(t) = \max_{D \in \mathcal{D}^*} \|q_{u \sim D, \cdot}^t - \pi^{G_2}(\cdot)\|_{\text{TV}} \geq \max_{D' \in \mathcal{D}} \|q_{u \sim D', \cdot}^t - \pi^{G_2}(\cdot)\|_{\text{TV}}.$$

Thus, for $\delta_v := |\hat{q}_{u,v}^t - q_{u,v}^t|$, we get by using triangle inequality,

$$\begin{aligned}
\frac{1}{2} \sum_{v \in V_2} \left| \hat{q}_{u \sim D', v}^t - \pi^{\hat{G}}(v) \right| &\leq \frac{1}{2} \sum_{v \in V_2} \left| \hat{q}_{u \sim D', v}^t - \pi^{G_2}(v) \right| + \frac{1}{2} \sum_{v \in V_2} \left| \pi^{G_2}(v) - \pi^{\hat{G}}(v) \right| \\
&\leq \frac{1}{2} \sum_{v \in V_2} \left| \hat{q}_{u \sim D', v}^t - \pi^{G_2}(v) \right| + \frac{1}{2} \sum_{v \in V_2} \pi^{G_2}(v) \zeta \\
&\leq \frac{1}{2} \sum_{v \in V_2} \left| q_{u \sim D', v}^t - \pi^{G_2}(v) \right| + \frac{1}{2} \sum_{v \in V_2} |\delta_v| + \frac{1}{2} \sum_{v \in V_2} \pi^{G_2}(v) |\zeta| \\
&\leq d^Q(t) + 1/32 + \frac{\zeta}{2}, \\
&\leq d^Q(t) + 1/32 + 1/32,
\end{aligned} \tag{13}$$

where the second-last inequality is due to [Lemma 3.8](#).(ii), $\frac{1}{2} \sum_{v \in V} |\delta_v| \leq t/(2\sqrt{n}) \leq \frac{1}{32}$. By definition of the $t_{\text{mix}}(G_2)$ and by sub-multiplicativity we have $d^Q(t) \leq d^Q(2t_{\text{mix}}(G_2)) \leq 1/e^2$.

The above equation (13) only consider the variation distance w.r.t. V_2 . For \hat{z} we have $\frac{1}{2} |\hat{q}_{u \sim D', \hat{z}}^t - \pi^{\hat{G}}(\hat{z})| \leq (2t_{\text{mix}}(G_2) + 7)/\sqrt{n} \leq 1/32$.

Putting everything together we get we get

$$\begin{aligned}
\|\hat{q}_{u \sim D', \cdot}^t - \pi^{\hat{G}}(\cdot)\|_{\text{TV}} &= \frac{1}{2} \sum_{v \in V_2} \left| \hat{q}_{u \sim D', v}^t - \pi^{\hat{G}}(v) \right| + \frac{1}{2} \left| \hat{q}_{u \sim D', \hat{z}}^t - \pi^{\hat{G}}(\hat{z}) \right| \\
&\leq d^Q(t) + 1/32 + 1/32 + 1/32 \leq 1/e^2 + 3/32 \\
&\leq 1/e.
\end{aligned} \tag{14}$$

$$\leq 1/e. \tag{15}$$

Consider the random walk starting at \hat{z} and let (X_0, X_1, \dots) denote its trajectory. Observe that at time 7 we have

$$\hat{q}_{\hat{z}, \hat{z}}^7 \leq \frac{1}{2^7} + \sum_{i \leq 7} \sum_{v \in N(\hat{z})} \hat{q}_{\hat{z}, v}^{i-1} \cdot \frac{1}{2(\sqrt{n} + 1)} \leq \frac{1}{2^7} + \frac{7^2}{\sqrt{n}} \leq 1/32.$$

The set of distribution for the position of the random walk at time 7 conditioning on $X_7 \neq \hat{z}$ gives the same distribution \mathcal{D} as defined in (12). Let $D_{\hat{z}} \in \mathcal{D}$ be distribution of the random at time 7 starting at \hat{z} . Hence, by(14), we get

$$\begin{aligned}
\|\hat{q}_{\hat{z}, \cdot}^{2t_{\text{mix}}(G_2)+7} - \pi^{\hat{G}}(\cdot)\|_{\text{TV}} &\leq \hat{q}_{\hat{z}, V(\hat{G}) \setminus \{\hat{z}\}} \cdot \|\hat{q}_{u \sim D_{\hat{z}}, \cdot}^{2t_{\text{mix}}(G_2)} - \pi^{\hat{G}}(\cdot)\|_{\text{TV}} + \hat{q}_{\hat{z}, \hat{z}} \cdot 1 \\
&\leq 1 \cdot (1/e^2 + 3/32) + 1/32 \leq 1/e.
\end{aligned} \tag{16}$$

$$\leq 1 \cdot (1/e^2 + 3/32) + 1/32 \leq 1/e. \tag{17}$$

Thus, for $t' = 2t_{\text{mix}}(G_2) + 7$ we have $\|\hat{q}_{\hat{z}, \cdot}^{t'} - \pi^{\hat{G}}(\cdot)\|_{\text{TV}} \leq 1/e$. Together with (14), we conclude that for all $u \in V'$, $\|\hat{q}_{u, \cdot}^{t'} - \pi^{\hat{G}}(\cdot)\|_{\text{TV}} \leq 1/e$ and by definition of t_{mix} and we get $t_{\text{mix}}(\hat{G}) \leq 2t_{\text{mix}} + 7 = O(1)$.

- (ii) To prove $\mathbb{P}[T_{\text{hit}}(u, \hat{z}) \geq n/c] \geq 1/2$ for $u \in V_2$ we show that the random walk restricted to \hat{G} does not hit \hat{z} after n/c_1 steps w.c.p. for some large enough constant c_1 . By the Union bound, for some large constants c_1, c_2 that

$$\begin{aligned}
\mathbb{P}[T_{\text{hit}}^{\hat{G}}(u, \hat{z}) \leq n/c_1] &= \mathbb{P}[T_{\text{hit}}^{\hat{G}}(u, \hat{z}) \leq n/c_1] \leq \sum_{t=1}^{n/c_1} \hat{q}_{u, \hat{z}}^t \leq \sum_{t=1}^{c_2 \log n} 1/\sqrt{n} + \sum_{t=c_2 \log n}^{n/c_1} \hat{q}_{u, \hat{z}}^t \\
&\leq o(1) + n/c_1 \cdot (\pi^{\hat{G}}(\hat{z}) + 1/n^2) \leq 1/2,
\end{aligned}$$

where we used $\hat{q}_{u,\hat{z}}^t \leq \pi^{\hat{G}}(\hat{z}) + \sqrt{\frac{\pi^{\hat{G}}(\hat{z})}{\pi^{\hat{G}}(u)}} \lambda_2(\hat{G})^t$ (Lemma A.1).

We proceed by bounding that $\mathbb{P}[T_{\text{hit}}(u, \hat{z}) \geq n/c_1] \geq 1/2$ for $u \in V_1$. Consider first a random walk $(\tilde{X}_t)_{t \geq 0}$ restricted to $G_1^1 = G_1$ that starts at vertex z^1 and let \tilde{P} denote the transition matrix. Furthermore, in order to couple the random walk \tilde{X}_t restricted to G_1 with a random walk in G , we will consider the random variable $\tilde{Z} := \sum_{t=0}^{t_{\text{sep}}^{G_1}} \mathbf{1}_{\tilde{X}_t = z^1}$. Since G_1 is a clique, $t_{\text{sep}}^{G_1} = O(1)$, and $\tilde{p}_{z^1, z^1}^t \leq \frac{1}{\sqrt{n}} + \lambda_2(G_1)^t$ by Lemma A.1, where $\lambda_2(G_1)$ is some constant bounded away from 1. Therefore, $\mathbb{E}[\tilde{Z}] = \sum_{t=0}^{n/c_1} \tilde{p}_{z^1, z^1}^t \leq 2\sqrt{n}/c_1$. Let $\gamma := 4 \cdot \mathbb{E}[\tilde{Z}]$. Then, by Markov's inequality

$$\mathbb{P}[\tilde{Z} \geq \gamma] \leq 1/4.$$

Consider now the straightforward coupling between a random walk $(X_t)_{t \geq 1}$ in G that starts at vertex z^1 and the random walk $(\tilde{X}_t)_{t \geq 1}$ restricted to G_1^i that starts at the same vertex. Whenever the random walk \tilde{X}_t is at a vertex different from z^1 , then the random walk X_t makes the same transition. If the random walk \tilde{X}_t is at vertex z^1 , then there is a coupling so that the random walk X_t makes the same transition as \tilde{X}_t with probability $\frac{2\sqrt{n}-1}{2\sqrt{n}}$. Conditional on the event $\tilde{Z} \leq \gamma$ occurring, the random walk \tilde{X}_t follows the random walk X_t up until step n/c_1 with probability at least

$$p := \left(\frac{2\sqrt{n}-1}{2\sqrt{n}} \right)^\gamma \geq 3/4,$$

since the random walk \tilde{X}_t has at most γ visits to z^1 . Therefore, by the Union bound,

$$\mathbb{P}[T_{\text{hit}}^G(u, \hat{z}) \geq n/c_1] \geq \mathbb{P}\left[\bigcup_{t=0}^{n/c_1} X_t = \tilde{X}_t\right] \geq 1 - \mathbb{P}[\tilde{Z} \geq \gamma] - (1-p) \geq 1/2$$

and the proof is complete.

(iii) We proceed by showing $t_{\text{hit}}(u, \hat{z}) = O(n)$ for $u \in V_2$.

Let Q be the transition matrix of the random walk restricted to G_2 . Let $u \in V_2$ and $S^* = N(\hat{z})$ be the neighbors of \hat{z} in G_2 . For every $v \in S^*$ we have $\pi^{G_2}(v) = \frac{\sqrt{n+1}}{\frac{n}{\sqrt{\alpha'}}\sqrt{n} + \frac{\sqrt{n}}{\sqrt{\alpha'}}} \geq \frac{\sqrt{\alpha'}}{1.2n}$. Hence, after $t = t_{\text{sep}}(G_2)$ we have that

$$q_{u, S^*}^t := \sum_{v \in S^*} q_{u,v}^t \geq \sum_{v \in S^*} \pi^{G_2}(v)(1 - e^{-1}) \geq \frac{\sqrt{n}}{\sqrt{\alpha'}} \cdot \frac{\sqrt{\alpha'}}{1.2n}(1 - e^{-1}) = \frac{1 - e^{-1}}{1.2\sqrt{n}}.$$

By Lemma 3.8, we have for any $u \in V_2$ that $\|p_{u,\cdot}^t - q_{u,\cdot}^t\|_{\text{TV}} \leq t_{\text{sep}}(G_2)/(2\sqrt{n})$. To bound $T_{\text{hit}}^G(u, \hat{z})$ we show that after $t_{\text{sep}} + 1 = O(1)$ steps the random walk hits \hat{z} w.p. $\Omega(1/n)$.

We distinguish between two cases.

(a) For all $i \leq t$ we have $p_{u,S^*}^t \leq 1/t_{\text{sep}}(G_2)$. Thus, by [Lemma 3.8.\(i\)](#)

$$\begin{aligned} p_{u,S^*}^t &= \sum_{v \in S^*} p_{u,v}^t \geq q_{u,S^*}^t - \|p_{u,\cdot}^t - q_{u,\cdot}^t\|_{\text{TV}} \\ &\geq \frac{1 - e^{-1}}{1.2\sqrt{n}} - \sum_{i=1}^{t-1} p_{u,S^*}^i / (2\sqrt{n}) \\ &\geq \frac{1 - e^{-1}}{1.2\sqrt{n}} - \frac{t_{\text{sep}}(G_2)}{t_{\text{sep}}(G_2)2\sqrt{n}} = \Omega(1/\sqrt{n}). \end{aligned}$$

Hence, the random walk hits \hat{z} after $t_{\text{sep}}(G_2) + 1$ w.p. at least $p_{u,S^*}^t \cdot \min_{v \in S^*} \{p_{v,\hat{z}}\} = \Omega(1/n)$.

(b) Otherwise there exists a t^* such that $p_{u,S^*}^{t^*} > 1/t_{\text{sep}}(G_2)$. Thus the random walk hits \hat{z} after $t_{\text{sep}}(G_2) + 1$ w.p. at least $p_{u,S^*}^{t^*} \cdot \min_{v \in S^*} \{p_{v,\hat{z}}\} = \Omega(1/n)$.

Thus after $O(1)$ steps the random walk hits \hat{z} w.p. $\Omega(1/n)$.

We now show a similar statement if $u \in V_1$. Let $(X_t)_{t \geq 0}$ be a random walk on G starting on u . Observe that X_t (the walk on G) hits \hat{z} with probability $p_{u,z^1}^1 \cdot p_{z^1,\hat{z}}^1 = \Omega(1/n)$ in 2 time steps. Hence, for any $u \in V$ we $\mathbb{P}[T_{\text{hit}}(u, \hat{z}) = O(1)] = \Omega(1/n)$. Thus, repeating this iteratively and using independence yields $t_{\text{hit}}(u, \hat{z}) = O(n)$ for $u \in V$. □

To establish a bound on the mixing time of G , we will make use of the following result of Peres and Sousi.

Theorem 3.10 ([\[PS15\]](#)). *For any $\beta < 1/2$, let $t_{\text{hit}}(\beta) = \max_{u,A:\pi(A) \geq \beta} t_{\text{hit}}(u, A)$. Then there exist positive constants c_β and c'_β such that*

$$c'_\beta \cdot t_{\text{hit}}(\beta) \leq t_{\text{mix}}(1/4) \leq c_\beta \cdot t_{\text{hit}}(\beta).$$

In the following we show for any β close enough to $1/2$, that any $A \subseteq V$ satisfying $\pi(A) \geq \beta$ must include at least a constant fraction of nodes from a constant fraction of copies of G_1 .

Claim 3.11. *Let $\beta = 1/2 - 10^{-3}$. For any $A \subseteq V$ with $\pi(A) \geq \beta$, define $H(A) = \{i \mid |G_1^i \cap A| \geq |V_1|/(2e)\}$. Then, $|H(A)| \geq \kappa/(2e)$.*

Proof. This follows from a simple pigeon-hole argument: Suppose $|H(A)| < \kappa/(2e)$ was true. Then,

$$\begin{aligned} \pi(A) &\leq |H(A)| \cdot \pi(V_1) + (\kappa - |H(A)|) \cdot \left(\frac{\pi(V_1)}{2e} + \pi(z^i) \right) + \pi(V_2) + \pi(\hat{z}) \\ &< \frac{\kappa}{2e} \cdot \pi(V_1) + \kappa \cdot \left(\frac{\pi(V_1)}{2e} + \pi(z^i) \right) + 1/20 < \beta \leq \pi(A), \end{aligned}$$

which is a contradiction and hence choice of A must fulfill $|H(A)| \geq \kappa/(2e)$. □

We are now ready to determine the mixing time of G . The lower bound is a simple application of Cheeger's inequality, while the upper bound combines the previous lemmas with [Theorem 3.10](#).

Lemma 3.12. *Let G be the graph described at the beginning of [Section 3.4](#). We have $t_{\text{mix}}(G) = \Theta(n)$.*

Proof. First we show $t_{\text{mix}} = \Omega(n)$. The *conductance* of $G = (V, E)$ is defined by $\Phi(G) = \min_{\substack{U \subseteq V, \\ 0 < \text{vol}(U) \leq \text{vol}(V)/2}} \frac{|E(U, V \setminus U)|}{\text{vol}(U)}$. In particular, for $U = V_1$ we get that $\Phi(G) \leq \frac{4}{n}$. Hence, by Cheeger's inequality and $\left(\frac{1}{1 - \lambda_2(G)} - 1\right) \cdot \log\left(\frac{e}{2}\right) \leq t_{\text{mix}}(1/e)$ (see, e.g., [LPW06, Chapter 12]),

$$\frac{n}{4} \leq \frac{1}{\Phi(G)} \leq \frac{2}{1 - \lambda_2(G)} = \frac{2}{1 - \lambda_2(G)} - 2 + 2 \leq \frac{2t_{\text{mix}}}{\log\left(\frac{e}{2}\right)} + 2.$$

Rearranging the terms yields $t_{\text{mix}} = \Omega(n)$.

We proceed with the upper bound on the mixing time. Let $\beta = 1/2 - 10^{-3}$ and let $A \subseteq V$ be an arbitrary set satisfying $\pi(A) \geq \beta$. First, we apply [Claim 3.11](#) to conclude that $|H(A)| \geq \kappa/(2e)$. This immediately implies that with $Z := \{z^i : i \in H(A)\}$, $|Z| \geq \kappa/(2e)$. The remainder of the proof is divided into the following three parts:

- (i) Starting from any vertex $u \in V$, with probability at least $1/2$, the random walk hits z^* after $2 \max_{u \in V} t_{\text{hit}}(u, \hat{z}) = O(n)$ steps.
- (ii) With constant probability $p_1 > 0$, the random walk moves from z^* to a vertex in Z .
- (iii) With constant probability $p_2 > 0$ a random walk starting from a vertex in Z will hit A after one step.

It is clear that combining these three results shows that with constant probability $\frac{1}{2}p_1p_2 > 0$, a random walk starting from an arbitrary vertex $u \in V$ hits a vertex in A after $O(n) + 1 + 1$ time-steps. Iterating this and using independence shows that $t_{\text{hit}}(u, A) = O(n)$, and hence by [Theorem 3.10](#), $t_{\text{mix}} = O(n)$ as needed.

Part (i). Consider $\max_{u \in V} t_{\text{hit}}(u, \hat{z})$. For $u \in V$, [Lemma 3.9.\(iii\)](#) implies $t_{\text{hit}}(u, \hat{z}) = O(n)$.

Part (ii). If the random walk is on z^* , then since $\deg(z^*) = \kappa + \sqrt{n/\alpha'}$, $|Z| \geq \kappa/(2e)$, it follows that the random walk hits a vertex in Z after one step with constant probability $p_1 := \frac{|Z|}{2(\kappa + \sqrt{n/\alpha'})} > 0$.

Part (iii). Finally, for any $z \in Z$ we have that $p_2 = p_{z,A} = \frac{|V_1|/(2e)}{2\sqrt{n}} > 0$ and the proof is complete. \square

In the following we establish the bound on the meeting time. As it turns out, any meeting is very likely to happen on V_2 and it takes about $\Theta(\alpha'n)$ time-steps until both walks reach V_2 simultaneously. The lower bound then follows from our common analysis method [\(1\)](#). The upper bound combines the mixing time bound of $O(n)$ ([Lemma 3.12](#)), and that once a random walk reaches a copy of G_1 , it stays there for $\Theta(n)$ steps with constant probability [Lemma 3.9.\(ii\)](#).

Lemma 3.13. *Let G be the graph described at the beginning of [Section 3.4](#). We have $t_{\text{meet}}(G) = \Theta(\alpha'n)$.*

Proof. We start by proving $t_{\text{meet}} = \Omega(\alpha'n)$: Consider two non-interacting, random walks with starting positions drawn from the stationary distribution π . Let $\ell = c'\alpha'n$, for some small enough constant $c' > 0$. Let Z_1 be the number of collisions of the two random walks on the nodes in $V_1^1 \cup V_1^2 \cup \dots \cup V_1^\kappa$. Let Z_2 be the number of collisions of the two random walks on the nodes in V_2 . Let Z_* be the number of collisions of the two random walks on the node \hat{z} .

Let Z be the number of collisions of the two walks during the first ℓ time steps, i.e., $Z = Z_1 + Z_2 + Z_*$. Using the Union bound we derive

$$\begin{aligned}
\mathbb{P}[Z \geq 1] &\leq \mathbb{P}[Z_1 \geq 1] + \mathbb{P}[Z_2 \geq 1] + \mathbb{P}[Z_* \geq 1] \\
&\leq \frac{\mathbb{E}[Z_1]}{\mathbb{E}[Z_1 | Z_1 \geq 1]} + \frac{\mathbb{E}[Z_2]}{\mathbb{E}[Z_2 | Z_2 \geq 1]} + \frac{\mathbb{E}[Z_*]}{\mathbb{E}[Z_* | Z_* \geq 1]}. \tag{18}
\end{aligned}$$

We have $\mathbb{E}[Z_1] \leq \ell n \left(\frac{2}{n}\right)^2$, $\mathbb{E}[Z_2] \leq \ell \frac{n}{\sqrt{\alpha'}} \left(\frac{2}{n}\right)^2$, and $\mathbb{E}[Z_*] \leq \ell \left(\frac{2}{n}\right)^2$, since $\max_u \pi(u) \leq 2/n$. Conditioning on $Z_1 \geq 1$ and since both random walks start from the stationary distribution, we have, by [Observation A.8](#), that the first meeting happens in the first $\ell/2$ time steps w.p. at least $1/2$.

Consider $\mathbb{E}[Z_1 | Z_1 \geq 1]$. Suppose the meeting occurred at node $u \in V_1$. Let \mathcal{E}_1 be the event that for $u \in V_1$ we have $T_{\text{hit}}(u, \hat{z}) \geq n/c$ for both walks, where $c > 0$ is a large enough constant. By [Lemma 3.9\(ii\)](#), we have that $\mathbb{P}[\mathcal{E}_1] \geq (1/2)^2 = 1/4$ due to independence of the walks. For any $t < n/c$ let $\hat{p}_{u,\cdot}^t$ be the distribution of the random walk on G_1 starting on u after t time steps under the conditioning \mathcal{E}_1 . Observe that $\sum_{v \in V_1} \hat{p}_{u,v}^t = 1$ implying that $\sum_{v \in V_1} (\hat{p}_{u,v}^t)^2 \geq \sum_{v \in V_1} \left(\frac{1}{|V_1|}\right)^2 = 1/|V_1|$. Hence, we get

$$\mathbb{E}[Z_1 | Z_1 \geq 1] \geq \mathbb{E}[Z_1 | Z_1 \geq 1, \mathcal{E}_1] \cdot \mathbb{P}[\mathcal{E}_1] \geq \frac{1}{2} \min_{u \in V_1} \sum_{t=0}^{n/c-1} \sum_{v \in V_1} (\hat{p}_{u,v}^t)^2 \geq \frac{1}{4} \sum_{t=0}^{n/c-1} 1/|V_1| = \frac{\sqrt{n}}{4c}.$$

Using an exactly analogous analysis for Z_2 we can upper bound $\mathbb{E}[Z_2 | Z_2 \geq 1]$ as follows:

$$\mathbb{E}[Z_2 | Z_2 \geq 1] \geq \mathbb{E}[Z_2 | Z_2 \geq 1, \mathcal{E}_2] \cdot \mathbb{P}[\mathcal{E}_2] \geq \frac{1}{4} \min_{u \in V_2} \sum_{t=0}^{n/c-1} \sum_{v \in V_2} (\hat{p}_{u,v}^t)^2 \geq \frac{1}{4} \sum_{t=0}^{n/c-1} 1/|V_2| = \frac{\sqrt{\alpha'}}{4c},$$

where \mathcal{E}_2 is the event that for $u \in V_2$ we have $T_{\text{hit}}(u, \hat{z}) \geq n/c$ for some large enough constant c . Plugging everything into [\(18\)](#) and using $\ell = c'\alpha'n$ yields

$$\begin{aligned}
\mathbb{P}[Z \geq 1] &\leq \frac{\mathbb{E}[Z_1]}{\mathbb{E}[Z_1 | Z_1 \geq 1]} + \frac{\mathbb{E}[Z_2]}{\mathbb{E}[Z_2 | Z_2 \geq 1]} + \frac{\mathbb{E}[Z_*]}{\mathbb{E}[Z_* | Z_* \geq 1]} \\
&\leq \frac{\ell n \left(\frac{2}{n}\right)^2}{\frac{\sqrt{n}}{4c}} + \frac{\ell \frac{n}{\sqrt{\alpha'}} \left(\frac{2}{n}\right)^2}{\frac{\sqrt{\alpha'}}{4c}} + \frac{\ell \left(\frac{2}{n}\right)^2}{1} \\
&\leq o(1) + 16c \cdot c' + o(1) \leq 1/2,
\end{aligned}$$

for any constant $c' \in (0, \frac{1}{33c}]$. This finishes the proof of $t_{\text{meet}} = \Omega(\alpha'n)$. In the remainder we prove $t_{\text{meet}} = O(\alpha'n)$. Consider two independent walks $(X_t)_{t \geq 0}$ and $(Y_t)_{t \geq 0}$ on G , both starting from arbitrary nodes. Note $t_{\text{sep}} = t_{\text{sep}}(G) \leq 4t_{\text{mix}} = O(n)$ by [Lemma 3.12](#), and

$$p_0 := \mathbb{P}[\{X_{t_{\text{sep}}} \in V_2\} \cap \{Y_{t_{\text{sep}}} \in V_2\}] \geq \left(\sum_{u \in V_2} (1-e)\pi(u) \right)^2 = \Omega\left((1/\sqrt{\alpha'})^2\right) = \Omega(1/\alpha').$$

We assume in the following that $\{X_{t_{\text{sep}}} \in V_2\} \cap \{Y_{t_{\text{sep}}} \in V_2\}$. We have $t_{\text{mix}}(G_2) = O(1)$, by [Lemma A.2](#). Consider a random walk $(\tilde{X}_t)_{t \geq t_{\text{sep}}}$ restricted to G_2 that starts at vertex $X_{t_{\text{sep}}} \in V_2$ and let \tilde{P} denote the transition matrix. Furthermore, in order to couple the random walk \tilde{X}_t restricted to G_2 with a random walk in G , we will consider the random variable

$$\tilde{Z} := \sum_{t=t_{\text{sep}}}^{t_{\text{sep}}+n/c-1} \sum_{z \in N(\hat{z})} \mathbf{1}_{\tilde{X}_t=z},$$

for $c = 32$. Thus, for any $z \in N(\hat{z})$,

$$\begin{aligned} \mathbb{E} \left[\tilde{Z} \right] &\leq t_{\text{mix}}(G_2) + \sum_{t=t_{\text{sep}}+t_{\text{mix}}(G_2)+1}^{t_{\text{sep}}+n/c-1} |N(\hat{z})|(\pi^{G_2}(z) + d^{\tilde{P}}(t)) \\ &\leq t_{\text{mix}}(G_2) + |N(\hat{z})|(n/c) + O(1) \leq (1 + 1/e)\sqrt{n}/c. \end{aligned}$$

Let $\gamma := 8(1 + 1/e)\sqrt{n}/c$. Then, by Markov's inequality

$$\mathbb{P} \left[\tilde{Z} \geq \gamma \right] \leq 1/8.$$

Consider now the straightforward coupling between a random walk $(X_t)_{t \geq t_{\text{sep}}}$ in G that starts at vertex $\tilde{X}_{t_{\text{sep}}} \in V_2$ and the random walk $(\tilde{X}_t)_{t \geq t_{\text{sep}}}$ restricted to G_2 that starts at the same vertex. Whenever the random walk \tilde{X}_t is at a vertex in $V_2 \setminus \{N(\hat{z})\}$, then the random walk X_t makes the same transition. If the random walk \tilde{X}_t is at vertex $z' \in N(\hat{z})$, then there is a coupling so that the random walk X_t makes the same transition as \tilde{X}_t with probability $\frac{2\sqrt{n}}{2\sqrt{n}+2}$. Conditional on the event $\{\tilde{Z} \leq \gamma\}$ occurring, the random walk \tilde{X}_t follows the random walk X_t up until step n/c with probability at least

$$p_1 := \left(\frac{2\sqrt{n}}{2\sqrt{n}+2} \right)^\gamma = \left(1 - \frac{1}{\sqrt{n}+1} \right)^\gamma \geq \frac{3}{4},$$

since the random walk \tilde{X}_t has at most γ visits to $N(\hat{z})$. Consider now the random walk $(\tilde{Y}_t)_{t \geq t_{\text{sep}}}$ using \tilde{P} (i.e., restricted to V_2) starting at $Y_{t_{\text{sep}}}$, i.e., $\tilde{Y}_{t_{\text{sep}}} = Y_{t_{\text{sep}}}$. By an analogous argument as before we can couple $(Y_t)_{t \geq t_{\text{sep}}}$ and $(\tilde{Y}_t)_{t \geq t_{\text{sep}}}$ for n/c time steps w.p. at least p_1 .

Furthermore, after $t_{\text{sep}}(G_2) = O(1)$ steps we can couple \tilde{X}_t and \tilde{Y}_t with nodes drawn independently from π^{G_2} . Hence,

$$p_2 := \mathbb{P} \left[\tilde{X}_{t+t_{\text{sep}}(G_2)} = \tilde{Y}_{t+t_{\text{sep}}(G_2)} \mid \mathcal{F}_t \right] \geq (1 - 1/e)^2 \|\pi^{G_2}\|_2^2 \geq \frac{\sqrt{\alpha'}}{8n}.$$

Recall that $\alpha' \geq 2^{20} t_{\text{sep}}(G_2)^2$ by definition. Therefore, the probability that \tilde{X}_t and \tilde{Y}_t do not meet in the time-interval $[t_{\text{sep}}(G_1), t_{\text{sep}}(G_1) + n/c - 1]$ is at most

$$p_3 := (1 - p_2)^{\lfloor n/(t_{\text{sep}}(G_2)c) \rfloor} \leq (1 - p_2)^{\lfloor 2^{10}n/(\sqrt{\alpha'}c) \rfloor} \leq 1/4.$$

Therefore, by the Union bound,

$$\mathbb{P} \left[\bigcup_{t=0}^{t_{\text{sep}}(G_1)+n/c-1} X_t = Y_t \right] \geq p_0 \cdot \left(1 - \mathbb{P} \left[\tilde{Z} \geq \gamma \right] - 2 \cdot (1 - p_1) - p_3 \right) = \Omega(\alpha').$$

Repeating this $O(1/p_3)$ times and using the independence yields that the expected meeting time is $O((t_{\text{sep}}(G_1) + n/c - 1)/p_3) = O(\alpha'n)$ and the proof is complete. \square

Finally, we analyze the coalescing time of G . The proof idea is to consider $\sqrt[5]{n}$ random walks starting from π and show that meetings only occur on V_2 and that at least one random walk requires $\Omega(\sqrt{\alpha'} \cdot n \log n)$ time-steps to reach V_2 .

Lemma 3.14. *Let G be the graph described at the beginning of Section 3.4. We have $t_{\text{coal}}(G) = \Omega(\sqrt{\alpha'} \cdot n \log n)$.*

Proof. Let $\varepsilon = 1/5$. We show that even the coalescing time of n^ε random walks requires $\Omega(\sqrt{\alpha'} \cdot n \log n)$ time-steps w.c.p.. Let R be a collection of n^ε independent, *i.e.*, non-interacting, random walks with starting positions drawn from the stationary distribution π . We define the following three bad events:

- (i) Let \mathcal{E}_1 be the event that any of the n^ε random walks meet on a node $V \setminus V_2$ in $\sqrt{\alpha'} \cdot n \log^2 n$ steps.
- (ii) Let \mathcal{E}_2 be the event that fewer than $n^\varepsilon/4$ random walks start on copies of G_1 , *i.e.*, on nodes $V \setminus (V_2 \cup \hat{z})$.
- (iii) Let \mathcal{E}_3 be the event that all random walks starting from a copy of G_1 require fewer than $c \cdot \sqrt{\alpha'} \cdot n \log n$ time-steps for leaving $V \setminus (V_2 \cup z^*)$ for some constant $c > 0$ to be determined later.

In the following we show that $\mathbb{P}[\mathcal{E}_1] = o(1)$, $\mathbb{P}[\mathcal{E}_2] = o(1)$, and $\mathbb{P}[\mathcal{E}_3 \mid \overline{\mathcal{E}_2}] < 1/e$, which implies, by union bound,

$$\mathbb{P}[\overline{\mathcal{E}_1} \cap \overline{\mathcal{E}_2} \cap \overline{\mathcal{E}_3}] \geq \mathbb{P}[\overline{\mathcal{E}_1}] - (1 - \mathbb{P}[\overline{\mathcal{E}_2} \cap \overline{\mathcal{E}_3}]) \geq 1 - o(1) - \left(1 - (1 - o(1)) \cdot \left(1 - \frac{1}{e}\right)\right) \geq 1 - \frac{1}{2e}.$$

Conditioning on $\overline{\mathcal{E}_1} \cap \overline{\mathcal{E}_2} \cap \overline{\mathcal{E}_3}$, none of the independent random walks meet on any node $V \setminus V_2$ and hence they are indistinguishable from coalescing random walks until they reach V_2 . Therefore, it is necessary for all random walks to reach G_2 in order to coalesce. Hence, we conclude that $t_{\text{coal}}(G) = \Omega(\sqrt{\alpha'} \cdot n \log n)$ yielding the lemma.

- (i) We now prove $\mathbb{P}[\mathcal{E}_1] = o(1)$. Consider any pair of the random walks R . Since both random walks start from the stationary distribution, the probability for them to meet on a node on \hat{z} in a fixed step $t \geq 0$ is at most $O(1/n^2)$.

Hence, by the Union bound over $\binom{n^\varepsilon}{2}$ pairs of random walks and $\sqrt{\alpha'} \cdot n \log^2 n \leq n \log^3 n$ steps, the probability of any two random walks meeting on \hat{z} is at most

$$p_1 := \binom{n^\varepsilon}{2} \cdot n \log^3 n \cdot O(1/n^2) = o(1),$$

since $\varepsilon = \frac{1}{5}$. Furthermore, the probability that no two walks start on the same copy of G_1 is at most $p_2 := n^\varepsilon \cdot \frac{n^\varepsilon}{\sqrt{n}} = o(1)$ by the Union bound.

Moreover, using a Chernoff bound together with [Lemma 3.9\(ii\)](#), it follows that a random walk visits the vertex z^* at most $10 \log^3 n$ times during $n \log^3 n$ steps with probability at least $1 - n^{-2}$. By the Union bound over all random walks, it follows that w.p. at least $1 - n^{-1}$, each random walk visits at most $10 \log^3 n$ different copies of G_1 , and by construction of G each such copy is chosen uniformly and independently at random among $G_1^1, G_1^2, \dots, G_1^\kappa$. Therefore, the probability that there exists a copy of G_1 which is visited by at least two random walks in $n \log^3 n$ steps is at most

$$p_3 := n^{-1} + n^\varepsilon (10 \log^3 n + 1) \cdot \frac{n^\varepsilon (10 \log^3 n + 1)}{\sqrt{n}} = o(1). \quad (19)$$

Putting everything together, using union bound, yields $\mathbb{P}[\mathcal{E}_1] \leq p_1 + p_2 + p_3 = o(1)$.

- (ii) We now prove $\mathbb{P}[\mathcal{E}_2] = o(1)$. The probability p for each random walk to start on a node of $V \setminus (V_2 \cup \hat{z})$ is $\pi(V \setminus (V_2 \cup \hat{z})) \geq 1/2$. For each of the random walks with label $1 \leq i \leq n^\varepsilon$ we define the indicator variable X_i to be one, if that random walk starts on $V \setminus (V_2 \cup \hat{z})$. Let $X = \sum_{i=1}^{n^\varepsilon} X_i$. We have $\mathbb{E}[X] = n^\varepsilon \cdot \mathbb{E}[X_i] \geq n^\varepsilon/2$. Since the starting positions of the n^ε random walks are drawn independently, by a Chernoff bound

$$\mathbb{P}[\mathcal{E}_2] = \mathbb{P}\left[X \leq \frac{1}{4}n^\varepsilon\right] \leq \mathbb{P}\left[X \leq \mathbb{E}[X]/2\right] \leq e^{-n^\varepsilon/16} = o(1).$$

- (iii) We now prove $\mathbb{P}[\mathcal{E}_3 \mid \overline{\mathcal{E}_2}] < 1/4$. From [Lemma 3.9\(ii\)](#) we get that w.p. at least $1/2$ a random walk starting at any node $u \in V_1$ does not leave G_1 , *i.e.*, does not reach z^* , after $c_1 n$ time-steps for some constant $c_1 > 0$. It is easy to see that the number of visits to \hat{z} required before the random walk hits G_2 instead of returning to G_1 is w.c.p. at least $\sqrt{\alpha'}/2$; this is because the fraction of edges from \hat{z} to G_2 is $\sqrt{n/\alpha'}/(\sqrt{n/\alpha'} + \sqrt{n})$. Using a Chernoff bound, we conclude that any random walk starting at G_1 doesn't hit G_2 during the first $T = c_1 \cdot \sqrt{\alpha'}n/2$ time-steps with constant probability $p > 0$. Thus the probability that a random walk does not reach G_2 after $\lambda \cdot T$ time-steps is at least p^λ , for any integer $\lambda \geq 1$. Setting $\lambda = \varepsilon \cdot \log(1/p) \cdot \log(n/4)$, the probability that all of the at least $\frac{1}{4}n^\varepsilon$ random walks starting from G_1 reach G_2 within $\lambda \cdot T = \Omega(\sqrt{\alpha'} \cdot n \log n)$ steps is

$$\mathbb{P}[\mathcal{E}_3 \mid \overline{\mathcal{E}_2}] \leq (1 - p^\lambda)^{\frac{1}{4}n^\varepsilon} \leq 1/e,$$

completing the proof. □

The following lemma establishes a bound on the average hitting time.

Lemma 3.15. *Let G be the graph described at the beginning of [Section 3.4](#). We have $t_{\text{avg-hit}} = \Omega(n^{3/2})$*

Proof. Consider a random walk that starts from an arbitrary vertex $u \in V$. By [Lemma 3.9\(ii\)](#), every time a vertex z^i is visited, with probability at least $c > 0$ it takes $\Omega(n)$ time-steps to visit another vertex z^j , $j \neq i$. Using a Chernoff bound, it follows that with probability larger than $1/2$ it takes at least $\Omega(n^{3/2})$ time-steps to visit at least half of the nodes in $\{z^1, z^2, \dots, z^\kappa\}$. By symmetry, it follows that for every vertex in a copy of G_1 there are $\Omega(n)$ vertices to which the hitting time is $\Omega(n^{3/2})$. Thus, by symmetry, $t_{\text{avg-hit}} = \sum_{u,v \in V} \pi(u) \cdot \pi(v) \cdot t_{\text{hit}}(u, v) = \Omega(n^2 \frac{1}{n^2} n^{3/2}) = \Omega(n^{3/2})$. □

4 Bounding $t_{\text{coal}} = O(t_{\text{hit}})$ for Almost-Regular Graphs

As mentioned in the introduction, the bound on t_{coal} in terms of t_{hit} will be based on the combination of two reduction results; the first result reduces the number of walks from $O(n)$ to $O(\log^3 n)$, while the second one reduces the number of walks from $O(\log^3 n)$ to $(\Delta/d)^{100}$; both taking $O(t_{\text{hit}})$ time. In [Section 4.1](#), we first develop concentration inequalities that will be needed for these reductions. Then in [Section 4.2](#), we present the first and technically simpler reduction to $O(\log^3 n)$ walks, which is stated in [Theorem 4.3](#). The proof basically combines the concentration inequalities with our well-known formula [\(1\)](#).

The derivation of the second reduction is done in [Section 4.3](#). It is based on identifying nearly-regular and dense subsets S , which will contain enough vertices visited by a random walk, even if

the walks only run for $o(t_{\text{hit}})$ steps (Lemma 4.7). The proof of this Lemma 4.7 also rests on the concentration inequalities we derive. The second reduction is then completed by Lemma 4.8, which uses the dense subsets S provided by Lemma 4.7 in order to prove that random walks are likely to collide. A more detailed proof outline can be found at the beginning of Section 4.3.

4.1 Concentration Inequalities for Random Walks

In this part, we derive several concentration inequalities for random walks that are new to the best of our knowledge. We point out that existing concentration inequalities tend to fail in our setting, since the events we are considering (like visits to a certain vertex or expected collisions with an unexposed walk) may only appear a small number of times during t_{hit} steps. Therefore, we have to develop new concentration inequalities that are parameterized by t_{hit} . Although the derivation is fairly elementary, the bounds are quite general and may complement existing bounds that are usually parameterized by the mixing time [CLLM12; Lez89]. In particular, our bounds are most useful when t_{mix} and t_{hit} are close, which is precisely the challenging regime for proving $t_{\text{coal}} = O(t_{\text{hit}})$. One limitation though is that our bounds only work for large deviations exceeding the expectation by a multiplicative factor.

Lemma 4.1. *Let $f : V \rightarrow [0, 1]$ be any function over the vertices and $\bar{f} = \sum_{u \in V} f(u) \cdot \pi(u)$. Then for any random walk starting from an arbitrary vertex X_0 and any number of steps $T \geq 0$,*

$$\mathbb{E} \left[\sum_{t=0}^{T-1} f(X_t) \right] \leq 8 \cdot T^+ \cdot \bar{f},$$

where $T^+ = \max\{t_{\text{hit}}, T\}$. Furthermore, for any integer $\lambda \geq 1$,

$$\mathbb{P} \left[\sum_{t=0}^{T-1} f(X_t) \geq \lambda \cdot (16 \cdot T^+ \cdot \bar{f} + 1) \right] \leq 2^{-\lambda}.$$

Moreover, suppose we have time-dependent functions, $f_t : V \rightarrow [0, 1]$, $0 \leq t \leq T$, where T may be any integer. Further assume that there is a universal bound $\Upsilon > 0$ so that for any $1 \leq s \leq T$ and any $w \in V$,

$$\mathbb{E} \left[\sum_{t=s}^{T-1} f_t(X_t) \mid X_s = w \right] \leq \Upsilon.$$

Then, again for any integer $\lambda \geq 1$,

$$\mathbb{P} \left[\sum_{t=s}^{T-1} f_t(X_t) \geq \lambda \cdot (2\Upsilon + 1) \right] \leq 2^{-\lambda}$$

Proof. We first prove that for all pairs of states $u, v \in V$ and any $T \geq t_{\text{hit}}$ that,

$$\sum_{t=0}^{T-1} p_{u,v}^t \leq 8 \cdot T^+ \cdot \pi(v).$$

Suppose for a sake of contradiction that $\sum_{t=0}^{T-1} p_{u,v}^t > 8 \cdot T^+ \cdot \pi(v)$. Then, for an arbitrary vertex $w \in V$, by Markov's inequality, $\mathbb{P}[T_{\text{hit}}(w, u) \leq 2t_{\text{hit}}] \geq \frac{1}{2}$, where we recall that $T_{\text{hit}}(w, u)$ is the

first time step at which a random walk starting at w hits u . We will use $N_t(u, v) = \sum_{i=0}^{t-1} \mathbf{1}_{X_i=v}$ to denote the number of visits to v up step $t - 1$ starting at u . Therefore,

$$\begin{aligned} \mathbb{E}[N_{3T^+}(X_0, v) \mid X_0 = w] &\geq \mathbb{P}[T_{\text{hit}}(w, u) \leq 2t_{\text{hit}}] \cdot \mathbb{E}[N_{T^+}(X_0, v) \mid X_0 = u] \\ &\geq \frac{1}{2} \cdot \sum_{t=0}^{T-1} p_{u,v}^t > 4 \cdot T^+ \cdot \pi(v). \end{aligned}$$

Since this holds for every vertex $w \in V$, we conclude $\mathbb{E}[N_{3T^+}(X_0, v) \mid X_0 \sim \pi] > 4T^+ \cdot \pi(v)$. However, by definition of the stationary distribution, we also have $\mathbb{E}[N_{3T^+}(X_0, v) \mid X_0 \sim \pi] = 3T^+ \cdot \pi(v)$, which yields the desired contradiction. Now the first statement of the lemma follows simply by linearity of expectations:

$$\begin{aligned} \mathbb{E}\left[\sum_{t=0}^{T-1} f(X_t) \mid X_0 = w\right] &= \mathbb{E}\left[\sum_{t=0}^{T-1} \sum_{u \in V} \mathbf{1}_{X_t=u} \cdot f(u) \mid X_0 = w\right] \\ &= \sum_{u \in V} \sum_{t=0}^{T-1} f(u) \cdot \mathbb{E}[\mathbf{1}_{X_t=u} \mid X_0 = w] \\ &= \sum_{u \in V} f(u) \cdot \sum_{t=0}^{T-1} p_{w,u}^t \\ &\leq \sum_{u \in V} f(u) \cdot 8T^+ \cdot \pi(u) = 8 \cdot T^+ \cdot \bar{f}. \end{aligned}$$

We now prove the second statement. By Markov's inequality, for every $w \in V$,

$$\mathbb{P}\left[\sum_{t=0}^{T-1} f(X_t) \geq 16 \cdot T^+ \cdot \bar{f} \mid X_0 = w\right] \leq \frac{1}{2}.$$

Hence with $\tau := \min\{s \in \mathbb{N} : \sum_{t=0}^s f(X_t) \geq 16 \cdot T^+ \cdot \bar{f}\}$ we have for every $w \in V$,

$$\mathbb{P}[\tau \leq T - 1 \mid X_0 = w] \leq \frac{1}{2}.$$

Since f is bounded by 1, we get $\sum_{t=0}^{\tau} f(X_t) \leq 16 \cdot T^+ \cdot \bar{f} + 1$ and therefore,

$$\mathbb{P}\left[\sum_{t=0}^{T-1} f(X_t) \geq \lambda \cdot (16 \cdot T^+ \cdot \bar{f} + 1) \mid X_0 = w\right] \leq \left(\max_{v \in V} \mathbb{P}[\tau \leq T - 1 \mid X_0 = v]\right)^\lambda \leq 2^{-\lambda}.$$

The third statement is derived in exactly the same way we proved the second statement. \square

The third statement of [Lemma 4.1](#) is very useful in that it can be used the following concentration inequality on \tilde{Z} . Notice that the variable random variable \tilde{Z} is defined using only one random walk $(X_t)_{t \geq 0}$, but it can be viewed as the expected number of collisions on the vertex set S of the random walk $(X_t)_{t \geq 0}$ with another (unexposed) random walk $(Y_t)_{t \geq 0}$, starting from the same vertex u .

Lemma 4.2. *Let S be any subset of vertices such that the degree of any pair of vertices in S differs by at most a factor of γ . Consider any random walk $(X_t)_{t \geq 0}$ that starts at an arbitrary vertex $u \in S$, and for any $T \geq 0$*

$$\tilde{Z} := \sum_{t=0}^{T-1} \mathbf{1}_{X_t \in S} \cdot p_{u, X_t}^t.$$

Then with $T^+ = \max\{t_{\text{hit}}, T\}$ it holds that

$$\mathbb{E} \left[\tilde{Z} \right] \leq 16\gamma \cdot T^+ \cdot \max_{w \in S} \pi(w) =: \Upsilon.$$

Furthermore, for any $\lambda \geq 1$,

$$\mathbb{P} \left[\tilde{Z} \geq \lambda \cdot (2\Upsilon + 1) \right] \leq 2^{-\lambda}.$$

Proof. First note that \tilde{Z} is a random variable over the walk $(X_0 = u, X_1, X_2, \dots, X_{T-1})$ with $u \in S$. Let us first upper bound the expectation of \tilde{Z} :

$$\mathbb{E} \left[\tilde{Z} \right] \leq \sum_{t=0}^{T-1} \sum_{v \in S} p_{u,v}^t \cdot p_{u,v}^t \leq \gamma \cdot \sum_{t=0}^{T-1} \sum_{v \in S} p_{u,v}^t \cdot p_{v,u}^t \leq \gamma \cdot \sum_{t=0}^{T-1} p_{u,u}^{2t} \leq 8\gamma \cdot T^+ \cdot \pi(u),$$

where the second inequality is due to reversibility, *i.e.*, $p_{u,s}^t \cdot \pi(u) = p_{s,u}^t \cdot \pi(s)$ and the fact that the degrees in S differ by a factor of at most γ , and the fourth inequality uses $p_{u,u}^{2t} \leq \pi(u)$ which hold since $p_{u,u}^t$ is non-decreasing ([Lemma A.1](#)) and the first statement of [Lemma 4.1](#).

Furthermore, suppose now that we condition on the walk $(X_t)_{t \geq 0}$ being at an arbitrary vertex $w \in S$ at step s , where $1 \leq s \leq T-1$. Then the remaining contribution towards \tilde{Z} is at most

$$\begin{aligned} \mathbb{E} \left[\sum_{t=s}^{T-1} \mathbf{1}_{X_t \in S} \cdot p_{u,X_t}^t \mid X_s = w \right] &= \sum_{t=s}^{T-1} \sum_{v \in S} p_{w,v}^{t-s} \cdot p_{u,v}^t \\ &\leq \gamma \cdot \sum_{t=s}^{T-1} \sum_{v \in S} p_{w,v}^{t-s} \cdot p_{v,u}^t \\ &\leq \gamma \cdot \sum_{t=s}^{T-1} \sum_{v \in V} p_{w,v}^{t-s} \cdot p_{v,u}^t \\ &= \gamma \cdot \sum_{t=s}^{T-1} p_{u,w}^{2t-s} \\ &\leq \gamma \cdot \sum_{t=0}^{2T-2} p_{u,w}^t \\ &\leq \gamma \cdot 16 \cdot T^+ \cdot \pi(w) \leq \Upsilon, \end{aligned}$$

where the penultimate inequality is due to the first statement of [Lemma 4.1](#), applied to the number of visits to w of a random walk of length T , *i.e.*, $f(v) := \mathbf{1}_{v=w}$. Finally, by the third statement of [Lemma 4.1](#), applied to the functions $f_t(v) := \mathbf{1}_{v \in S} \cdot p_{u,v}^t$, $0 \leq t \leq T-1$,

$$\mathbb{P} \left[\tilde{Z} \geq \lambda \cdot \left(32\gamma \cdot T^+ \cdot \max_{w \in S} \pi(w) + 1 \right) \right] \leq 2^{-\lambda}.$$

□

4.2 Reducing the Walks from n to $O(\log^3 n)$ in $O(t_{\text{hit}})$

We now present our first reduction result that reduces the number of walks from n to $O(\log^3 n)$ in $O(t_{\text{hit}})$ time.

Theorem 4.3. *Let $G = (V, E)$ be an arbitrary, possibly non-regular, graph. Then after $O(t_{\text{hit}})$ steps, the number of walks can be reduced from n to $O(\log^3 n)$ with probability at least $1 - n^{-1}$.*

Thanks to $t_{\text{meet}} \leq 4t_{\text{hit}}$ (Proposition B.9) and $t_{\text{coal}}(S_0) = O(t_{\text{meet}} \log |S_0|)$ (Proposition 3.4), the result of Theorem 4.3 implies, among other things, a bound of $t_{\text{coal}} = O(t_{\text{hit}} \cdot \log \log n)$ for any graph. The proof idea is as follows. First, we use the concentration inequalities of the previous section to show for a given random walk $(X_t)_{t \geq 0}$, there exists w.h.p. a set $S' = S'((X_t)_{t \geq 0})$ of nodes where (i) all nodes have up to a factor of 2 the same degree, (ii) the stationary mass of that set is at least $\pi(S') \geq 1/\log^3 n$, and (iii) the nodes S' receives at least $\Omega(t_{\text{hit}}/\log n)$ visits during the interval $[t_{\text{sep}}, t_{\text{sep}} + 2t_{\text{hit}}]$. From this we will be able conclude that any random walk collides with $(X_t)_{t \geq 0}$ w.p. at least $p = \Omega(1/\log^2 n)$. Second, we consider the process P_{imm} of Section 3.1 and make use of the majorization given by Proposition 3.1. We divide the walks into two sets \mathcal{G}_1 and \mathcal{G}_2 with $|\mathcal{G}_1| = \Theta(\log^3 n)$. We show, using the first part, that w.h.p. each walks $(X_t)_{t \geq 0}$ of \mathcal{G}_2 will vanish due to its frequent visits to S' and the fact that each independent random walk $(Y_t)_{t \geq 0}$ of \mathcal{G}_1 intersects with $(X_t)_{t \geq 0}$ on S' w.p. at least p : Using independence, the probability for each walk of \mathcal{G}_2 to survive is $(1 - p)^{|\mathcal{G}_1|} \leq n^{-2}$. The claim then follows by the Union bound.

Proof of Theorem 4.3. First consider any random walk $(X_t)_{t=0}^{t_{\text{sep}}+2t_{\text{hit}}-1}$, that reaches an arbitrary vertex u at time t_{sep} . Next divide all vertices in V into buckets $S_i = \{v \in V : \deg(v) \in [2^{i-1}, 2^i]\}$, where $1 \leq i \leq \log_2 n$. For any bucket i with $\pi(S_i) \leq 1/\log^3 n$, let $Z_i := \sum_{t=t_{\text{sep}}}^{t_{\text{sep}}+2t_{\text{hit}}-1} \mathbf{1}_{X_t \in S_i}$ count the number of visits to S_i . Then by the first statement of Lemma 4.1, $\mathbb{E}[Z_i] \leq 16t_{\text{hit}} \cdot \pi(S_i) \leq 16t_{\text{hit}}/\log^3 n$. By the second statement of Lemma 4.1, it follows that

$$\mathbb{P}[Z_i \geq t_{\text{hit}}/(2 \log_2 n)] \leq \mathbb{P}[Z_i \geq \log^{1.5} n \cdot (32t_{\text{hit}}/\log^3 n + 1)] \leq 2^{-\log^{1.5} n} \leq n^{-2}/(2 \log_2 n),$$

where we used the fact that $t_{\text{hit}} \geq n$ (Lemma A.7). Hence with

$$S := \bigcup_{\substack{1 \leq i \leq \log_2 n : \\ \pi(S_i) \geq 1/\log^3 n}} S_i,$$

it follows by the Union bound that S gets at least $t_{\text{hit}} - \log_2 n \cdot t_{\text{hit}}/(2 \log_2 n) = t_{\text{hit}}/2$ visits with probability at least $1 - n^{-2}/2$.

Let us now consider any S_i with $\pi(S_i) \geq 1/\log^3 n$, and define

$$\tilde{Z}_i(s) := \sum_{t=s}^{t_{\text{sep}}+2t_{\text{hit}}-1} \mathbf{1}_{X_t \in S_i} \cdot p_{X_s, X_t}^{t-s}, \text{ for any } s \in [t_{\text{sep}}, t_{\text{sep}} + 2t_{\text{hit}} - 1].$$

Notice that by Lemma 4.2, $\kappa = 2$, setting $\Upsilon = 16 \cdot \kappa \cdot \max\{t_{\text{hit}}, 2t_{\text{hit}} - 1 - s\} \cdot \max_{w \in S_i} \pi(w) \leq 64t_{\text{hit}} \cdot \max_{w \in S_i} \pi(w)$ we have

$$\mathbb{P}\left[\tilde{Z}_i(s) \geq 520 \log n \cdot t_{\text{hit}} \cdot \max_{w \in S_i} \pi(w)\right] \leq \mathbb{P}\left[\tilde{Z}_i(s) \geq 8 \log n (2\Upsilon + 1)\right] \leq n^{-8},$$

having used the fact that $t_{\text{hit}} \geq 1/\min_{v \in V} \pi(v)$ due to Lemma A.7. Since $t_{\text{hit}} = O(n^3)$ and there are at most $\log_2 n$ buckets, by the Union bound,

$$\mathbb{P}\left[\bigcup_{\substack{1 \leq i \leq \log_2 n : \\ \pi(S_i) \geq 1/\log^3 n}} \bigcup_{t_{\text{sep}} \leq s \leq t_{\text{sep}} + 2t_{\text{hit}} - 1} \left\{\tilde{Z}_i(s) \geq c \log n \cdot t_{\text{hit}} \cdot \max_{w \in S_i} \pi(w)\right\}\right] \leq n^{-2}/2.$$

Hence by the Union bound, with probability at least $1-n^{-2}$, the trajectory $(x_0, x_1, \dots, x_{t_{\text{sep}}+2t_{\text{hit}}-1})$ of $(X_t)_{t=0}^{t_{\text{sep}}+2t_{\text{hit}}}$ is *good*, *i.e.*, its trajectory (i) makes at least $t_{\text{hit}}/2$ visits to S during the steps $[t_{\text{sep}}, t_{\text{sep}} + 2t_{\text{hit}}]$ and (ii) all $\tilde{Z}_i(s)$ are bounded by $c \log n \cdot t_{\text{hit}} \cdot \max_{w \in S_i} \pi(w)$.

In the following, condition on $(X_t)_{t=0}^{t_{\text{sep}}+2t_{\text{hit}}-1}$ being a good random walk, and let us denote by $(x_0, x_1, \dots, x_{t_{\text{sep}}+2t_{\text{hit}}-1})$ the deterministic trajectory. Since S gets at least $t_{\text{hit}}/2$ visits in the time-interval $[t_{\text{sep}}, t_{\text{sep}} + 2t_{\text{hit}} - 1]$, by the pigeonhole principle, there must be at least one bucket S_j with $\pi(S_j) \geq 1/\log^3 n$ so that bucket S_j gets at least $(1/8) \cdot t_{\text{hit}}/\log n$ visits in that time-interval. We shall now prove that any other random walk $(Y_t)_{t=0}^{t_{\text{sep}}+2t_{\text{hit}}-1}$, starting from an arbitrary vertex $w \in V$, collides with this deterministic trajectory on a vertex in S_j in the time-interval $[t_{\text{sep}}, t_{\text{sep}} + 2t_{\text{hit}} - 1]$ with probability at least $\Omega(1/\log^2 n)$. To this end, let us define

$$Z := \sum_{t=t_{\text{sep}}}^{t_{\text{sep}}+2t_{\text{hit}}-1} \mathbf{1}_{x_t \in S_j} \cdot \mathbf{1}_{Y_t = x_t}.$$

Since for any $t \geq t_{\text{sep}}$ we have $p_{w,v}^t \geq \frac{1}{2} \cdot \pi(v)$ for any pair of vertices $w, v \in V$, it follows that

$$\begin{aligned} \mathbb{E} [Z \mid (x_0, x_1, \dots, x_{t_{\text{sep}}+2t_{\text{hit}}-1})] &\geq \sum_{t=t_{\text{sep}}}^{t_{\text{sep}}+2t_{\text{hit}}-1} \sum_{v \in S_j} \mathbf{1}_{x_t = v} \cdot p_{w,v}^t \\ &\geq \sum_{v \in S_j} \sum_{t=t_{\text{sep}}}^{t_{\text{sep}}+2t_{\text{hit}}-1} \mathbf{1}_{x_t = v} \cdot \frac{1}{2} \cdot \pi(v) \\ &\geq \frac{1}{2} \cdot \min_{w \in S_j} \pi(w) \cdot \sum_{v \in S_j} \sum_{t=t_{\text{sep}}}^{t_{\text{sep}}+2t_{\text{hit}}-1} \mathbf{1}_{x_t = v} \\ &\geq \frac{1}{16} \cdot \min_{w \in S_j} \pi(w) \cdot \frac{1}{\log n} \cdot t_{\text{hit}}, \end{aligned}$$

where the last inequality holds because the deterministic path $(x_0, x_1, \dots, x_{t_{\text{sep}}+2t_{\text{hit}}-1})$ makes at least $(1/8) \cdot t_{\text{hit}}/\log n$ visits to S_j .

Furthermore, since the deterministic walk $(x_0, x_1, \dots, x_{t_{\text{sep}}+2t_{\text{hit}}-1})$ satisfies invariant (ii), conditional on Y_t having its first collision with X_t at step s on a vertex $x_s \in S_j$,

$$\begin{aligned} \mathbb{E} [Z \mid (x_0, x_1, \dots, x_{t_{\text{sep}}+2t_{\text{hit}}-1}), Z \geq 1] &\leq \max_{t_{\text{sep}} \leq s \leq t_{\text{sep}}+2t_{\text{hit}}-1} \sum_{t=s}^{t_{\text{sep}}+2t_{\text{hit}}-1} \mathbf{1}_{x_t \in S_j} \cdot p_{x_s, x_t}^{t-s} \\ &= O(\log n \cdot t_{\text{hit}} \cdot \max_{w \in S_j} \pi(w)), \end{aligned}$$

by part (ii) of the definition of a good walk. Combining our last two bounds yields

$$\begin{aligned} \mathbb{P} [Z \geq 1 \mid (x_0, x_1, \dots, x_{t_{\text{sep}}+2t_{\text{hit}}-1})] &= \frac{\mathbb{E} [Z \mid (x_0, x_1, \dots, x_{t_{\text{sep}}+2t_{\text{hit}}-1})]}{\mathbb{E} [Z \mid (x_0, x_1, \dots, x_{t_{\text{sep}}+2t_{\text{hit}}-1}), Z \geq 1]} \\ &= \Omega(1/\log^2 n) =: p \end{aligned} \tag{20}$$

To complete the proof of the theorem, divide the $k > 2/p \cdot \ln n$ random walk arbitrarily into two disjoint groups \mathcal{G}_1 and \mathcal{G}_2 such that $|\mathcal{G}_1| = 2/p \cdot \ln n$. We will analyze the P_{imm} process defined in [Section 3.1](#) in which random walks from \mathcal{G}_1 are immortal. By making use of the majorization

given in [Proposition 3.1](#), to show the claim it suffices to bound the time it takes in P_{imm} for all walks of \mathcal{G}_2 be eliminated.

By the above argument, any random walk $(X_t)_{t=0}^{t_{\text{sep}}+2t_{\text{hit}}-1}$ in \mathcal{G}_2 will be good with probability at least $1 - n^{-2}$. Hence by Markov's inequality, all trajectories of the random walks in \mathcal{G}_2 are good with probability at least $1 - |\mathcal{G}_2|n^{-2}$. Conditioning on this event, [\(20\)](#) shows that any from the random walks in \mathcal{G}_1 collides with the trajectory of any fixed good random walk in \mathcal{G}_2 in the time-interval $[t_{\text{sep}}, t_{\text{sep}} + 2t_{\text{hit}} - 1]$ with probability at least p . By independence of these events across random walks in \mathcal{G}_1 , random walk $(X_t)_{t=0}^{t_{\text{sep}}+2t_{\text{hit}}-1}$ is not eliminated with probability at most

$$(1 - p)^{|\mathcal{G}_1|} = (1 - p)^{2^{\frac{1}{p}} \ln n} \leq n^{-2}.$$

Note, that we neglected the fact that random walks of \mathcal{G}_2 can eliminate each other, which only further decreases the probability of a walks of \mathcal{G}_2 to survive. Combining everything, and using Union bound, it follows that with probability at least $1 - n^{-1}$ all random walks in \mathcal{G}_2 are eliminated. The claim follows by noting that the number of steps used is $t_{\text{sep}} + 2t_{\text{hit}} = O(t_{\text{mix}} + t_{\text{hit}}) = O(t_{\text{hit}})$. \square

4.3 Reducing the Walks from $\log^4 n$ to $(\Delta/d)^{O(1)}$ in $O(t_{\text{hit}})$

Theorem 4.4. *Let $G = (V, E)$ be any graph with maximum degree Δ and average degree d . Then the expected time to reduce the number of walks from $\log^4 n$ to $(\Delta/d)^{100}$ is at most $O(t_{\text{hit}})$.*

4.3.1 Proof Overview

In comparison with [Theorem 4.3](#), the reduction in [Theorem 4.4](#) is more subtle, as there might be a sub-logarithmic number of random walks preventing us from using the simple bucketing-argument into “nearly-regular” partitions used in [Theorem 4.3](#). Furthermore, [Theorem 4.3](#) achieved the reduction in just a single phase of $\Theta(t_{\text{hit}})$ steps: All random walks have w.h.p. a distribution of visits to nodes which is reasonably close to the expectation of visits to these nodes when starting from the stationary distribution.

Here however, we are only able to prove an “exponential” progress and consider periods which can be much shorter than the mixing time. This means that we need to cope with random walks whose distribution may be far from the stationary distribution. Specifically, if there are k random walks left, we will analyze a phase of length $\approx t_{\text{hit}}/\kappa$, with $k = \kappa^{100}$, and show that a constant fraction of random walks will be eliminated. To account for the fact that the random walks are not mixed, we will identify certain “dense” subsets D_0 having the crucial property that each node in D_0 has a sufficiently large stationary mass and all nodes together have a stationary mass which is close to 1.

We then show the existence of a subset of $D_1(i) \subseteq D_0$ which random walk i will pay enough visits to within t_{hit}/κ steps (see first part of the proof of [Lemma 4.7](#)). This is derived via our new concentration inequality ([Lemma 4.1](#)) to show that (i) random walk i does not spend too many steps outside D_0 and (ii) most vertices do not receive much more visits than predicted by the stationary distribution. Thus we end up in a favourable situation where for most walks $1 \leq i \leq k$, we have a subset $D_1(i) \subseteq D_0$ with $|D_1(i)| \geq 2n/\kappa^8$. Since we have $\approx k = \kappa^{100}$ of such walks, an overwhelmingly large fraction of these subsets $S(i)$, $1 \leq i \leq k$, must overlap.

Unfortunately, we are still not done since in order to reduce the number of random walks, we also need to consider *when* the visits to $D_1(i)$ occur. Specifically, if random walk i makes a visit to a vertex $u \in D_1(i)$ at time, say, t , then we need to ensure that there are enough other random walks j which could potentially also visit vertex u at time t . To ensure this, we will discard “surprising” visits, which are visits to vertices when the probability for this to happen at this step or before is

at most κ^{-23} . It is worth pointing out that the property of a visit to u being surprising, depends not only on the vertex but also on the start vertex of the walk. The second part of the proof of [Lemma 4.7](#) deals with this issue and shows that for most walks, there is a subset $D_2(i) \subseteq D_1(i)$ with $|D_2(i)| \geq n/\kappa^8$ containing only vertices which receive enough “unsurprising” visits.

Equipped with these subsets $D_2(i)$, we regard the “unsurprising” visits as a balls-into-bins configuration, where each ball on a bin (vertex) is associated to a walk i which may visit this vertex (we refer to [Figure 5](#) for an illustration). Through a series of counting arguments [Lemma 4.8](#), we establish that for most random walk i there is a subset $D_4(i) \subseteq D_2$ of vertices, so that each vertex receives enough visits and for each such visit at some time t , there are sufficiently many other walks j which have a probability of at least κ^{-23} each to visit the same vertex at time t .

After all these preparations, we analyze the coalescing process and achieve the desired reduction in the number of the random walks in [Lemma 4.9](#). Similarly to previous analyses, we use a division of random walks into groups \mathcal{G}_1 and \mathcal{G}_2 . The roles of \mathcal{G}_1 and \mathcal{G}_2 are as before; walks in \mathcal{G}_1 are merely used to eliminate walks in \mathcal{G}_2 . This time, however, the division into \mathcal{G}_1 and \mathcal{G}_2 is completely uniformly at random, in particular, this means that \mathcal{G}_1 and \mathcal{G}_2 are roughly of the same size. We establish that for most fixed random walks $1 \leq i \leq k$, conditional on being in group \mathcal{G}_2 , there is a constant probability of picking a trajectory that will likely lead to an intersection with any of the other $k - 1$ unexposed random walk.

Combining the two steps of the proof, the structural result in [Lemma 4.7](#) with the probabilistic analysis in [Lemma 4.8](#), it immediately follows that the number of walks can be reduced by a constant factor within $O(t_{\text{hit}}/\kappa)$ steps, yielding [Theorem 4.4](#).

4.3.2 Definitions and Lemmas required to prove [Theorem 4.4](#)

Before giving the formal proof of [Theorem 4.10](#), we introduce additional notation. Recall that k is the number of random walks at a certain time, w.l.o.g., say $t = 0$. Consider a fixed random walk $(X_t)_{t=0}^\infty$ with label $1 \leq i \leq k$, where $k = \kappa^{100}$ that is run for

$$\tau := 4t_{\text{hit}}/\kappa$$

steps, and starts at an arbitrary vertex $X_0 = u_i$. Since we seek to reduce the number of random walks to $(\Delta/d)^{100}$, we assume the following.

Assumption 4.5. *Throughout this section we assume $k^{1/100} = \kappa > \max\{2^{10}, \Delta/d\}$.*

Note that if $\kappa \leq \max\{2^{10}, \Delta/d\}$ the claim follows immediately from [Proposition 3.4](#). We define a “dense” subset of nodes as

$$D_0 := \{v \in V : \pi(v) \geq 1/(n\kappa^2)\}.$$

Clearly,

$$\pi(D_0) \geq 1 - |V \setminus D_0| \cdot 1/(\kappa^2 n) \geq 1 - \frac{1}{\kappa^2}. \quad (21)$$

$$\pi(u) \leq \pi_{\max} \leq \frac{\Delta}{nd} \leq k^{1/100} \cdot \frac{1}{n} = \kappa \cdot \frac{1}{n}, \quad (22)$$

where the penultimate inequality holds since $k^{1/100} > \Delta/d$ by assumption. Hence the degree of any two vertices in D_0 differ by at most a factor of κ^3 .

Before proceeding further, we introduce another piece of notation. For any random walk $1 \leq i \leq k$ denoted by $(X_t)_{t \geq 0}$, we call a time-step $s \in [0, \tau - 1]$ is *bad* if $X_s \in D_0$ and additionally,

$$\sum_{t=s}^{\tau-1} \mathbf{1}_{X_t \in D_0} \cdot p_{X_s, X_t}^{t-s} \geq 16 \cdot \kappa^{10} \cdot t_{\text{hit}}/n.$$

Intuitively, a time-step s is bad, if the expected number of collisions for another random walk starting at vertex $X_s \in D_0$ at step s with the walk (X_s, X_{s+1}, \dots) is too large.

Lemma 4.6. *Consider a random walk $(X_t)_{t=0}^{\tau-1}$ of length $\tau = 4t_{\text{hit}}/\kappa$. Then with probability at least $1 - 2^{-\kappa}$, there are at most $\tau \cdot 2^{-\kappa}$ bad time-steps $t \in [0, \tau]$. Consequently, for a collection of κ^{100} random walks with $\kappa \geq 2^{10}$, all of these walks have at most $\tau \cdot 2^{-\kappa}$ bad time-steps $t \in [0, \tau - 1]$ w.p. at least $1 - 1/\kappa$.*

Proof. First, let us fix any step $s \in [0, \tau - 1]$, and following the notation of Lemma 4.2, let

$$\tilde{Z}(s) := \sum_{t=s}^{\tau-1} \mathbf{1}_{X_t \in D_0} \cdot p_{X_s, X_t}^{t-s}.$$

Then, since $\tau - s \leq t_{\text{hit}}$, by Lemma 4.2, for any vertex $u \in D_0$,

$$\mathbb{E} \left[\tilde{Z}(s) \mid X_s = u \right] \leq 8\kappa^3 \cdot t_{\text{hit}} \cdot \max_{w \in D_0} \pi(w) \leq 8\kappa^4 \cdot t_{\text{hit}}/n =: \Upsilon.$$

Since $t_{\text{hit}} \geq n$ by Lemma A.7, we have $\Upsilon \geq 2$ and the concentration inequality in Lemma 4.2 implies

$$\mathbb{P} \left[\tilde{Z}(s) \geq 16 \cdot \kappa^{10} \cdot t_{\text{hit}}/n \right] \leq \mathbb{P} \left[\tilde{Z}(s) \geq 2\kappa \cdot (2\Upsilon + 1) \right] \leq 2^{-2\kappa}.$$

Now let B denote the number of bad time-steps, *i.e.*,

$$B := \left| \left\{ 0 \leq s \leq \tau - 1 : \tilde{Z}(s) \geq 16 \cdot \kappa^{10} \cdot t_{\text{hit}}/n \right\} \right|.$$

Then, by linearity of expectation

$$\mathbb{E}[B] = \sum_{s=0}^{\tau-1} \mathbb{P} \left[\tilde{Z}(s) \geq 16 \cdot \kappa^{10} \cdot t_{\text{hit}}/n \right] \leq \tau \cdot 2^{-2\kappa},$$

and a simple application of Markov's inequality implies the first part of the claim. For the second part we simply take Union bound over all κ^{100} walks and using that $1 - \kappa^{100}/2^\kappa \geq 1 - 1/\kappa$. \square

Let $N_t(i, v)$ denote the number of visits of the random walk i to v within t time-steps. Let

$$D_1(i) := \{v \in D_0 : N_\tau(i, v) \geq 2\tau/\kappa^4 \cdot \pi(v)\},$$

i.e., D_1 are all vertices $v \in D_0$ that are visited at least $2\tau/\kappa^4 \cdot \pi(v)$ times before time-step τ . Notice that D_1 is a (random) set that depends on the realization of the walk.

Furthermore, let

$$t(i, v) := \min \{t : \mathbb{P}[T_{\text{hit}}(u_i, v) \leq t] \geq \kappa^{-23}\}.$$

Basically $t(i, v)$ is the “smallest” step t so that the probability that random walk i visits vertex v at step t or earlier is bounded below by κ^{-23} . Note that $t(i, v)$ is a deterministic integer that does

not depend on the realization of the walk. With reference to our proof outline, we regard any visit before $t(i, v)$ as a “surprising” visit, while visits at step $t(i, v)$ or later as an “unsurprising” visit.

Let $D_2(i)$ be the set of vertices v that get at least $\tau/\kappa^4 \cdot \pi(v)$ visits between the time steps $t(i, v)$ and τ ; in symbols

$$D_2(i) := \{v \in D_1(i) : N_\tau(i, v) - N_{t(i, v)}(i, v) \geq \tau/\kappa^4 \cdot \pi(v)\}. \quad (23)$$

We now to state a structural lemma, providing lower bounds on the size of $D_2(i)$. Recall that $D_2(i)$ is a subset of the “dense” set D_0 that has a large stationary mass and contains only vertices with sufficiently high degree. This “projection” is not required on regular graphs, where we could simply work with all vertices, *i.e.*, $D_0 = V$. However, for non-regular graphs, the projection on D_0 is essential since on the set D_0 , the random walk will behave sufficiently similar to a random walk on a regular graph.

Lemma 4.7. *Let $G = (V, E)$ be an arbitrary graph and let $\kappa \geq 2^{10}$ be any integer. Consider any random walk $(X_t)_{t \geq 0}$ with label i . Then, we have that $|D_2(i)| \geq n/\kappa^8$ w.p. at least $1 - 6/\kappa$.*

Proof of Lemma 4.7. First we bound the number of visits of walk i to $V \setminus D_0$. To this end, let \tilde{C} be the hits from i to the vertices which are not in D_0 before time step τ , in symbols,

$$\tilde{C} := \sum_{t=0}^{\tau-1} \mathbf{1}_{X_t \notin D_0}.$$

By Lemma 4.1 and using $\tau \leq t_{\text{hit}}$, we derive

$$\mathbb{E} \left[\tilde{C} \right] \leq 8 \max\{t_{\text{hit}}, \tau\} \cdot \pi(V \setminus D_0) \leq 8t_{\text{hit}}/\kappa^2 = 2\tau/\kappa,$$

where we used the fact that $\pi(V \setminus D_0) \leq 1/\kappa^2$ by (21). Hence by Markov’s inequality,

$$\mathbb{P} \left[\tilde{C} \geq \tau/2 \right] \leq \mathbb{P} \left[\tilde{C} \geq \kappa/4 \cdot \mathbb{E} \left[\tilde{C} \right] \right] \leq \frac{4}{\kappa}. \quad (24)$$

Next for any fixed vertex $v \in V$, we know that for any $\lambda \geq 1$ the probability that the random walk makes more than $32\lambda \cdot t_{\text{hit}} \cdot \pi(v)$ visits to v is

$$\mathbb{P} [N_\tau(i, v) \geq 32\lambda \cdot t_{\text{hit}} \cdot \pi(v)] \leq \mathbb{P} [N_\tau(i, v) \geq \lambda \cdot (16 \cdot t_{\text{hit}} \cdot \pi(v) + 1)] \leq 2^{-\lambda}, \quad (25)$$

where we used the facts that $t_{\text{hit}} \cdot \pi(v) \geq 1$ (Lemma A.7) and the second statement of Lemma 4.1. Recall that $N_t(i, v)$ denotes the number of visits of the random walk i to v within t time-steps. Define

$$B := \sum_{v \in V} \mathbf{1}_{N_\tau(i, v) \geq \kappa^2 \cdot t_{\text{hit}} \cdot \pi(v)} \cdot N_\tau(i, v).$$

Then

$$\begin{aligned}
\mathbb{E}[B] &= \sum_{v \in V} \sum_{j=\kappa^2}^{\infty} j \cdot \mathbb{P}[N_\tau(i, v) = j] \\
&\leq \sum_{v \in V} \sum_{\sigma=\kappa^2}^{\infty} (\sigma+1) \cdot t_{\text{hit}} \cdot \pi(v) \cdot \sum_{j=\sigma \cdot t_{\text{hit}} \cdot \pi(v)}^{(\sigma+1) \cdot t_{\text{hit}} \cdot \pi(v) - 1} \mathbb{P}[N_\tau(i, v) = j] \\
&\leq \sum_{v \in V} \sum_{\sigma=\kappa^2}^{\infty} (\sigma+1) \cdot t_{\text{hit}} \cdot \pi(v) \cdot \mathbb{P}[N_\tau(i, v) \geq \sigma \cdot t_{\text{hit}} \cdot \pi(v)] \\
&= \sum_{v \in V} t_{\text{hit}} \cdot \pi(v) \cdot \sum_{\sigma=\kappa^2}^{\infty} (\sigma+1) \cdot \mathbb{P}[N_\tau(i, v) \geq \sigma \cdot t_{\text{hit}} \cdot \pi(v)] \\
&\stackrel{(25)}{\leq} \sum_{v \in V} t_{\text{hit}} \cdot \pi(v) \cdot \sum_{\sigma=\kappa^2}^{\infty} 2\sigma \cdot 2^{-\sigma/32} \\
&= t_{\text{hit}} \cdot \sum_{\sigma=\kappa^2}^{\infty} 2\sigma \cdot 2^{-\sigma/32} \\
&\leq 2t_{\text{hit}} \cdot \kappa^{-10} \leq \tau \cdot \kappa^{-9},
\end{aligned}$$

where we used that $\kappa \geq 2^{10}$ and $\sum_{v \in V} \pi(v) = 1$. Hence, by Markov's inequality,

$$\mathbb{P}[B \geq \tau \cdot \kappa^{-8}] \leq \frac{1}{\kappa}. \quad (26)$$

Suppose now that $\{\tilde{C} \leq \tau/2\}$ and $\{B \leq \tau \cdot \kappa^{-8}\}$ both occur. Conditioning on $B \leq \tau/\kappa^8$ and $\tilde{C} \leq \tau/2$, we will show by pigeonhole principle

$$|D_1(i)| \geq 2n/\kappa^8. \quad (27)$$

Suppose for the sake of contradiction that $|D_1(i)| < 2n/\kappa^8$. Then, using $\pi_{\max} \leq \kappa/n$ by (22), we have that the total number of visits to nodes in D is at most

$$\begin{aligned}
\sum_{u \in D_0} N_\tau(i, u) &= \sum_{u \in D_0 \setminus D_1(i)} N_\tau(i, u) + \sum_{u \in D_1(i)} N_\tau(i, u) \\
&\leq \sum_{u \in D_0 \setminus D_1(i)} 2\pi(u) \cdot \tau/\kappa^4 + \sum_{u \in D_1(i): N_\tau(i, u) \leq \kappa^2 \cdot t_{\text{hit}} \cdot \pi(u)} N_\tau(i, u) + \sum_{u \in D_1(i): N_\tau(i, u) > \kappa^2 \cdot t_{\text{hit}} \cdot \pi(u)} N_\tau(i, u) \\
&\leq \sum_{u \in D_0 \setminus D_1(i)} 2\pi(u) \cdot \tau/\kappa^4 + |D_1(i)| \cdot \max_{u \in D_0} \pi(u) \cdot t_{\text{hit}} \cdot \kappa^2 + B \\
&\leq 2(|D_0| - |D_1(i)|) \cdot \max_{u \in D_0} \pi(u) \cdot \tau/\kappa^4 + |D_1(i)| \cdot \max_{u \in D_0} \pi(u) \cdot t_{\text{hit}} \cdot \kappa^2 + B \\
&\leq 2(|D_0| - |D_1(i)|) \cdot (\kappa/n) \cdot \tau/\kappa^4 + |D_1(i)| \cdot (\kappa/n) \cdot t_{\text{hit}} \cdot \kappa^2 + B \\
&< 2(|D_0| - |D_1(i)|) \cdot \tau/(n \cdot \kappa) + 2n/\kappa^8 \cdot (\kappa/n) \cdot t_{\text{hit}} \cdot \kappa^2 + B \\
&\leq 2\tau/\kappa + 2n/\kappa^8 \cdot (\kappa/n) \cdot \tau \cdot \kappa \cdot \kappa^2 + \tau/\kappa^8 \\
&= 2\tau/\kappa + \tau/\kappa^2 + \tau/\kappa^8 \leq \frac{1}{2}\tau.
\end{aligned}$$

Thus, $\tilde{C} > \tau - \frac{1}{2}\tau = \frac{1}{2}\tau$ which is a contradiction to the assumption that the event $\{\tilde{C} \leq \tau/2\}$ occurs.

Finally, we will upper bound the number of “surprising” visits, which are visits to vertices that happen too early. That is, we will upper bound the number of visits to vertices v before time $t(i, v)$; in symbols,

$$\tilde{B} := \sum_{v \in V} \sum_{0 \leq t < t(i, v)} \mathbf{1}_{X_t = v}.$$

By definition of $t(i, v)$, with probability at most κ^{-23} the random walk visits the vertex v before $t(i, v)$. Conditional on this event occurring, the expected number of visits is at most $\sum_{t=0}^{t(i, v)-1} p_{v, v}^t$. Hence by linearity of expectations,

$$\mathbb{E} [\tilde{B}] \leq \sum_{v \in V} \left(\kappa^{-23} \cdot \sum_{t=0}^{t(i, v)-1} p_{v, v}^t \right)$$

Since $t(i, v) \leq 2t_{\text{hit}}$, it follows by the first statement of [Lemma 4.1](#) that $\sum_{t=0}^{t(i, v)-1} p_{v, v}^t \leq 16t_{\text{hit}} \cdot \pi(v)$ and hence

$$\mathbb{E} [\tilde{B}] \leq \kappa^{-23} \cdot 16t_{\text{hit}} \cdot \sum_{v \in V} \pi(v) \leq \tau/2 \cdot \kappa^{-16},$$

and thus by Markov’s inequality,

$$\mathbb{P} [\tilde{B} \geq \tau/2 \cdot \kappa^{-15}] \leq \mathbb{P} [\tilde{B} \geq \kappa \cdot \mathbb{E} [\tilde{B}]] \leq \frac{1}{\kappa}. \quad (28)$$

Hence, the total number of visits to vertices in V before time $t(i, v)$ is at most $\tau/2 \cdot \kappa^{-15}$ with probability at least $1 - 1/\kappa$.

Hence, by [\(24\)](#), [\(26\)](#), and [\(28\)](#) and by taking Union bound, we have

$$\begin{aligned} p &:= \mathbb{P} \left[\left\{ \tilde{C} \geq \tau/2 \right\} \cap \left\{ B \geq \tau \cdot \kappa^{-8} \right\} \cap \left\{ \tilde{B} \geq \tau/2 \cdot \kappa^{-15} \right\} \right] \\ &\geq 1 - \mathbb{P} \left[\left\{ \tilde{C} \geq \tau/2 \right\} \cup \left\{ B \geq \tau \cdot \kappa^{-8} \right\} \cup \left\{ \tilde{B} \geq \tau/2 \cdot \kappa^{-15} \right\} \right] \\ &\geq 1 - \frac{6}{\kappa}. \end{aligned}$$

In particular, w.p. p and by [\(27\)](#), we have $|D_1(i)| \geq 2n/\kappa^8$. Observe that, by definition of the sets, each vertex $v \in D_1(i) \setminus D_2(i)$ is visited at least $2\tau/\kappa^4 \cdot \pi(v) - \tau/\kappa^4 \cdot \pi(v) = \pi(v) \cdot \tau/\kappa^4$ times before time-step $t(i, v)$ and thus,

$$|D_2(i)| = |D_1(i)| - |D_1(i) \setminus D_2(i)| \geq |D_1(i)| - \frac{\tilde{B}}{\min_{w \in S} \pi(w) \cdot \tau/\kappa^4} \geq \frac{2n}{\kappa^8} - \frac{\tau/2 \cdot \kappa^{-15}}{1/(\kappa^3 n) \cdot \tau/\kappa^4} \geq \frac{n}{\kappa^8}.$$

□

The next step in the proof is to elaborate on the sets $D_2(i)$ from [Lemma 4.7](#) in order to analyse collisions on this set. Before doing this, we need to introduce additional notation in order to define a balls-into-bins configuration.

Let us denote the random walk with label i by $(X_t^i)_{t=0}^\tau$. The random walk may start from an arbitrary vertex $X_0^i = u_i$ and is run for $\tau = 4t_{\text{hit}}/\kappa$ steps. Recall

$$t(i, v) := \min \left\{ t : \mathbb{P} [T_{\text{hit}}(u_i, v) \leq t] \geq \kappa^{-23} \right\}.$$




Figure 5: Illustration of the balls-into-bins configuration of $k = 4$ walks labeled 1, 2, 3, 4 into 6 vertices v_1, v_2, \dots, v_6 . In the illustration, a ball $(i, t(i, u))$ on bin u is good if there is at least one other ball with $t(i, u) < t(j, u)$ (or $t(i, u) = t(j, u)$ and $i < j$). For each of the random walks 1, 2 and 3, there are at least two good balls, while for random walk 4 only one ball is good.

We now consider the following balls-into-bins configuration, where we emphasise that the balls-into-bins configuration is completely deterministic (for fixed start vertices at time 0) and does not depend on the realization of any of the random walks. Every vertex in V corresponds to a bin. For every walk j and $u \in D_0(j)$, we place deterministically a ball with label $(j, t(j, u))$ into bin u . We call a ball with label $(j, t(j, u))$ in bin u *bad* if there are fewer than κ^{55} other balls i in the same bin such that either (i) $t(i, u) < t(j, u)$ or (ii) $t(i, u) = t(j, u)$ and $i < j$. Next define a random walk j to be *bad* if at least n/κ^9 bad balls have label j and otherwise we call j *good*. Since there are at most $\kappa^{55} \cdot n$ bad balls, it follows that the number of *bad* walks $j \in \{1, \dots, k\}$ is at most

$$\frac{\kappa^{55} \cdot n}{n/\kappa^9} = \kappa^{64}. \quad (29)$$

In the following, we will focus on the $\kappa^{100} - \kappa^{64}$ good walks and ignore all other walks. Recall that any fixed good walk i has at most n/κ^9 bad balls. We now make another central definition of a random subset:

$$D_4(i) := \{v \in D_0 : \text{the ball } (i, t(i, v)) \text{ is good and } N^G(i, v) \geq \tau/(\kappa^9 n)\},$$

where $N^G(i, v)$ denotes the number of times $v \in D_0$ is visited by walk i on a good time step in the interval $[t(i, v), \tau - 1]$. Intuitively, every such visit of a random walk i to a vertex $v \in D_4(i)$ at a time t is very helpful for the following reason: Since the ball $(i, t(i, v))$ is good, there are at least κ^{55} other random walks $j \neq i$ with $t(j, v) \leq t(i, v) \leq t$, and thus each walk j has a probability of at least κ^{-23} to visit vertex v at a time $t(j, v)$ and potentially collide with random walk i at time t later. The next lemma provides a lower bound on the size of $|D_4(i)|$.

Lemma 4.8. *Consider any random walk $(X_t)_{t \geq 0}$ with label i . Then $\mathbb{P}[|D_4(i)| \geq n/(4\kappa^8)] \geq 1 - 8/\kappa$.*

Proof. Recall that, (see (23))

$$D_2(i) = \{v \in D_1(i) : N_\tau(i, v) - N_{t(i, v)}(i, v) \geq \tau/\kappa^4 \cdot \pi(v)\}.$$

Let us now define

$$D_3(i) = \{v \in D_2(i) : \text{the ball } (i, t(i, v)) \text{ is good}\}. \quad (30)$$

By Lemma 4.7, we have $\mathbb{P}[|D_2(i)| \geq n/\kappa^8] \geq 1 - 6/\kappa$. In the following assume that the event $\{|D_2(i)| \geq n/\kappa^8\}$ occurs. Since by definition a good walk has fewer than n/κ^9 bad balls, we have

$$|D_3(i)| = |D_2(i)| - \text{number of bad balls of } i \geq n/\kappa^8 - n/\kappa^9 \geq \frac{n}{2\kappa^8} \quad (31)$$

Further, by definition of $D_2(i)$ each vertex in $v \in D_3(i)$ is visited at least

$$\tau/\kappa^4 \cdot \pi(v) \geq \tau/(\kappa^7 n) \quad (32)$$

times during the interval $[t(i, v), \tau - 1]$, where the inequality is due to the definition of D_0 and the fact that $D_2(i) \subseteq D_0$. We now define the following random variables. Let

1. $N^G(i, u)$ be the number of times $u \in D_3(i)$ is visited by walk i on a good time step in the interval $[t(i, u), \tau - 1]$ (as defined previously).
2. $N^B(i, u)$ be the number of times $u \in D_3(i)$ is visited by walk i on a bad time step in the interval $[t(i, u), \tau - 1]$.
3. $N^E(i)$ be the set of nodes $u \in D_3(i)$ that are visited by walk i prior to $t(i, u)$, *i.e.*, $N^E(i) := \{u \in D_3(i) : T_{\text{hit}}(u_i, v) < t(i, v)\}$.

We have

$$\begin{aligned} \mathbb{E}[N^E(i)] &= \mathbb{E}\left[\sum_{v \in D_3(i)} \mathbf{1}_{T_{\text{hit}}(u_i, v) < t(i, v)}\right] = \sum_{v \in D_3(i)} \mathbb{E}[\mathbf{1}_{T_{\text{hit}}(u_i, v) < t(i, v)}] \\ &= \sum_{v \in D_3(i)} \mathbb{P}[T_{\text{hit}}(u_i, v) < t(i, v)] \leq \sum_{v \in D_3(i)} 1/\kappa^{23} = |D_3(i)|/\kappa^{23}, \end{aligned}$$

where the inequality comes from the definition of $t(i, v)$. By Markov inequality,

$$\mathbb{P}\left[N^E(i) \leq \frac{|D_3(i)|}{\kappa^{22}}\right] \geq 1 - \frac{1}{\kappa}. \quad (33)$$

In the remainder we condition on $N^E(i) \leq \frac{|D_3(i)|}{\kappa^{22}}$. By [Lemma 4.6](#), with probability at least $1 - \kappa$, all random walks have at most $\tau \cdot 2^{-\kappa} \leq t_{\text{hit}}/\kappa^{20}$ bad time-steps $s \in [0, \tau]$, where we recall t is *bad* if $X_s \in D_0$ and $\sum_{t=s}^{\tau} \mathbf{1}_{X_t \in D_0} \cdot p_{X_s, X_t}^{t-s} \geq 16 \cdot \kappa^{10} \cdot t_{\text{hit}}/n$. In the following we condition on the number of bad time steps being bounded by $t_{\text{hit}}/\kappa^{20}$. We claim that

$$|D_4(i)| \geq |D_3(i)|/2. \quad (34)$$

Assume, for the sake of contradiction, that $|D_4(i)| < \frac{1}{2}|D_3(i)|$. Let

$$\tilde{D} := \{v \in D_3(i) \setminus D_4(i) : T_{\text{hit}}(u_i, v) \geq t(i, v)\}.$$

We have, using $|D_3(i) \setminus D_4(i)| \geq |D_3(i)| - |D_4(i)| \geq \frac{1}{2}|D_3(i)|$ that

$$|\tilde{D}| \geq |D_3(i) \setminus D_4(i)| - N^E(i) \geq \frac{1}{2}|D_3(i)| - \frac{|D_3(i)|}{\kappa^{22}} \geq \frac{1}{4}|D_3(i)|.$$

For each vertex $u \in \tilde{D} \subseteq D_3(i) \setminus D_4(i)$, we have $N^G(i, u) < \tau/(2n\kappa^9)$ and thus

$$\begin{aligned} N^B(i, u) &\stackrel{u \in \tilde{D}}{=} N_\tau(i, u) - N_{t(i, v)}(i, u) - N^G(i, u) \\ &\geq \tau/\kappa^4 \cdot \pi(u) - \tau(2n\kappa^9) \geq \tau/(\kappa^7 n) - \tau/(2n\kappa^9) \geq \tau/(2n\kappa^7), \end{aligned}$$

where the first inequality follows from (32) and $u \in D_2(i)$. In words, at least $\tau/(2n\kappa^7)$ visits to u happened on a bad time step during the interval $[t(i, u), \tau]$. Thus, the total number of visits to nodes of \tilde{D} at bad time steps is at least

$$\sum_{v \in \tilde{D}} N^B(i, v) \geq |\tilde{D}| \cdot \tau/(2n\kappa^7) \geq \frac{1}{8} |D_3(i)| \cdot \frac{\tau}{n\kappa^7} \stackrel{(31)}{>} \frac{t_{\text{hit}}}{\kappa^{20}}.$$

This contradicts the assumption that there are at most $t_{\text{hit}}/\kappa^{20}$ bad time steps in total. Thus, (34) holds and we derive using (31)

$$|D_4(i)| \geq n/(4\kappa^8). \quad (35)$$

As shown above, this lower bound on $|D_4(i)|$ holds whenever the following three events all occur: (i) $|D_2(i)| \geq n/\kappa^8$ (which holds with probability at least $1 - 6/\kappa$ by Lemma 4.7), (ii) $N^E(i) \leq \frac{|D_3(i)|}{\kappa^{22}}$ occurs (which holds with probability $1 - 1/\kappa$ by (33)) and (iii) the number of bad time steps of random walk i is at most $t_{\text{hit}}/\kappa^{20}$ (which holds with probability $1 - 1/\kappa$ by Lemma 4.6). Hence by the Union bound,

$$\mathbb{P} [|D_4(i)| \geq n/(4\kappa^8)] \geq 1 - 6/\kappa - 1/\kappa - 1/\kappa = 1 - 8/\kappa. \quad \square$$

The previous lemma established that with reasonably large probability, any fixed good random walk i satisfies $|D_4(i)| \geq n/\kappa^{10}$. In the next lemma we show that, conditioning on this event occurring, that random walk i is eliminated by any of the other random walks with some constant probability > 0 .

Lemma 4.9. *Assume that a good random walk i has a trajectory $(x_0 = u_i, x_1, \dots, x_{\tau-1})$ satisfying $|D_4(i)| \geq n/\kappa^{10}$. Then random walk i will be eliminated before time τ with probability at least $1/10$.*

Proof. Consider now another random walk $j \neq i$ starting from an arbitrary vertex u_j . Define

$$D_5(i, j) := \{u \in D_4(i) : t(j, u) \leq t(i, u)\}. \quad (36)$$

With reference to Figure 5, $u \in D_5(i, j)$ if the green ball $(i, t(i, u))$ lies above the ball $(j, t(j, u))$. Intuitively, $D_5(i, j)$ contains all the vertices in $u \in D_4(i)$ so that each time random walk i visits u , also random walk j could visit that vertex with sufficiently large probability. For $u \in D_4(i)$ we have that $(u, t(i, u))$ is good and for each such good ball, by definition there at least κ^{55} other random walks j such that $t(j, u) < t(i, u)$ (or $t(j, u) = t(i, u)$ and $j < i$). Hence by considering all bins we conclude

$$\sum_{j=1, j \neq i}^{\kappa^{100}} |D_5(i, j)| \geq \kappa^{55} \cdot |D_4(i)| \geq \kappa^{55} \cdot n/(4\kappa^8) = n\kappa^{47}/4, \quad (37)$$

where the second inequality holds by our assumption $|D_4(i)| \geq n/(4\kappa^8)$.

We are now in a position to apply our common analysis method. We consider the P_{imm} process of Section 3.1 and make use of the majorization of Proposition 3.1. We assign each random walk $1 \leq i \leq k = \kappa^{100}$ into group \mathcal{G}_1 and \mathcal{G}_2 independently and uniformly at random. Recall that walks of \mathcal{G}_1 cannot be eliminated. In the following, we define the following event \mathcal{E}_i for walk i : $\mathcal{E}_i = \{i \in \mathcal{G}_2\}$. Clearly, $\mathbb{P}[\mathcal{E}_i] = 1/2$. We will prove that conditional on this event occurring, random walk i is eliminated by one of the walks in \mathcal{G}_1 with at least constant probability > 0 . Let

$Z(i, j)$ denote the number of collisions between random walk i , denoted by $(X_t)_{t \geq 0}$, and j denoted by $(Y_t)_{t \geq 0}$, that happen on a vertex in $D_5(i, j)$ at a good time step, in symbols,

$$Z(i, j) := \mathbf{1}_{i \in \mathcal{G}_2} \cdot \mathbf{1}_{j \in \mathcal{G}_1} \cdot \sum_{u \in D_5(i, j)} \sum_{\substack{t(i, u) \leq t \leq \tau - 1: \\ t \text{ is good}}} \mathbf{1}_{X_t = u} \cdot \mathbf{1}_{Y_t = u}$$

By conditioning on \mathcal{E}_i and the trajectory of $(x_0, x_1, \dots, x_{\tau-1})$ of the good random walk i ,

$$\mathbb{E}[Z(i, j) \mid \text{trajectory } i = (x_0, x_1, \dots, x_{\tau-1}), \mathcal{E}_i] = \frac{1}{2} \cdot \sum_{u \in D_5(i, j)} \sum_{\substack{t(i, u) \leq t \leq \tau - 1: \\ t \text{ is good and } x_t = u}} p_{u_j, u}^t.$$

We now would like to derive a lower bound on $p_{u_j, u}^t$, where $t(i, u) \leq t \leq \tau - 1$ is a good time-step with $X_t^i = u$, $u \in D_5(i, j)$. By definition of $D_5(i, j)$ we have $t(i, u) \geq t(j, u)$. By conditioning on the first visit of random walk j to u , we obtain

$$\begin{aligned} p_{u_j, u}^t &= \sum_{s=0}^t \mathbb{P}[T_{\text{hit}}(u_j, u) = s] \cdot p_{u, u}^{t-s} \stackrel{\text{Lemma A.1}}{\geq} \pi(u) \cdot \sum_{s=0}^t \mathbb{P}[T_{\text{hit}}(u_j, u) = s] \\ &\stackrel{u \in D_0}{\geq} 1/(\kappa^3 \cdot n) \cdot \mathbb{P}[T_{\text{hit}}(u_j, u) \leq t] \stackrel{t \geq t(i, u) \geq t(j, u)}{\geq} 1/(\kappa^3 \cdot n) \cdot \mathbb{P}[T_{\text{hit}}(u_j, u) \leq t(j, u)] \\ &\stackrel{\text{def. of } t(j, u)}{\geq} 1/(\kappa^{26} \cdot n). \end{aligned}$$

By definition of $D_4(i)$, for any vertex $u \in D_5(i, j) \subseteq D_4(i)$, u is visited at least $\tau/(\kappa^9 n)$ times during the interval $[t(i, u), \tau - 1]$. Therefore,

$$\begin{aligned} &\mathbb{E}[Z(i, j) \mid \text{trajectory } i = (x_0, x_1, \dots, x_{\tau-1}), \mathcal{E}_i] \\ &\geq \frac{1}{2} \cdot |D_5(i, j)| \cdot \frac{\tau}{\kappa^9 \cdot n} \cdot \frac{1}{\kappa^{26} \cdot n} = \frac{1}{2} \cdot |D_5(i, j)| \cdot \frac{\tau}{\kappa^{35}} \cdot \frac{1}{n^2}. \end{aligned}$$

Recall that if a time step $s \in [0, \tau - 1]$ is good (*i.e.*, not bad), then $X_s \in D_0$ implies

$$\sum_{t=s}^{\tau-1} \mathbf{1}_{X_t \in D_0} \cdot p_{X_s, X_t}^{t-s} < 16 \cdot \kappa^{10} \cdot t_{\text{hit}}/n.$$

Since $Z(i, j)$ sums only over good time steps and using $D_5(i, j) \subseteq D_0$ we conclude that

$$\begin{aligned} &\mathbb{E}[Z(i, j) \mid \text{trajectory } i = (x_0, x_1, \dots, x_{\tau-1}), \mathcal{E}_i, Z(i, j) \geq 1] \\ &\leq \max_{s: X_s \in D_0} \mathbb{E} \left[\sum_{\substack{s \leq t \leq \tau - 1: \\ t \text{ is good}}} \mathbf{1}_{X_t \in D_5(i, j)} \cdot p_{X_s, X_t}^{t-s} \right] \leq 16\kappa^{10} \cdot \frac{t_{\text{hit}}}{n} \end{aligned}$$

Combining the last two inequalities yields

$$\begin{aligned} &\mathbb{P}[Z(i, j) \geq 1 \mid \text{trajectory } i = (x_0, x_1, \dots, x_{\tau-1}), \mathcal{E}_i] \\ &= \frac{\mathbb{E}[Z(i, j) \mid \text{trajectory } i = (x_0, x_1, \dots, x_{\tau-1}), \mathcal{E}_i]}{\mathbb{E}[Z(i, j) \mid \text{trajectory } i = (x_0, x_1, \dots, x_{\tau-1}), \mathcal{E}_i, Z \geq 1]} \\ &\geq \frac{|D_5(i, j)|}{n \cdot \kappa^{47}} \end{aligned}$$

We are interested in the probability for i being eliminated. Neglecting the possibility that i might even be eliminated by another rand walk of \mathcal{G}_2 , which can only increase the probability of i being eliminated, we derive

$$\begin{aligned}
& \mathbb{P}[\text{Walk } i \text{ is eliminated} \mid \text{trajectory } i = (x_0, x_1, \dots, x_{\tau-1}), \mathcal{E}_i] \\
& \geq \mathbb{P}\left[\bigcup_{j=1, j \neq i}^{\kappa^{100}} \{Z(i, j) \geq 1\} \mid \text{trajectory } i = (x_0, x_1, \dots, x_{\tau-1}), \mathcal{E}_i\right] \\
& \geq 1 - \prod_{j=1, j \neq i}^{\kappa^{100}} \left(1 - \frac{|D_5(i, j)|}{n \cdot \kappa^{47}}\right) \\
& \geq 1 - \exp\left(-\frac{1}{n \cdot \kappa^{47}} \cdot \sum_{j=1, j \neq i}^{\kappa^{100}} |D_5(i, j)|\right) \\
& \stackrel{(37)}{\geq} 1 - \exp(-1/4).
\end{aligned}$$

Note that the above derivation was conditional on \mathcal{E}_i , but this event holds with probability $1/2$. Hence with probability at least $1/2 \cdot (1 - \exp(-1/4)) > 1/10$, the trajectory of i meets with that of a random walk in \mathcal{G}_1 and hence the random walk i is eliminated before time step τ . \square

4.3.3 Completing the Proof of Theorem 4.4

We are now ready to complete the proof of Theorem 4.4 by combining Lemma 4.7, Lemma 4.8 and Lemma 4.9.

Proof of Theorem 4.4. Let $k = \kappa^{100}$ be the number of random walks. Since we seek to reduce the number of random walks to $(\Delta/d)^{100}$, we assume in the following that $k^{1/100} = \kappa > \Delta/d$ and $\kappa \geq 2^{10}$ (cf. Assumption 4.5). Otherwise, if $\kappa < 2^{10}$, then with V_0 denoting the set of start vertices of the k walks, we have $t_{\text{coal}}(V_0) = O(t_{\text{meet}} \cdot \log(|V_0|)) = O(t_{\text{hit}})$, by Proposition 3.4 and Proposition B.9. As derived in (29) we have that out of the κ^{100} random walks at least $\kappa^{100} - \kappa^{64}$ random walks are good. By Lemma 4.8, any good random walk $(X_t)_{t \geq 0}$ with label i satisfies $\mathbb{P}[|D_4(i)| \geq n/(4\kappa^8)] \geq 1 - 8/\kappa$. Conditioning on the trajectory $(x_0, x_1, \dots, x_{\tau-1})$ satisfying $|D_4(i)| \geq n/(4\kappa^8)$, Lemma 4.9 shows that with probability at least $1/10$ the random walk i will be eliminated before time step $\tau = O(t_{\text{hit}}/\kappa)$. Hence a constant fraction of all $k = \kappa^{100}$ random walks are eliminated in a single phase of $O(t_{\text{hit}}/\kappa)$ steps with constant probability > 0 , provided that $\kappa > \Delta/d$.

In conclusion, for any $k' > (\Delta/d)^{100}$, there exists a constant $c > 0$ such that the expected time required to reduce the number of walks from k' to $\max\{k'/2, (\Delta/d)^{100}\}$ is bounded by $c \cdot t_{\text{hit}} / \sqrt[100]{k'}$, by Lemma A.9. Therefore, the expected time to reduce the number of walks from $k' \leq \log^4 n$ to $(\Delta/d)^{100}$ is upper bounded by

$$\sum_{i=0}^{\log(\log^4 n)} \frac{c \cdot t_{\text{hit}}}{\sqrt[100]{2^i}} \leq c \cdot t_{\text{hit}} \sum_{i=0}^{\infty} \left(\frac{1}{\sqrt[100]{2}}\right)^i = \frac{c \cdot t_{\text{hit}}}{1 - \sqrt[100]{1/2}} = O(t_{\text{hit}}).$$

\square

4.4 Bounding t_{coal} in terms of t_{hit}

In this subsection we prove the following theorem relating t_{coal} to t_{hit} . Recall that $t_{\text{meet}} \leq 4t_{\text{hit}}$ (Proposition B.9) for any graph.

Theorem 4.10. *Let $G = (V, E)$ be any graph with maximum degree Δ and average degree d . Then $t_{\text{coal}} = O(t_{\text{hit}} + t_{\text{meet}} \cdot \log(\Delta/d))$. So in particular, for any almost-regular graph, $t_{\text{coal}} = O(t_{\text{hit}})$.*

Theorem 4.10 follows almost immediately from the previous two reductions in Theorem 4.3 and Theorem 4.4.

Proof of Theorem 4.10. By Theorem 4.3, we can reduce the number of walks from n to $O(\log^3 n)$ in $O(t_{\text{hit}})$ steps with probability at least $1 - n^{-1}$. Then, using Theorem 4.4, we can reduce the number of walks from $O(\log^3 n)$ to $(\Delta/d)^{100}$ in $O(t_{\text{hit}})$ expected time. Finally, we apply Proposition 3.4 to reduce the number of walks from $(\Delta/d)^{100}$ to 1 in $O(t_{\text{meet}} \cdot \log(\Delta/d))$ expected time to obtain the result. \square

4.5 Proof of Theorem 1.3

Part (i) follows from Theorem 4.3 together with $t_{\text{meet}} \leq 4t_{\text{hit}}$ (Proposition B.9) and $t_{\text{coal}}(S_0) = O(t_{\text{meet}} \log |S_0|)$ (Proposition 3.4).

Part (ii) is the statement of Theorem 4.10. To prove Part (iii) follows from the following three facts. First, $t_{\text{hit}} = \Theta(t_{\text{meet}})$ by Proposition B.9. Second, $t_{\text{coal}} = O(t_{\text{hit}} + t_{\text{meet}})$ by Part (ii). Third, $t_{\text{coal}} \geq t_{\text{meet}}$. Finally, Part (iv) follows from the results presented in Appendix C.

4.6 Conjecture and a Possible Improvement for Non-Regular Graphs

Before concluding this section, we mention an intriguing conjecture that might be useful to improve our bound on t_{coal} when $\Delta \gg d$.

Conjecture 4.11. *There exists a universal constant $C > 0$ so that for any graph $G = (V, E)$, any vertex $u \in V$ and any path of vertices $(x_0, x_1, \dots, x_{t_{\text{hit}}})$, i.e., either $x_i = x_{i+1}$ or $\{x_i, x_{i+1}\} \in E(G)$,*

$$\sum_{t=0}^{t_{\text{hit}}} p_{u, x_t}^t \leq C \cdot \sum_{t=0}^{t_{\text{hit}}} \pi(x_t).$$

Note that the inequality is a stronger version than the one given in the first statement of Lemma 4.1 or Lemma 4.2. We do not know whether the conjecture is actually true in this generality. However, if it is true, it would imply that any random walk of length $4t_{\text{hit}}$ starting from an arbitrary vertex meets with any deterministic path of length $4t_{\text{hit}}$ with constant probability > 0 . This would then result in a simple proof that $t_{\text{coal}} = O(t_{\text{hit}} \cdot \log^* n)$ for any graph, since each phase of $O(t_{\text{hit}})$ steps would reduce the number of walks from k to $O(\log k)$.

Conjecture 4.11 can be also seen as the optimization problem of “predicting” a random walk $(Y_t)_{t \geq 0}$ for t_{hit} time steps. More precisely, we are given the start vertex of the random walk $Y_0 = u$ and for each time step $1 \leq t \leq t_{\text{hit}}$, we have to specify a vertex x_t that acts as a predictor of the random location of the random walk at step t . The goal is to maximize the (expected) number of correct predictions, which is equal to

$$\sum_{t=0}^{t_{\text{hit}}} p_{u, x_t}^t.$$

The conjecture states that regardless which prediction, i.e., which path $(x_0, x_1, \dots, x_{t_{\text{hit}}})$ is picked, the expected number of correct predictions cannot be made larger than in the setting where the random walks starts from stationarity (and the start vertex is unknown).

Graph	t_{meet}		t_{coal}	
General Graphs	$\Omega(1), O(n^3)$	<i>Thm. B.1</i>	$\Omega(\log n), O(n^3)$	<i>Lem. 5.2 & Thm. 1.3</i>
Regular Graphs	$\Omega(n), O(n^2)$	<i>Thm. B.1 & Thm. 1.3</i>	$\Omega(n), O(n^2)$	<i>Thm. B.1 & Thm. 1.3</i>
Vertex-Trans. Graphs	$\Omega(n), O(n^2)$	<i>Thm. B.1 & Thm. 1.3</i>	$\Omega(n), O(n^2)$	<i>Thm. B.1 & Thm. 1.3</i>

Table 2: A summary of bounds on the meeting and coalescence times graph classes. All bounds are easily shown to be tight: For general graphs the meeting time and coalescence bounds are matched by the star and the barbell graph. For vertex-transitive and regular graphs the bounds are matched by the clique and the cycle.

One specific strategy would be to choose $x_0 = x_1 = \dots = x_{t_{\text{hit}}} = v$ for some vertex v . In that case we know by [Lemma 4.1](#) that

$$\sum_{t=0}^{t_{\text{hit}}} p_{u,v}^t \leq 4 \cdot (t_{\text{hit}} + 1) \cdot \pi(v),$$

so the conjecture holds in this case.

It is also worth mentioning that we cannot replace t_{hit} by a smaller value, say, t_{mix} . Indeed if G is a two-dimensional grid, then $t_{\text{mix}} = \Theta(n)$ and choosing $x_0 = x_1 = \dots = x_{t_{\text{hit}}} = u$, we obtain $\sum_{t=0}^{t_{\text{mix}}} p_{u,u}^t = \Omega(\log n)$, while, $\sum_{t=0}^{t_{\text{mix}}} \pi(x_t) = O(1)$.

Finally, there is some resemblance to the meeting-time-lemma in the continuous-time setting [[Oli12](#)], however, one important difference is that in [Conjecture 4.11](#), the right hand side depends on the actual path $(x_0, x_1, \dots, x_{t_{\text{hit}}})$.

5 Bounding $t_{\text{coal}} \in [\Omega(\log n), O(n^3)]$

Given that worst-case upper and lower bounds have long been known for $t_{\text{mix}}, t_{\text{hit}}$ and t_{cov} , it is very natural to pose the same question for t_{meet} and t_{coal} . In the following we determine the correct asymptotic worst-case upper and lower bounds for t_{meet} and t_{coal} on (i) general graphs, (ii) regular graphs and (iii) vertex-transitive graphs. We refer to [Table 2](#) for an overview.

5.1 General Upper Bound $t_{\text{coal}} = O(n^3)$

In this section we establish that $t_{\text{coal}} = O(n^3)$ on all graphs, which is matched for instance by the Barbell graph.

Theorem 5.1. *For any graph G we have $t_{\text{coal}} = O(n \cdot |E| \cdot \log(|E|/n))$, so in particular, $t_{\text{coal}} = O(n^3)$.*

Proof of Theorem 5.1. It is well-known that $t_{\text{hit}} \leq n \cdot 2|E|$ (cf. [[AKL+79](#)]). From [Proposition B.9](#) and [Theorem 4.10](#) we derive

$$\begin{aligned} t_{\text{coal}} &= O(t_{\text{hit}} + t_{\text{meet}} \cdot \log(\Delta/d)) \\ &= O(t_{\text{hit}} + t_{\text{hit}} \log(n^2/|E|)) \\ &= O(n \cdot |E| + n \cdot |E| \cdot \log(n^2/|E|)) = O(n^3), \end{aligned}$$

where the last inequality holds since $|E| < n^2$. □

5.2 General Lower Bound $t_{\text{coal}} = \Omega(\log n)$

In this section, we prove that the coalescing time of any graph is $\Omega(\log n)$. We consider a process P' where there is exactly one random walk starting at each node in the graph. For every node $u \in V$ and every time step $t \in \mathbb{N}$ we draw an independent random variable $Z_{u,t} \in \{0, 1\}$ with $\mathbb{P}[Z_{u,t} = 1] = 1/2$ and $\mathbb{P}[Z_{u,t} = 0] = 1/2$. If $Z_{u,t} = 1$, then the random walk on u at time t (if there is any), moves to a neighbouring node chosen u.a.r.. Otherwise ($Z_{u,t} = 0$), the random walk on u at time t (if there is any) stays on the same node. It is straightforward to show that the set of nodes which have an active random walk according to this process can be coupled with the coalescence process defined in [Section 2](#).

We show that after $c \log n$ steps, for a sufficiently small c , there are at least two surviving walks in this process. In order to do this, we simply argue that there must be at least two walks that have not left their starting position. Note that there is no way for these walks to be eliminated, because even if other walks visited one of their starting nodes, there are two nodes from which no walks can have left. The formal proof follows.

Lemma 5.2. *For any graph $G = (V, E)$, $|V| = n$ we have $t_{\text{coal}} = \Omega(\log n)$.*

Proof. Consider the process P' defined above. Let T be the coalescence time. Note that coalescence at time τ in P' requires that for $n - 1$ nodes $u \in V$ there exists $t_u \leq \tau$ such $Z_{u,t_u} = 1$. In symbols, let T be the first point in time where all walks coalesced, then $T \geq T'$, with $T' := \min\{t' \in \mathbb{N} : |\{u : \exists t_u \leq t' \text{ s.t. } Z_{u,t_u} = 1\}| \geq n - 1\}$. Let Y_u be the indicator variable which is 1 if $Z_{u,t} = 0$ for all $t \leq \tau := \log n/2$. The process ensures independence of the Y_u . Due to the laziness of the random walk, $\mathbb{P}[Y_u = 1] = 1/2^\tau = 1/\sqrt{n}$. Thus, using the independence of the Y_u ,

$$\mathbb{P}[T \geq \tau] \geq \mathbb{P}[T' \geq \tau] \geq \mathbb{P}\left[\sum_{u \in V} Y_u \geq 2\right] = \mathbb{P}[\text{Binomial}(n, 1/\sqrt{n}) \geq 2] = 1 - o(1),$$

where $\text{Binomial}(n, p)$ denotes the binomial distribution with parameters n and p . We conclude that $\mathbb{E}[T] = \Omega(\log n)$ which yields the claim. \square

5.3 Proof of [Theorem 1.5](#)

We are now ready to put all the pieces together. The upper bound on general graphs follows directly from $t_{\text{meet}} \leq t_{\text{coal}} = O(n^3)$, by [Theorem 5.1](#). The lower bound on the meeting time holds by definition and the lower bound on the coalescing time follows from [Lemma 5.2](#). For the upper bound on regular graphs we have $t_{\text{meet}} \leq t_{\text{coal}} = O(t_{\text{hit}}) = O(n^2)$ due to [Theorem 1.3](#), having used the standard bound $t_{\text{hit}} = O(n^2)$ for regular graphs (see [\[AF02\]](#)). The lower bound follows from $t_{\text{coal}} \geq t_{\text{meet}} \geq t_{\text{meet}}^\pi = \Omega(n)$, by [Theorem B.1](#).

Acknowledgements.

The authors would like to thank Petra Berenbrink, Robert Elsässer, Nikolaos Fountoulakis, Peter Kling, Roberto Oliveira and Perla Sousi for helpful discussions and in particular Yuval Peres for pointing out how to further improve the bound on the hitting time ([Theorem 1.4](#)). Moreover, the authors would like to thank the anonymous reviewer for pointing out a mistake in the domination described in [Section 3.1](#) in an earlier version of this work.

References

- [AF02] D. Aldous and J. Fill. *Reversible Markov Chains and Random Walks on Graphs*. Unpublished. <http://www.stat.berkeley.edu/~aldous/RWG/book.html>. 2002 (cit. on pp. 2–5, 10, 23, 52, 59, 60, 64, 66, 68–70, 72).
- [Ald91] D. Aldous. “Meeting times for independent Markov chains”. In: *Stochastic Processes and their Applications* 38.2 (1991), pp. 185–193 (cit. on p. 1).
- [AKL+79] R. Aleliunas, R. Karp, R. Lipton, L. Lovász, and C. Rackoff. “Random Walks, Universal Traversal Sequences, and the Complexity of Maze Problems”. In: *FOCS*. 1979, pp. 218–223 (cit. on pp. 9, 51).
- [AAK+11] N. Alon, C. Avin, M. Koucký, G. Kozma, Z. Lotker, and M. R. Tuttle. “Many Random Walks Are Faster Than One”. In: *Combinatorics, Probability & Computing* (2011) (cit. on p. 6).
- [BPS12] M. Barlow, Y. Peres, and P. Sousi. “Collisions of Random Walks”. In: *Annales de l’Institut Henri Poincaré* 48.4 (2012), pp. 922–946 (cit. on p. 7).
- [BCN+16] L. Becchetti, A. E. F. Clementi, E. Natale, F. Pasquale, and L. Trevisan. “Stabilizing Consensus with Many Opinions”. In: *SODA*. 2016, pp. 620–635 (cit. on p. 1).
- [BCN+15] L. Becchetti, A. Clementi, E. Natale, F. Pasquale, and R. Silvestri. “Plurality Consensus in the Gossip Model”. In: *SODA*. 2015, pp. 371–390 (cit. on p. 1).
- [BCE+17] P. Berenbrink, A. E. F. Clementi, R. Elsässer, P. Kling, F. Mallmann-Trenn, and E. Natale. “Ignore or Comply?: On Breaking Symmetry in Consensus”. In: *PODC*. 2017, pp. 335–344. URL: <http://doi.acm.org/10.1145/3087801.3087817> (cit. on p. 1).
- [BGKM16] P. Berenbrink, G. Giakkoupis, A.-M. Kermarrec, and F. Mallmann-Trenn. “Bounds on the Voter Model in Dynamic Networks”. In: *ICALP*. 2016 (cit. on pp. 1–3, 5, 65).
- [BK89] A. Z. Broder and A. R. Karlin. “Bounds on the cover time”. In: *Journal of Theoretical Probability* 2.1 (1989), pp. 101–120 (cit. on pp. 5, 9, 66, 71).
- [CG97] F. R. Chung and F. C. Graham. *Spectral graph theory*. 92. American Mathematical Soc., 1997 (cit. on p. 67).
- [CLLM12] K.-M. Chung, H. Lam, Z. Liu, and M. Mitzenmacher. “Chernoff-Hoeffding Bounds for Markov Chains: Generalized and Simplified”. In: *STACS*. 2012, pp. 124–135 (cit. on pp. 8, 34).
- [CFR09a] C. Cooper, A. Frieze, and T. Radzik. “Multiple Random Walks in Random Regular Graphs”. In: *SIAM J. Discret. Math.* 23.4 (2009), pp. 1738–1761 (cit. on pp. 1, 3, 5).
- [CEOR13] C. Cooper, R. Elsässer, H. Ono, and T. Radzik. “Coalescing Random Walks and Voting on Connected Graphs”. In: *SIAM Journal on Discrete Mathematics* 27.4 (2013), pp. 1748–1758 (cit. on pp. 1–3, 5, 9, 60, 63, 64, 66, 71).
- [CER14] C. Cooper, R. Elsässer, and T. Radzik. “The Power of Two Choices in Distributed Voting”. In: *Proc. ICALP*. 2014, pp. 435–446 (cit. on p. 1).
- [CER+15] C. Cooper, R. Elsässer, T. Radzik, N. Rivera, and T. Shiraga. “Fast Consensus for Voting on General Expander Graphs”. In: *DISC*. 2015, pp. 248–262 (cit. on p. 1).
- [CF05] C. Cooper and A. M. Frieze. “The Cover Time of Random Regular Graphs”. In: *SIAM J. Discrete Math.* 18.4 (2005), pp. 728–740 (cit. on pp. 7, 60).

- [CFR09b] C. Cooper, A. Frieze, and T. Radzik. “Multiple Random Walks and Interacting Particle Systems”. In: *Automata, Languages and Programming*. Ed. by S. Albers, A. Marchetti-Spaccamela, Y. Matias, S. Nikolettseas, and W. Thomas. Berlin, Heidelberg: Springer Berlin Heidelberg, 2009, pp. 399–410 (cit. on p. 1).
- [CRRS16] C. Cooper, T. Radzik, N. Rivera, and T. Shiraga. “Fast plurality consensus in regular expanders”. In: *CoRR* abs/1605.08403 (2016). URL: <http://arxiv.org/abs/1605.08403> (cit. on p. 1).
- [CR16] C. Cooper and N. Rivera. “The Linear Voting Model”. In: *ICALP*. Vol. 55. 2016, 144:1–144:12 (cit. on p. 2).
- [Cox89] J. Cox. “Coalescing Random Walks and Voter Model Consensus Times on the Torus in \mathbb{Z}^d ”. In: *The Annals of Probability* 17.4 (1989), pp. 1333–1366 (cit. on pp. 2, 70).
- [CS10] A. Czumaj and C. Sohler. “Testing Expansion in Bounded-Degree Graphs”. In: *Combinatorics, Probability & Computing* 19.5-6 (2010), pp. 693–709 (cit. on p. 6).
- [DGM+11] B. Doerr, L. A. Goldberg, L. Minder, T. Sauerwald, and C. Scheideler. “Stabilizing Consensus With the Power of Two Choices”. In: *Proc. SPAA*. 2011, pp. 149–158 (cit. on p. 1).
- [ER09] K. Efremenko and O. Reingold. “How Well Do Random Walks Parallelize?” In: *RANDOM*. 2009, pp. 476–489 (cit. on p. 6).
- [EFK+16] R. Elsässer, T. Friedetzky, D. Kaaser, F. Mallmann-Trenn, and H. Trinker. “Efficient k-Party Voting with Two Choices”. In: *CoRR* abs/1602.04667 (2016). URL: <http://arxiv.org/abs/1602.04667> (cit. on p. 1).
- [ES11] R. Elsässer and T. Sauerwald. “Tight bounds for the cover time of multiple random walks”. In: *Theor. Comput. Sci.* 412.24 (2011), pp. 2623–2641 (cit. on p. 57).
- [GORN17] M. Gauci, M. E. Ortiz, M. Rubenstein, and R. Nagpal. “Error Cascades in Collective Behavior: A Case Study of the Gradient Algorithm on 1000 Physical Agents”. In: *Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems, AAMAS 2017, São Paulo, Brazil, May 8-12, 2017*. 2017, pp. 1404–1412. URL: <http://dl.acm.org/citation.cfm?id=3091319> (cit. on p. 1).
- [GL17] M. Ghaffari and J. Lengler. “Tight Analysis for the 3-Majority Consensus Dynamics”. In: *CoRR* abs/1705.05583 (2017). URL: <http://arxiv.org/abs/1705.05583> (cit. on p. 1).
- [GMS03] C. Gkantsidis, M. Mihail, and A. Saberi. “Conductance and congestion in power law graphs”. In: *SIGMETRICS*. 2003, pp. 148–159 (cit. on p. 71).
- [HP01] Y. Hassin and D. Peleg. “Distributed Probabilistic Polling and Applications to Proportionate Agreement”. In: *Information and Computation* 171.2 (2001), pp. 248–268 (cit. on pp. 1–3, 17).
- [LPW06] D. Levin, Y. Peres, and E. Wilmer. *Markov chains and mixing times*. American Mathematical Society, 2006 (cit. on pp. 24, 25, 29, 55, 56, 58, 64, 70).
- [Lez89] P. Lezaud. “Chernoff-type bound for finite Markov chains”. In: *Ann. Appl. Probab.* 8.3 (1989), pp. 849–867 (cit. on pp. 6, 8, 34).
- [Lov93] L. Lovász. “Random Walks on Graphs: A Survey”. In: *Combinatorics, Paul Erdős is Eighty* 2 (1993), pp. 1–46 (cit. on pp. 70, 74).

- [MSS15] A. W. Marcus, D. A. Spielman, and N. Srivastava. “Interlacing Families IV: Bipartite Ramanujan Graphs of All Sizes”. In: *FOCS*. 2015, pp. 1358–1377 (cit. on pp. 23, 56).
- [MPS06] M. Mihail, C. H. Papadimitriou, and A. Saberi. “On certain connectivity properties of the internet topology”. In: *J. Comput. Syst. Sci.* 72.2 (2006), pp. 239–251 (cit. on p. 72).
- [Moo73] J. Moon. “Random walks on random trees”. In: *J. Austral. Math. Soc.* 15 (1973) (cit. on p. 73).
- [Oli12] R. I. Oliveira. “On the coalescence time of reversible random walks.” In: *Trans. Am. Math. Soc.* 364.4 (2012), pp. 2109–2128 (cit. on pp. 2, 11, 51).
- [Oli13] R. I. Oliveira. “Mean field conditions for coalescing random walks”. In: *Ann. Probab.* 41.5 (2013), pp. 3420–3461 (cit. on p. 2).
- [Pel02] D. Peleg. “Local majorities, coalitions and monopolies in graphs: a review”. In: *Theoretical Computer Science* 282.2 (2002), pp. 231–257 (cit. on p. 1).
- [PS15] Y. Peres and P. Sousi. “Mixing Times are Hitting Times of Large Sets”. In: *Journal of Theoretical Probability* 28.2 (2015), pp. 488–519 (cit. on p. 28).
- [PVV09] E. Perron, D. Vasudevan, and M. Vojnović. “Using Three States for Binary Consensus on Complete Graphs”. In: *Proc. INFOCOM*. 2009, pp. 2527–2535 (cit. on p. 1).

A Basic Results about Markov Chains

We will frequently use the following basic fact about lazy random walks, which in fact also holds for arbitrary reversible Markov chains:

Lemma A.1 (cf. [LPW06, Chapter 12]). *Let P be the transition matrix of a reversible Markov chain with state space Ω . Then the following statements hold:*

(i) *If P is irreducible, then for any two states $x, y \in \Omega$,*

$$p_{x,y}^t \leq \pi(y) + \sqrt{\frac{\pi(y)}{\pi(x)}} \cdot \lambda^t,$$

where $\lambda := \max\{\lambda_2, |\lambda_n|\}$ and $\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_n$ are the n real eigenvalues of the matrix P .

(ii) *If the Markov chain is a non-lazy random walk on a bipartite regular graph with two partitions V_1 and V_2 , then for any pair of states x, y in the same partition*

$$p_{x,y}^t \leq \frac{2}{n} \cdot (1 + (-1)^{t-1}) + 2(\max\{\lambda_2, |\lambda_{n-1}|\})^t.$$

Similarly, if x and y are in opposite partitions,

$$p_{x,y}^t \leq \frac{2}{n} \cdot (1 + (-1)^t) + 2(\max\{\lambda_2, |\lambda_{n-1}|\})^t.$$

(iii) *If the Markov chain is lazy, then for any state $x \in \Omega$, $p_{x,x}^t$ is non-increasing in t . In particular, $p_{x,x}^t \geq \pi(x)$.*

Proof. The first statement can be found in [LPW06, Equation 12.11].

For the second statement, recall the spectral representation [LPW06, Lemma 12.2 (iii)]

$$p_{x,y}^t = \pi(y) + \pi(y) \cdot \sum_{k=2}^n u_k(x) \cdot u_k(y) \cdot \lambda_k^t, \quad (38)$$

where u_k is the corresponding eigenvector to λ_k . Since all eigenvalues are non-negative, we conclude from (38) that $p_{x,x}^t$ is non-increasing in t as needed. Since G is bipartite and regular, it is not difficult to verify that $\lambda_n = -1$ and $u_n(x) = \sqrt{1/n}$ if $x \in V_1$ and $u_n(x) = -\sqrt{1/n}$ if $x \in V_2$ is the corresponding eigenvector. Hence,

$$\begin{aligned} \left| p_{x,y}^t - \frac{2}{n} \cdot (1 + (-1)^{t-1}) \right| &\leq \pi(y) \cdot \left| \sum_{k=2}^{n-1} u_k(x) \cdot u_k(y) \cdot \lambda_k^t \right| \\ &\leq \frac{2}{n} \cdot \max_{2 \leq k \leq n-1} |\lambda_k^t| \cdot \frac{1}{n} \cdot \sum_{k=2}^{n-1} |u_k(x) \cdot u_k(y)| \\ &\leq \frac{2}{n} \cdot \max_{2 \leq k \leq n-1} |\lambda_k^t| \cdot \sqrt{\sum_{k=2}^{n-1} u_k(x)^2 \cdot \sum_{k=2}^{n-1} u_k(y)^2} \end{aligned}$$

As in [LPW06, Proof of Theorem 12.3], using the orthonormality of the eigenvectors, we have

$$\sum_{k=2}^{n-1} u_k(x)^2 \leq \sum_{k=2}^n u_k(x)^2 \leq n,$$

and the second statement follows if u and v are in the same partition. The case where u and v are in different partitions follows analogously.

For the third statement, first note that by [LPW06, Exercise 12.3], all eigenvalues of the transition matrix P are non-negative. Since all eigenvalues are non-negative, we conclude from (38) that $p_{x,x}^t$ is non-increasing in t as needed. Due to this and the fact that $p_{x,x}^t$ converges to π_x , we get that $p_{x,x}^t \geq \pi_x$. \square

The following is a simple corollary from a recent work by Marcus et al. [MSS15] on the existence of Ramanujan graphs.

Lemma A.2 (cf. Marcus et al. [MSS15]). *For any integer $d \geq 3$, there are d -regular bipartite Ramanujan graph $H = (V, E)$ with $t_{\text{mix}} = O(\log n / \log d)$.*

Proof. Marcus et al. [MSS15] show that the existence of a d -regular bipartite Ramanujan graph H such that $\max\{\lambda_2(\widehat{Q}), |\lambda_{n-1}(\widehat{Q})|\} = O(1/\sqrt{d})$, where $\widehat{Q} = \frac{1}{d}A$ is the transition matrix of a non-lazy random walk where A is the adjacency matrix. By the second statement of Lemma A.1, for any pair of states x, y in the same partition

$$\hat{q}_{x,y}^t \leq \frac{2}{n} \cdot (1 + (-1)^{t-1}) + 2 (\max\{\lambda_2, |\lambda_{n-1}|\})^t.$$

Similarly, x and y are in opposite partitions,

$$\hat{q}_{x,y}^t \leq \frac{2}{n} \cdot (1 + (-1)^t) + 2 (\max\{\lambda_2, |\lambda_{n-1}|\})^t.$$

Furthermore note that $q_{x,y}^t \geq 2/n$ due to [Lemma A.1](#).(iii) for even (or odd) t depending on whether x and y are in the same partitions.

Fix $t = O(\log n / \log d)$ such that $2 \left(\max\{\lambda_2(\widehat{Q}), |\lambda_{n-1}(\widehat{Q})|\} \right)^t \leq \frac{1}{20n}$, where we note that such a t exists due to $\max\{\lambda_2(\widehat{Q}), |\lambda_{n-1}(\widehat{Q})|\} = O(1/\sqrt{d})$. We choose s to be the smallest odd integer being greater than $20t$. To translate from the non-lazy random walk \widehat{Q} to a lazy-random walk P , let Z denote the number of non-loops performed by a lazy random walk of length s . Since, the probability for a self-loop is $1/2$ and the number of self-loops is binomially distributed, we have

$$\mathbb{P}[Z \geq t] \geq 19/20.$$

By symmetry and the fact that s is odd, $\mathbb{P}[Z \text{ is even}] = \frac{1}{2}$. Hence, by the Union bound,

$$\mathbb{P}[Z \text{ is even} \mid Z \geq t] \geq \mathbb{P}[Z \text{ is even} \cap Z \geq t] \geq \mathbb{P}[Z \text{ is even}] - \mathbb{P}[Z < t] \geq \frac{9}{20},$$

and similarly, $\mathbb{P}[Z \text{ is odd} \mid Z \geq t] \geq \frac{9}{20}$. Let V_1 and V_2 be the bipartite partition of V .

$$\begin{aligned} \|p_{u,\cdot}^s - \pi\|_{\text{TV}} &\leq \mathbb{P}[Z < t] \cdot 1 + \mathbb{P}[Z \geq t] \cdot \left(\sum_{v \in V_1} \left| \frac{11}{20} \hat{q}_{u,v}^t - \frac{1}{n} \right| + \sum_{v \in V_2} \left| \frac{11}{20} \hat{q}_{u,v}^t - \frac{1}{n} \right| \right) \\ &\leq \mathbb{P}[Z < t] \cdot 1 + \mathbb{P}[Z \geq t] \cdot \left(\sum_{v \in V} \left| \frac{11}{20} \left(\frac{2}{n} + \frac{1}{20n} \right) - \frac{1}{n} \right| \right) \\ &\leq \mathbb{P}[Z < t] \cdot 1 + \mathbb{P}[Z \geq t] \cdot \left(\sum_{v \in V} \left| \frac{22}{20n} - \frac{1}{n} \right| + \frac{11}{400} \right) \\ &\leq \frac{1}{20} \cdot 1 + \frac{19}{20} \left(\frac{2}{20} + \frac{11}{400} \right) < 1/e, \end{aligned}$$

where the first inequality follows from the equations for $p_{x,y}^t$ above. \square

Corollary A.3. *Let n_0 be a sufficiently large constant. Let H_n be the graph of [Lemma A.2](#) with n nodes and $d = \lceil \sqrt{n} \rceil$ for $n \geq n_0$. There exists a universal constant C such that $\max_{n \geq n_0} \{t_{\text{sep}}(H_n)\} \leq C$.*

The corollary follows directly from [Lemma A.2](#) and $t_{\text{sep}} \leq 4t_{\text{mix}}$.

The following lemma will be helpful to define a coupling between distributions that are close to the stationary distribution and the exact stationary distribution. (A very similar lemma has been derived in [\[ES11, Lemma 2.8\]](#))

Lemma A.4. *Let $\varepsilon \in (0, 1]$ be an arbitrary value. Let Z_1 and Z_2 be two probability distributions over $\{1, \dots, n\}$ so that $\mathbb{P}[Z_1 = i] \geq \varepsilon \cdot \mathbb{P}[Z_2 = i]$ for every $1 \leq i \leq n$. Then, there is a coupling $(\tilde{Z}_1, \tilde{Z}_2)$ of (Z_1, Z_2) and an event \mathcal{E} with $\mathbb{P}[\mathcal{E}] \geq \varepsilon$ so that*

$$\mathbb{P}[\tilde{Z}_1 = i \mid \mathcal{E}] = \mathbb{P}[\tilde{Z}_2 = i] \quad \text{for every } 1 \leq i \leq n.$$

Proof. Let $U \in [0, 1]$ be a uniform random variable. We next define our coupling $(\tilde{Z}_1, \tilde{Z}_2)$ of Z_1 and Z_2 that will depend on the outcome of U . First, if $U \in [0, \varepsilon)$, then we set

$$\tilde{Z}_1 = \tilde{Z}_2 = i, \quad \text{if } i \text{ satisfies } \varepsilon \sum_{k=1}^{i-1} \mathbb{P}[Z_2 = k] \leq U < \varepsilon \sum_{k=1}^i \mathbb{P}[Z_2 = k].$$

For the case where $U \in (\varepsilon, 1)$, it is clear that the definition of U can be extended in a way so that \tilde{Z}_1 has the same distribution as Z_1 , and \tilde{Z}_2 has the same distribution as Z_2 . Furthermore, notice that if $U \in [0, \varepsilon]$ happens, then \tilde{Z}_1 has the same distribution as Z_2 , and $\tilde{Z}_1 = \tilde{Z}_2$. Observing that $\mathbb{P}[U \in [0, \varepsilon]] = \varepsilon$ completes the proof. \square

The following lemma is an immediate consequence of [Lemma A.4](#).

Lemma A.5. *Consider a random walk $(X_t)_{t \geq 0}$, starting from an arbitrary but fixed vertex x_0 . Then with probability at least $1 - 1/e$, we can couple $X_{4t_{\text{mix}}}$ with the stationary distribution.*

Proof. Consider the random walk $(X_t)_{t \geq 0}$ after step $s := t_{\text{sep}} \leq 4t_{\text{mix}}$. By definition of t_{sep} , $p_{u,v}^s \geq (1 - 1/e)\pi(v)$. Applying [Lemma A.4](#), where Z_1 is the distribution given by $p_{u,v}^t$ and Z_2 is the stationary distribution shows that with probability at least $1 - 1/e$, X_s has the same distribution as π . If this is the case, then the same holds for $X_{4t_{\text{mix}}}$ as well. \square

The lemma above shows that for t_{meet} and t_{coal} it suffices to consider the stationary case:

Lemma A.6. *For any graph G ,*

$$\max\{(1/e)t_{\text{mix}}, t_{\text{meet}}^\pi\} \leq t_{\text{meet}} \leq \frac{2}{(1 - 1/e)^2} \cdot (4t_{\text{mix}} + 2t_{\text{meet}}^\pi),$$

and similarly, $t_{\text{coal}} \leq 4 \cdot (4t_{\text{mix}} + 2t_{\text{coal}}^\pi)$.

Proof. We begin by proving the lower bound on t_{meet} . First, consider two independent random walks $(X_t)_{t \geq 0}$ and $(Y_t)_{t \geq 0}$ that are run for $t = e \cdot t_{\text{meet}}$ time-steps. Then, we have

$$\bar{d}(t) = \max_{u,v} \|p_{u,\cdot}^t - p_{v,\cdot}^t\|_{\text{TV}} \leq \mathbb{P}[\cup_{s=0}^t X_s = Y_s] \leq \frac{1}{e},$$

where the first inequality is due to the coupling method [[LPW06](#), Theorem 5.3] and the second inequality follows by Markov's inequality. The above inequality implies $t_{\text{mix}} \leq e \cdot t_{\text{meet}}$. Furthermore, $t_{\text{meet}}^\pi \leq t_{\text{meet}}$ holds by definition, and the lower bound follows.

For the upper bound, we divide the two random walks into consecutive epochs of length $\ell := 4t_{\text{mix}} + 2t_{\text{meet}}^\pi$. For the statement it suffices to prove that in each such epoch, regardless of the start vertices of the two random walks, a meeting occurs with probability at least $(1 - 1/e)^2 \cdot 1/2$.

Consider the first random walk $(X_t)_{t \geq 0}$ starting from an arbitrary vertex after $s := 4t_{\text{mix}}$ steps. By [Lemma A.5](#), we obtain that with probability at least $1 - 1/e$, the distribution of X_s is equal to that of a stationary random walk. Similarly, we obtain that with probability at least $1 - 1/e$, the distribution of Y_s is equal to that of a stationary distribution. Hence with probability $(1 - 1/e)^2$, X_s and Y_s are drawn independently from the stationary distribution. In this case, it follows by Markov's inequality that the two random walks meet before step $s + 2t_{\text{meet}}^\pi$ with probability at least $1/2$. Overall, we have shown that with probability at least $(1 - 1/e)^2 \cdot 1/2$, a meeting occurs in a single epoch. Since this lower bound holds for every epoch, independent of the outcomes in previous epochs, the upper bound on the expected time t_{meet} follows. The upper bound on t_{coal} in terms of t_{coal}^π is shown in exactly the same way. \square

Lemma A.7. *For a lazy random walk on an n -vertex graph with $n \geq 2$, we have $t_{\text{hit}} \geq \frac{2}{\pi_{\text{min}}} - 2$. In particular for $n \geq 2$, we have $t_{\text{hit}} \geq 1/\pi_{\text{min}} \geq n$.*

Note that $t_{\text{hit}} \geq \frac{2}{\pi_{\text{min}}} - 2$ is tight in the sense that the hitting time of the clique is indeed $2(n - 1) = \frac{2}{\pi_{\text{min}}} - 2$ since the random walk moves w.p. $1/2$ and when it moves the probability to hit the target node is $1/(n - 1)$ (assuming that the random walk is not on the target node).

Proof. Let u be a vertex attaining $\pi_{\min} = \pi(u)$. Consider the random walks $(X_t)_{t \geq 0}$ starting at u . Then it is well-known (cf. [AF02]) that for the first return $\tau^+(u, u) := \min\{t > 0 : X_t = u, X_0 = u\}$, we have $\mathbb{E}[\tau^+(u, u)] = 1/\pi(u) = 1/\pi_{\min}$. By conditioning on the first step of the random walk, we obtain

$$\frac{1}{\pi_{\min}} = \mathbb{E}[\tau^+(u, u)] = 1 + \frac{1}{2} \cdot 0 + \frac{1}{2} \sum_{v \in N(u)} \frac{1}{\deg(u)} \cdot t_{\text{hit}}(v, u),$$

and rearranging yields

$$\frac{1}{\deg(u)} \cdot \sum_{v \in N(u)} t_{\text{hit}}(v, u) = \frac{2}{\pi_{\min}} - 2.$$

Now by the pigeonhole principle there exists a vertex $v \in N(u)$ with $t_{\text{hit}}(v, u) \geq \frac{2}{\pi_{\min}} - 2$, and the first claim follows. The second part follows from observing that if $n \geq 2$ we have $\pi_{\min} \leq 1/2$ and thus $t_{\text{hit}} \geq \frac{2}{\pi_{\min}} - 2 \geq \frac{1}{\pi_{\min}} \geq n$, where the last inequality follows from the simple pigeon hole principle. \square

Observation A.8. Consider two random walks $(X_t)_{t \geq 0}$ and $(Y_t)_{t \geq 0}$ starting on nodes drawn from the stationary distribution. Fix an arbitrary $t \in \mathbb{N}$. Define the collision-counting random variables $Z_1 = \sum_{i=0}^{\lceil t/2 \rceil} \mathbf{1}_{X_t=Y_t}$, $Z_2 = \sum_{i=\lceil t/2 \rceil+1}^t \mathbf{1}_{X_t=Y_t}$, and $Z = Z_1 + Z_2$. Then $\mathbb{P}[Z_1 \geq 1 \mid Z \geq 1] \geq \frac{1}{2}$.

Proof. Since both nodes start from the stationary distribution, $\mathbb{P}[Z_1 \geq 1] \geq \mathbb{P}[Z_2 \geq 1]$. By the Union bound, $\mathbb{P}[Z \geq 1] \leq \mathbb{P}[Z_1 \geq 1] + \mathbb{P}[Z_2 \geq 1] \leq 2 \cdot \mathbb{P}[Z_1 \geq 1]$. By law of total probability, $\mathbb{P}[Z_1] = \mathbb{P}[Z_1 \geq 1 \mid Z \geq 1] \cdot \mathbb{P}[Z \geq 1]$. Putting everything together yields $\mathbb{P}[Z_1 \geq 1 \mid Z \geq 1] \geq \frac{1}{2}$. \square

Lemma A.9. Let $(X_t)_{t \geq 0}$ be a stochastic process satisfying (i) $\mathbb{E}[X_t \mid \mathcal{F}_{t-1}] \leq \beta \cdot X_{t-1}$, for some $\beta < 1$, and (ii) $X_t \geq 0$ for all $t \geq 0$. Let $\tau(g) := \min\{t \geq 0 \mid X_t \leq g\}$ for $g \in (0, |X_0|)$, then

$$\mathbb{E}[\tau(g)] \leq 2 \cdot \lceil \log_{\beta}(g/(2X_0)) \rceil.$$

Proof. By the iterative law of expectation, we have

$$\mathbb{E}[X_t] \leq \beta^t \cdot X_0.$$

Furthermore, by Markov's inequality, for any $\lambda \geq 1$

$$\mathbb{P}[\tau(g) > \lambda \cdot \lceil \log_{\beta}(g/(2X_0)) \rceil] \leq \mathbb{P}\left[X_{\lambda \cdot \lceil \log_{\beta}(g/(2X_0)) \rceil} > g\right] \leq \frac{\mathbb{E}\left[X_{\lambda \cdot \lceil \log_{\beta}(g/(2X_0)) \rceil}\right]}{g} \leq 2^{-\lambda}.$$

Therefore,

$$\begin{aligned} \mathbb{E}[\tau(g)] &= \sum_{i=1}^{\infty} \mathbb{P}[\tau(g) \geq i] \\ &\leq \lceil \log_{\beta}(g/(2X_0)) \rceil + \sum_{\lambda=1}^{\infty} \lceil \log_{\beta}(g/(2X_0)) \rceil \cdot \mathbb{P}[\tau(g) > \lambda \cdot \lceil \log_{\beta}(g/(2X_0)) \rceil] \\ &\leq 2 \cdot \lceil \log_{\beta}(g/(2X_0)) \rceil. \end{aligned}$$

\square

B Bounding t_{meet} and Implications for t_{coal} , t_{hit} and t_{cov}

Although the focus of this work is on understanding the coalescence time, in order to apply our general results, we need to devise some tools to obtain lower and upper bounds on t_{meet} . In [Theorem B.1 \(Section B.1\)](#) we establish upper and lower bounds on the meeting time in terms of $\|\pi\|_2^2 = \sum_{u \in V} \pi(u)^2$. [Section B.1](#) contains several additional upper bounds on t_{meet} and t_{hit} . Through combination with other results, we also obtain new bounds on t_{coal} and t_{cov} . A common feature of many of these bounds is a sub-linear dependence on the spectral gap $1/(1 - \lambda_2)$, which we obtain by an application of short-term bounds on the t -step transition probabilities.

In [Proposition B.9 \(Section B.2\)](#) we establish a discrete-time counterpart of [[AF02](#), Proposition 14.5], albeit with worse constants, stating that the meeting time is at most of the order of the hitting time; on vertex transitive graphs these quantities are asymptotically of the same order.

B.1 Relating Meeting Time to t_{mix} and $\frac{1}{1-\lambda_2}$

We first state some basic bounds on t_{hit} and t_{meet} , which mostly follow directly from [\(1\)](#) and its counterpart for the hitting times (cf. Cooper, Frieze [[CF05](#)]). In these bounds, we will use the following notation:

$$C_{\max} := \max_{u \in V} \sum_{t=0}^{t_{\text{mix}}-1} \sum_{v \in V} (p_{u,v}^t)^2,$$

$$C_{\min} := \min_{u \in V} \sum_{t=0}^{t_{\text{mix}}-1} \sum_{v \in V} (p_{u,v}^t)^2.$$

Note that C_{\max} and C_{\min} provide worst-case upper respective lower bounds on the expected collisions of two independent random walks of length t_{mix} , starting from the same vertex u . Similarly, we define

$$R_{\max} := \max_{u \in V} \sum_{t=0}^{t_{\text{mix}}-1} p_{u,u}^t.$$

Note that R_{\max} is the number of expected returns of a random walk to u during t_{mix} steps. This quantity is more convenient to bound than C_{\max} , for instance, it can be easily bounded by $\max_{u \in V} \pi(u) \cdot t_{\text{mix}} + \frac{1}{1-\lambda_2}$ (cf. [Lemma A.1](#), or also [[CEOR13](#)]).

Theorem B.1. *For any graph $G = (V, E)$, the following statements hold:*

(i) *For any pair of vertices $u, v \in V$,*

$$t_{\text{hit}}(u, v) \leq \frac{5e \cdot (\sum_{t=0}^{t_{\text{mix}}-1} p_{v,v}^t)}{\pi(v)}.$$

In particular, if the graph G is Γ -approximative regular, then $t_{\text{hit}}(u, v) \leq 5e \cdot \Gamma \cdot n \cdot \sum_{t=0}^{t_{\text{mix}}-1} p_{v,v}^t$.

(ii) *For any pair of vertices $u, v \in V$,*

$$t_{\text{meet}}(u, v) \leq \frac{5e^2 \cdot C_{\max}}{\|\pi\|_2^2}.$$

In particular, if the graph G is Γ -approximative regular, then

$$t_{\text{meet}}(u, v) \leq \frac{10e^2 \cdot (4 + \log_2(\Gamma)) \cdot R_{\max}}{\|\pi\|_2^2}.$$

(iii) It holds that,

$$t_{\text{meet}}^\pi \geq \frac{C_{\min}}{64\|\pi\|_2^2}.$$

In particular, if the graph G is Γ -approximative regular, then $t_{\text{meet}}^\pi = \Omega(n/\Gamma)$.

Since $C_{\min} \geq 1$, the last statement of the lemma implies also $t_{\text{meet}}^\pi = \Omega(\frac{1}{\|\pi\|_2^2})$. We remark that the second upper bound on $t_{\text{meet}}(u, v)$ depends only logarithmically on Γ .

Proof. We begin by proving the first part. Consider one random walk $(X_t)_{t \geq 0}$, starting from an arbitrary vertex. Divide the time-interval into consecutive epochs of length $t_{\text{sep}} + t_{\text{mix}}$, and let

$$Z := \sum_{t=t_{\text{sep}}}^{t_{\text{sep}}+t_{\text{mix}}-1} \mathbf{1}_{X_t=v}$$

denote the number of visits. Then, by the separation time, $\mathbb{E}[Z] \geq t_{\text{mix}} \cdot \frac{1}{e} \cdot \pi(v)$, and (1) yields

$$\mathbb{P}[Z \geq 1] \geq \frac{t_{\text{mix}} \cdot \frac{\pi(v)}{e}}{\mathbb{E}[Z \mid Z \geq 1]}.$$

Clearly, $\mathbb{E}[Z \mid Z \geq 1] \leq \max_{t=0}^{t_{\text{mix}}-1} \sum_{s=t}^{t_{\text{mix}}} p_{v,v}^{s-t} \leq \sum_{t=0}^{t_{\text{mix}}-1} p_{v,v}^t$. Hence,

$$\mathbb{P}[Z \geq 1] \geq \frac{t_{\text{mix}} \cdot \frac{\pi(v)}{e}}{\sum_{t=0}^{t_{\text{mix}}-1} p_{v,v}^t} =: p.$$

This means that in every epoch of length $t_{\text{sep}} + t_{\text{mix}} \leq 5t_{\text{mix}}$, the random walk has a probability of at least p to visit vertex v , and this is independent of any previous epoch. Therefore, the expected number of steps until v is visited is upper bounded by

$$t_{\text{hit}}(u, v) \leq 5t_{\text{mix}} \cdot \frac{1}{p} \leq \frac{5e \cdot \sum_{t=0}^{t_{\text{mix}}-1} p_{v,v}^t}{\pi(v)}.$$

The claim for Γ -approximative regular graphs follows from the observation that $\min_{u \in V} \pi(u) \geq 1/(\Gamma n)$. We continue with the second part. Consider two independent random walks, $(X_t)_{t \geq 0}$, $(Y_t)_{t \geq 0}$ of length $t_{\text{sep}} + t_{\text{mix}}$ with arbitrary start vertices. Let Z be the random variable counting the number of collisions between steps t_{sep} and $t_{\text{sep}} + t_{\text{mix}} - 1$, *i.e.*,

$$Z := \sum_{t=t_{\text{sep}}}^{t_{\text{sep}}+t_{\text{mix}}-1} \mathbf{1}_{X_t=Y_t}.$$

By linearity of expectation,

$$\mathbb{E}[Z] = \sum_{t=t_{\text{sep}}}^{t_{\text{sep}}+t_{\text{mix}}-1} \sum_{u \in V} \mathbb{P}[X_t = u] \cdot \mathbb{P}[Y_t = u] \geq t_{\text{mix}} \cdot \frac{1}{e^2} \cdot \|\pi\|_2^2. \quad (39)$$

Let us now consider $\mathbb{E}[Z \mid Z \geq 1]$ and recall that conditioning on $Z \geq 1$ can be regarded as jumping to the first step $\tau := \min\{t: t_{\text{sep}} \leq t \leq t_{\text{mix}} - 1, X_\tau = Y_\tau\}$ without knowing anything about the future steps $t > \tau$ of both walks. Therefore,

$$\mathbb{E}[Z \mid Z \geq 1] \leq \max_{u \in V} \sum_{t=0}^{t_{\text{mix}}-1} \sum_{v \in V} (p_{u,v}^t)^2 = C_{\max}.$$

Plugging this into (1) and using (39) we finally arrive at

$$\mathbb{P}[Z \geq 1] = \frac{\mathbb{E}[Z]}{\mathbb{E}[Z \mid Z \geq 1]} \geq \frac{t_{\text{mix}} \cdot \frac{1}{e^2} \cdot \|\pi\|_2^2}{C_{\text{max}}} =: p.$$

Hence,

$$t_{\text{meet}}(u, v) \leq 5t_{\text{mix}} \cdot \frac{1}{p} \leq \frac{5e^2 \cdot C_{\text{max}}}{\|\pi\|_2^2}.$$

Let us derive the result for Γ -approximate regular graphs. To this end, define

$$S_i := \{u \in V : \deg(u) \in [2^{i-1}, 2^i]\},$$

and note that $S_0, \dots, S_{\log_2 n}$ forms a partition of V . Since the graph is Γ -approximate regular, at most $4 + \log_2(\Gamma)$ of the S_i 's are non-empty. Hence there exists a set S_j with

$$\sum_{j \in S_i} \pi(j)^2 \geq \frac{1}{4 + \log_2(\Gamma)} \cdot \|\pi\|_2^2.$$

Let us now by Z_j denote the collisions on the set S_j , *i.e.*,

$$Z_j := \sum_{t=t_{\text{sep}}}^{t_{\text{sep}}+t_{\text{mix}}-1} \mathbf{1}_{\{X_t=Y_t\} \cap \{X_t \in S_j\}}.$$

Then,

$$\mathbb{E}[Z_j] = \sum_{t=t_{\text{sep}}}^{t_{\text{sep}}+t_{\text{mix}}} \sum_{u \in S_j} \mathbb{P}[X_t = u] \cdot \mathbb{P}[Y_t = u] \geq t_{\text{mix}} \cdot \sum_{j \in S_i} \frac{1}{e^2} \cdot \pi(j)^2 \geq \frac{t_{\text{mix}}}{e^2} \cdot \frac{1}{4 + \log_2(\Gamma)} \cdot \|\pi\|_2^2.$$

Furthermore,

$$\mathbb{E}[Z_j \mid Z_j \geq 1] \leq \max_{u \in S_i} \sum_{t=0}^{t_{\text{mix}}} \sum_{v \in S_i} (p_{u,v}^t)^2 \leq \max_{u \in S_i} \sum_{t=0}^{t_{\text{mix}}} \sum_{v \in S_i} p_{u,v}^t \cdot 2p_{v,u}^t,$$

having used reversibility, *i.e.*, $p_{u,v}^t \pi(u) = p_{v,u}^t \pi(v)$ and $\pi(v)/\pi(u) \leq 2$ by definition of S_i . Further,

$$\mathbb{E}[Z_j \mid Z_j \geq 1] \leq 2 \cdot \max_{u \in S_i} \sum_{t=0}^{t_{\text{mix}}-1} \sum_{v \in V} p_{u,v}^t \cdot p_{v,u}^t \leq 2 \cdot \max_{u \in V} \sum_{t=0}^{t_{\text{mix}}-1} p_{u,u}^{2t} \leq 2 \cdot \max_{u \in V} \sum_{t=0}^{t_{\text{mix}}-1} p_{u,u}^t = 2 \cdot R_{\text{max}},$$

where the last inequality holds since $p_{u,u}^t$ is non-increasing by [Lemma A.1](#). Hence, similarly as before,

$$t_{\text{meet}}(u, v) \leq 5t_{\text{mix}} \cdot \frac{2 \cdot R_{\text{max}}}{\frac{t_{\text{mix}}}{e^2} \cdot \frac{1}{4 + \log_2(\Gamma)} \cdot \|\pi\|_2^2} \leq \frac{10e^2 \cdot (4 + \log_2(\Gamma)) \cdot R_{\text{max}}}{\|\pi\|_2^2}.$$

Finally, for the third statement, let $(X_t)_{t \geq 0}, (Y_t)_{t \geq 0}$ be two random walk starting from stationarity. Let \tilde{Z} be the random variable counting the number of collisions between steps 0 and $2t_{\text{mix}}$, *i.e.*,

$$\tilde{Z} := \sum_{t=0}^{2t_{\text{mix}}-1} \mathbf{1}_{X_t=Y_t}.$$

Then,

$$\mathbb{E} \left[\tilde{Z} \right] = 2t_{\text{mix}} \cdot \|\pi\|_2^2.$$

In order to lower bound $\mathbb{E} \left[\tilde{Z} \mid \tilde{Z} \geq 1 \right]$, let us write $\tilde{Z} = \tilde{Z}_1 + \tilde{Z}_2$ with $\tilde{Z}_1 := \sum_{t=0}^{t_{\text{mix}}} \mathbf{1}_{X_t=Y_t}$ and $\tilde{Z}_2 := \sum_{t=t_{\text{mix}}+1}^{2t_{\text{mix}}-1} \mathbf{1}_{X_t=Y_t}$. By [Observation A.8](#), $\mathbb{P} \left[\tilde{Z}_1 \geq 1 \mid \tilde{Z} \geq 1 \right] \geq \frac{1}{2}$. Therefore, by law of total expectation,

$$\mathbb{E} \left[\tilde{Z} \mid \tilde{Z} \geq 1 \right] \geq \frac{1}{2} \cdot \mathbb{E} \left[\tilde{Z} \mid \tilde{Z}_1 \geq 1 \right] \geq \frac{1}{2} \cdot \min_{u \in V} \sum_{t=0}^{t_{\text{mix}}} \sum_{v \in V} (p_{u,v}^t)^2 = \frac{1}{2} \cdot C_{\text{min}}.$$

Hence,

$$\mathbb{P} \left[\tilde{Z} \geq 1 \right] \leq \frac{2t_{\text{mix}} \cdot \|\pi\|_2^2}{\frac{1}{2} \cdot C_{\text{min}}} =: p.$$

Consider now $\lceil \frac{C_{\text{min}}}{16\|\pi\|_2^2 t_{\text{mix}}} \rceil$ consecutive time-intervals of length $2t_{\text{mix}}$ each. Note that if $2t_{\text{mix}} \geq \frac{1}{2} \cdot \frac{C_{\text{min}}}{16\|\pi\|_2^2}$, then we have

$$t_{\text{meet}}^\pi \geq t_{\text{mix}} \geq \frac{C_{\text{min}}}{64\|\pi\|_2^2},$$

and the claim follows immediately. Hence we may assume for the remainder of the proof that $2t_{\text{mix}} < \frac{1}{2} \cdot \frac{C_{\text{min}}}{16\|\pi\|_2^2}$ and we conclude that, if B denotes the total number of collisions between the walks across all the intervals,

$$\mathbb{E} [B] \leq \left\lceil \frac{C_{\text{min}}}{16\|\pi\|_2^2 \cdot t_{\text{mix}}} \right\rceil \cdot p \leq \frac{C_{\text{min}}}{8\|\pi\|_2^2 \cdot t_{\text{mix}}} \cdot \frac{2t_{\text{mix}} \cdot \|\pi\|_2^2}{\frac{1}{2} \cdot C_{\text{min}}} \leq \frac{1}{2}.$$

Hence by Markov's inequality, $\mathbb{P} [B \geq 1] \leq \frac{1}{2}$ and thus $t_{\text{meet}}^\pi \geq \frac{1}{2} \cdot \frac{C_{\text{min}}}{8\|\pi\|_2^2}$ in this case. The claim for Γ -approximate regular graph follows immediately since $\|\pi\|_2^2 \leq \max_{u \in V} \pi(u) \leq \Gamma/n$. Together with $C_{\text{min}} \geq 1$ this completes the proof of the theorem. \square

It is interesting to compare the upper bound on t_{meet} in [Theorem B.1](#) with the bound $t_{\text{meet}} = O(\frac{1}{1-\lambda_2} \cdot (\frac{1}{\|\pi\|_2^2} + \log n))$ from Cooper et al. [[CEOR13](#), Theorem 2]. Using the trivial bound $C_{\text{max}} \leq t_{\text{mix}}$ and $t_{\text{mix}} = O(\frac{\log n}{1-\lambda_2})$, we obtain $t_{\text{meet}} = O(\frac{1}{1-\lambda_2} \cdot \frac{\log n}{\|\pi\|_2^2})$, which is at most a $\log n$ -factor worse. However, for certain graphs like grids or tori one may have a better control on the t -step probabilities, so that $C_{\text{max}} \ll t_{\text{mix}}$ could be established.

Proposition B.2. *Combining the upper bound on t_{meet} in [Theorem B.1](#) with the bound $t_{\text{meet}} = O(\frac{1}{1-\lambda_2} \cdot (\frac{1}{\|\pi\|_2^2} + \log n))$ from Cooper et al. [[CEOR13](#), Theorem 2] together with [Theorem 1.1](#) we derive*

$$t_{\text{coal}} = O \left(\frac{1}{1-\lambda_2} \cdot \frac{1}{\|\pi\|_2^2} + t_{\text{mix}} \cdot \log^3 n \right),$$

which is at least as good as the bound of [[CEOR13](#), Theorem 1] and equally good if one uses the trivial bound $t_{\text{mix}} = O(\frac{\log n}{1-\lambda_2})$.

Proof. First assume $t_{\text{meet}}/t_{\text{mix}} \geq \log^2 n$. In this case, by [Theorem 1.1](#),

$$t_{\text{coal}} = O(t_{\text{meet}}) = O\left(\frac{1}{1-\lambda_2} \cdot \left(\frac{1}{\|\pi\|_2^2} + \log n\right)\right)$$

follows immediately. Next assume $t_{\text{meet}}/t_{\text{mix}} \leq \log^2 n$, so $t_{\text{meet}} \leq t_{\text{mix}} \cdot \log^2 n$. By [Proposition 3.4](#), $t_{\text{coal}} = O(t_{\text{meet}} \log n) = O(t_{\text{mix}} \log^3 n)$. Using $t_{\text{mix}} = O\left(\frac{\log n}{1-\lambda_2}\right)$ we derive indeed $O\left(\frac{1}{1-\lambda_2} \cdot \left(\frac{1}{\|\pi\|_2^2} + \log^4 n\right)\right)$ the same bound as [\[CEOR13, Theorem 1\]](#). \square

In the following, we will try to get more concrete estimates than the ones in [Theorem B.1](#) by expressing the number of expected returns or C_{max} through $1-\lambda_2$ and t_{mix} . To this end, we define

$$\beta := \min\left\{\frac{\log(1/(1-\lambda_2))}{1-\lambda_2}, t_{\text{mix}}\right\}.$$

Note that since $t_{\text{mix}} = \Omega(1/(1-\lambda_2))$ (e.g., [\[AF02\]](#)), we have $\beta \geq \Omega(1/(1-\lambda_2))$. Further, $\beta \leq t_{\text{mix}} = O(\log n/(1-\lambda_2))$. Hence β is always sandwiched between the relaxation time $1/(1-\lambda_2)$ and mixing time.

We will frequently make use of the following result, which is a straightforward generalization of a result in the textbook by Aldous and Fill [\[AF02\]](#) from regular to Γ -approximate regular graph.

Lemma B.3 ([\[AF02, Proposition 6.16 \(iii\)\]](#)). *Let G be any Γ -approximate regular graph. Then for any $\tau \leq 5n^2$,*

$$\sum_{t=0}^{\tau-1} p_{u,u}^t \leq 2\Gamma \cdot \sqrt{5\tau}.$$

Since $p_{u,u}^t$ is non-increasing, this implies for any $1 \leq t \leq 5n^2$, $p_{u,u}^t \leq \frac{20\Gamma}{\sqrt{t}}$.

Theorem B.4. *For any regular graph we have*

$$t_{\text{hit}} = O\left(\frac{n}{\sqrt{1-\lambda_2}}\right),$$

and by Cheeger's inequality we obtain $t_{\text{hit}} = O\left(\frac{n}{\Phi}\right)$, where Φ is the conductance of G . Furthermore, for any non-regular graph with maximum degree Δ , average degree d and minimum degree δ , we have

$$t_{\text{hit}} = O\left(\frac{\sqrt{\Delta d}}{\delta} \cdot \frac{n}{\sqrt{1-\lambda_2}}\right)$$

Proof. By [\[LPW06, Lemma 10.2\]](#) and [\[LPW06, Proposition 10.19\]](#),

$$\begin{aligned} t_{\text{hit}} &\leq 2 \max_u \sum_v t_{\text{hit}}(v, u) \pi(v) = \frac{2}{\pi(u)} \sum_{t=0}^{\infty} (p_{u,u}^t - \pi(u)) \\ &\leq \frac{2}{\pi(u)} \sum_{t=0}^{2t_{\text{rel}}} (p_{u,u}^t) + \frac{2}{\pi(u)} \sum_{t=2t_{\text{rel}}+1}^{\infty} (p_{u,u}^t - \pi(u)) \\ &\leq \frac{O(1)}{\pi(u)} \sqrt{t_{\text{rel}}} + \frac{2}{\pi(u)} 2t_{\text{rel}} \sum_{k=0}^{\infty} \left(p_{u,u}^{2(\sqrt{t_{\text{rel}}}+kt_{\text{rel}})} - \pi(u)\right), \end{aligned} \tag{40}$$

where the bound on the first term of the last inequality follows from $p_{u,u}^{\sqrt{t_{\text{rel}}}} \leq O\left(\frac{1}{\sqrt{t_{\text{rel}}}}\right)$ (Lemma B.3).

We can bound the sum as follows using that $\pi(v) = 1/n$ and that for regular graphs any τ it holds that $p_{u,u}^{2\tau} = \sum_{v \in V} (p_{u,v}^\tau)^2$ as follows

$$\begin{aligned}
& \sum_{k=0}^{\infty} \left(p_{u,u}^{2(\sqrt{t_{\text{rel}}+kt_{\text{rel}})}} - \frac{1}{n} \right) = \sum_{k=0}^{\infty} \left(\sum_{v \in V} \left(p_{u,v}^{\sqrt{t_{\text{rel}}+kt_{\text{rel}}}} \right)^2 - \frac{1}{n} \right) \\
&= \sum_{k=0}^{\infty} \left(\sum_{v \in V} \left(p_{u,v}^{\sqrt{t_{\text{rel}}+kt_{\text{rel}}}} \right)^2 - 2 \sum_{v \in V} p_{u,v}^{\sqrt{t_{\text{rel}}+kt_{\text{rel}}}} \cdot \frac{1}{n} + \sum_{v \in V} \frac{1}{n^2} \right) \\
&= \sum_{k=0}^{\infty} \left\| p_{u,\cdot}^{\sqrt{t_{\text{rel}}+kt_{\text{rel}}}} - \frac{1}{n} \right\|^2 = \sum_{k=0}^{\infty} \left\| P^{k \cdot t_{\text{rel}}} p_{u,\cdot}^{\sqrt{t_{\text{rel}}}} - \frac{1}{n} \right\|^2 \leq \sum_{k=0}^{\infty} \lambda_2^{k \cdot t_{\text{rel}}} \left\| p_{u,\cdot}^{\sqrt{t_{\text{rel}}}} - \frac{1}{n} \right\|^2 \\
&= \sum_{k=0}^{\infty} \lambda_2^{k \cdot t_{\text{rel}}} \left(\sum_{v \in V} p_{u,v}^{2\sqrt{t_{\text{rel}}}} - 2 \sum_{v \in V} p_{u,v}^{\sqrt{t_{\text{rel}}}} + \sum_{v \in V} \frac{1}{n^2} \right) \\
&\leq \sum_{k=0}^{\infty} \lambda_2^{k \cdot t_{\text{rel}}} p_{u,u}^{2\sqrt{t_{\text{rel}}}} \stackrel{(*)}{=} O\left(\sum_{k=0}^{\infty} \lambda_2^{k \cdot t_{\text{rel}}} \frac{1}{\sqrt{t_{\text{rel}}}} \right) \stackrel{(**)}{=} O\left(\frac{1}{\sqrt{t_{\text{rel}}}} \right), \tag{41}
\end{aligned}$$

where $(*)$ follows from $p_{u,u}^{\sqrt{t_{\text{rel}}}} \leq O\left(\frac{1}{\sqrt{t_{\text{rel}}}}\right)$ (Lemma B.3) and $(**)$ follows since $f(y) = y^{1/(1-y)}$ is bounded from above by $1/e$ for any $y \in (0, 1)$ and hence the sum is a geometric series. Combining (40) into (41) yields the claim.

To obtain the result for non-regular graphs, we consider the modified Markov chain with transition matrix Q where the loop probability of every vertex is $1 - \frac{\deg(u)}{\Delta}$. As a result, every transition of the walk to another vertex is made with probability $1/\Delta$. Thus Q is symmetric and the stationary distribution π_Q is uniform. We can apply the result from the first statement to Q and it only remains to relate $\lambda_2(Q)$ to $\lambda_2(P)$. The variational characterization of $\lambda_2(P)$ gives:

$$\begin{aligned}
1 - \lambda_2(Q) &= \inf_{\varphi: V \rightarrow \mathbb{R}, \varphi \text{ non-constant}} \frac{\sum_{u,v \in V} (\varphi(u) - \varphi(v))^2 \pi_Q(u) Q_{u,v}}{\sum_{u,v \in V} (\varphi(u) - \varphi(v))^2 \pi_Q(u) \pi_Q(v)} \\
&= \inf_{\varphi: V \rightarrow \mathbb{R}, \varphi \text{ non-constant}} \frac{\sum_{u,v \in V} (\varphi(u) - \varphi(v))^2 \frac{1}{n\Delta}}{\sum_{u,v \in V} (\varphi(u) - \varphi(v))^2 \frac{1}{n^2}}
\end{aligned}$$

Similarly,

$$1 - \lambda_2(P) = \inf_{\varphi: V \rightarrow \mathbb{R}, \varphi \text{ non-constant}} \frac{\sum_{u,v \in V} (\varphi(u) - \varphi(v))^2 \frac{1}{2|E|}}{\sum_{u,v \in V} (\varphi(u) - \varphi(v))^2 \cdot \frac{\deg(u)}{2|E|} \cdot \frac{\deg(v)}{2|E|}}.$$

Comparing the two equations, we can see that

$$1 - \lambda_2(Q) \geq (1 - \lambda_2(P)) \cdot \frac{\Delta}{d} \cdot \left(\frac{d}{\delta} \right)^2.$$

□

It turns out that the hitting time bound of Theorem B.4 is tight in the sense that for any Φ there exists a graph with conductance Φ and hitting time and coalescence time of order $\Omega(n/\Phi)$.

Proposition B.5 ([BGKM16]). *For every $n, d \geq 3$, and constant Φ , there exists a d -regular graph G with n nodes and a constant conductance such that the expected consensus time on G is $\Omega(n)$.*

Furthermore, for every even n , $\Phi > 1/n$, and constant d , there exists a d -regular graph G with $\Theta(n)$ nodes and a conductance of $\Theta(\Phi)$ such that the meeting time on G is $\Omega(n/\Phi)$. Therefore, the coalescence time and hitting time are of order $\Omega(n/\Phi)$.

Theorem B.6. *Let G be any non-regular graph with maximum degree Δ , average degree d and minimum degree δ , we have*

$$t_{\text{meet}} = O\left(\frac{\sqrt{\Delta d}}{\delta} \cdot \frac{n}{\sqrt{1-\lambda_2}}\right).$$

In particular,

$$t_{\text{coal}} = O(t_{\text{hit}} \log(\Delta/\delta)) = O\left(\frac{\sqrt{\Delta d}}{\delta} \cdot \frac{n}{\sqrt{1-\lambda_2}} \log(\Delta/\delta)\right)$$

Furthermore,

$$t_{\text{cov}} = O\left(\frac{\sqrt{\Delta d}}{\delta} \cdot \frac{n}{\sqrt{1-\lambda_2}} \log n\right)$$

The upper bound on t_{meet} and t_{coal} gives $t_{\text{meet}} = O(n^2)$ for cycles and paths, and $t_{\text{meet}} = O(n)$ on regular expanders (since $1/(1-\lambda_2) = O(1)$). It thus improves the bound by Cooper *et al.* [CEOR13, Theorem 1], which states that for any regular graph, $t_{\text{meet}} = O(n/(1-\lambda_2))$.

Proof. The proof of the first part follows from Theorem B.4 and $t_{\text{meet}} \leq 4t_{\text{hit}}$ (Proposition B.9). The second part is due to Theorem B.4 and Theorem 1.3. The last statement follows from Theorem B.6 and the well-known trivial bound $t_{\text{cov}} = O(t_{\text{hit}} \cdot \log n)$. \square

For any $\Gamma = O(1)$ -approximate regular graph, we also improve the best-known bound on the cover time t_{cov} in terms of the eigenvalue gap, which is $t_{\text{cov}} = O(n \log n / (1-\lambda_2))$ established by Broder and Karlin in 1989 [BK89].

As mentioned earlier, $\sum_{t=0}^{t_{\text{mix}}-1} p_{u,u}^t \leq t_{\text{mix}} \cdot \pi(u) + \frac{1}{1-\lambda_2}$ are well-known bounds. The next corollary provides an improvement in many cases:

Corollary B.7. *For any Γ -approximate regular graph $G = (V, E)$*

$$R_{\text{max}} = \max_{u \in V} \sum_{t=0}^{t_{\text{mix}}-1} p_{u,u}^t = O\left(\min\left\{\Gamma \sqrt{\frac{d \log(1/(1-\lambda_2))}{\delta(1-\lambda_2)}}, \Gamma^{3/2} \sqrt{t_{\text{mix}}}\right\}\right) = O\left(\Gamma^{3/2} \cdot \sqrt{\beta}\right).$$

Proof. Thus, using that $p_{u,u}^t$ is non-increasing (e.g., Lemma A.1) and Lemma B.3, we derive for any τ

$$\sum_{t=0}^{\tau-1} p_{u,u}^t \leq 2\Gamma \sqrt{5\tau} + \sum_{t=5n^2}^{\tau-1} p_{u,u}^{5n^2} \leq 2\Gamma \sqrt{5\tau} + \tau \cdot \frac{20\Gamma}{\sqrt{5n}}, \quad (42)$$

where we used that $p_{u,u}^{5n^2} \leq \frac{20\Gamma}{\sqrt{5n}}$, by Lemma B.3. In particular, using $t_{\text{mix}} \leq t_{\text{hit}} = O(\Gamma \cdot n^2)$ ([AF02, Corollary 6.9])

$$\begin{aligned} R_{\text{max}} &\leq 2\Gamma \sqrt{5t_{\text{mix}}} + t_{\text{mix}} \cdot \frac{20\Gamma}{\sqrt{5n}} \\ &= O\left(\Gamma \cdot \sqrt{t_{\text{mix}}} (1 + \sqrt{t_{\text{mix}}/n})\right) \stackrel{t_{\text{mix}}=O(\Gamma \cdot n^2)}{=} O\left(\Gamma^{3/2} \sqrt{t_{\text{mix}}}\right) \end{aligned}$$

In the remainder we derive a bound in terms on R_{\max} in terms of $\frac{\log(1/(1-\lambda_2))}{1-\lambda_2}$. We split the expected number of returns to u at time $x \leq t_{\text{mix}}$ and obtain

$$\begin{aligned} \sum_{t=0}^{t_{\text{mix}}-1} p_{u,u}^t &= \sum_{t=0}^{x-1} p_{u,u}^t + \sum_{t=x}^{t_{\text{mix}}-1} p_{u,u}^t \stackrel{\text{Lemma A.1} \ \& \ (42)}{\leq} \left(2\Gamma \cdot \sqrt{5x} + x \cdot 20\Gamma/n\right) + \sum_{t=x}^{t_{\text{mix}}-1} (\pi(u) + \lambda_2^t) \\ &\leq 2\Gamma \cdot \sqrt{5x} + 21\Gamma \cdot t_{\text{mix}}/n + \frac{\lambda_2^x}{1-\lambda_2}. \end{aligned}$$

Next choose $x = \frac{\ln(1/(1-\lambda_2))}{1-\lambda_2}$. Since $f(y) = y^{1/(1-y)}$ is bounded from above by $1/e$ for any $y \in (0, 1)$, we have

$$\sum_{t=0}^{t_{\text{mix}}-1} p_{u,u}^t \leq 10\Gamma \cdot \sqrt{\frac{d \ln(1/(1-\lambda_2))}{\delta(1-\lambda_2)}} + 21\Gamma \cdot t_{\text{mix}}/n + 1.$$

We next prove that the second term in the bound above is always asymptotically upper bounded by the first one. This is established via a simple case distinction. First, if $1/(1-\lambda_2) \leq n^{1/2}$, then the claim holds because of $t_{\text{mix}} = O(\log n/(1-\lambda_2)) = O(\sqrt{n} \log n)$ and hence $\Gamma t_{\text{mix}}/n = o(\Gamma)$ whereas the first term is $\Omega(\Gamma)$. Secondly, if $1/(1-\lambda_2) \geq n^{1/2}$, then using the same bound on t_{mix} along with the fact that $1-\lambda_2 \geq \frac{\delta}{dn^2}$, where d is the average degree and δ the minimum degree:

By [CG97, Lemma 1.9], we have $1-\lambda_2 \geq \frac{1}{\text{diam } dn}$, and since the diameter of a graph is at most n/δ , we get $1-\lambda_2 \geq \frac{\delta}{dn^2}$.

$$t_{\text{mix}} = O\left(\frac{\log n}{\sqrt{1-\lambda_2} \cdot \sqrt{1-\lambda_2}}\right) = O\left(\sqrt{\frac{dn^2}{\delta}} \frac{2 \log(n^{1/2})}{\sqrt{1-\lambda_2}}\right) = O\left(n \cdot \sqrt{\frac{d \log(1/(1-\lambda_2))}{\delta(1-\lambda_2)}}\right).$$

□

We now derive an extension of [Theorem B.6](#) that is more suited for graphs with a very high degree discrepancy.

Theorem B.8. *Let $G = (V, E)$ be any Γ -approximate regular graph. Then,*

$$t_{\text{meet}} = O\left(\frac{\Gamma^{3/2} \log_2(\Gamma) \cdot \sqrt{\beta}}{\|\pi\|_2^2}\right).$$

Furthermore,

$$t_{\text{meet}} = O\left(\frac{\frac{\log(\Gamma)}{1-\lambda_2} + \log_2(\Gamma) \cdot t_{\text{mix}} \cdot \max_{u \in V} \pi(u)}{\|\pi\|_2^2}\right).$$

We point out that for constant Γ , the first statement of the theorem recovers the second statement of [Theorem B.6](#).

Proof. Similar to the proof of [Theorem B.1](#), we define $S_i := \{u \in V : \deg(u) \in (2^{i-1}, 2^i]\}$, and note that $S_0, \dots, S_{\log_2 n}$ forms a partition of V . Since the graph is Γ -approximate regular, at most $4 + \log_2(\Gamma)$ of the S_i 's are non-empty. Hence there exists a set S_j with

$$\sum_{j \in S_j} \pi(j)^2 \geq \frac{1}{4 + \log_2(\Gamma)} \cdot \|\pi\|_2^2.$$

We will only count collisions on vertices in that bucket, *i.e.*, $Z := \sum_{t=t_{\text{sep}}}^{t_{\text{sep}}+t_{\text{mix}}-1} \mathbf{1}_{X_t=Y_t} \cdot \mathbf{1}_{X_t \in S_j}$, where $(X_t)_{t \geq 0}$ and $(Y_t)_{t \geq 0}$ are two arbitrary walks. Then,

$$\mathbb{E}[Z] \geq \frac{t_{\text{mix}} \cdot \|\pi\|_2^2}{e^2 \cdot (4 + \log(\Gamma))}.$$

Furthermore,

$$\mathbb{E}[Z \mid Z \geq 1] \leq \max_{u \in S_j} \sum_{t=0}^{t_{\text{mix}}-1} \sum_{v \in S_j} (p_{u,v}^t)^2 \leq 4 \max_{u \in S_j} \sum_{t=0}^{t_{\text{mix}}-1} p_{u,u}^{2t} = O\left(\Gamma^{3/2} \cdot \sqrt{\beta}\right)$$

where the second inequality holds due to the fact that vertices in S_j have the same degree up to a factor of 2 and the final inequality holds due to [Corollary B.7](#). Plugging the two bounds into (1) yields

$$\mathbb{P}[Z \geq 1] = \Omega\left(\frac{t_{\text{mix}} \cdot \|\pi\|_2^2}{\Gamma^{3/2} \log_2(\Gamma) \cdot \sqrt{\beta}}\right) =: p.$$

Hence, by iterating over consecutive time-intervals of length $t_{\text{sep}} + t_{\text{mix}} \leq 5t_{\text{mix}}$ that are independent, we conclude

$$t_{\text{meet}} \leq \frac{1}{p} \cdot 5t_{\text{mix}} = O\left(\frac{\Gamma^{3/2} \log_2(\Gamma) \cdot \sqrt{\beta}}{\|\pi\|_2^2}\right).$$

For the second statement, we also have, by [Lemma A.1](#),

$$4 \max_{u \in S_i} \sum_{t=0}^{t_{\text{mix}}-1} p_{u,u}^t \leq 4 \cdot \frac{1}{1 - \lambda_2} + t_{\text{mix}} \cdot \max_{u \in V} \pi(u),$$

and the bound on t_{meet} is derived in exactly the same way as before. \square

B.2 Relating Meeting Time to t_{hit}

In this section we prove the following proposition which can be seen as an analogous version of [\[AF02, Proposition 14.5\]](#) in discrete time.

Proposition B.9. *For any graph $G = (V, E)$ and $u, v \in V$ we have*

$$\left(\min_{u' \in V} t_{\text{hit}}(\pi, u') + t_{\text{hit}}(u, v) - t_{\text{hit}}(\pi, v)\right)/2 \leq t_{\text{meet}}(u, v) \leq 2\left(\max_{u' \in V} t_{\text{hit}}(\pi, u') + t_{\text{hit}}(u, v) - t_{\text{hit}}(\pi, v)\right).$$

Consequently, for any graph we have $t_{\text{meet}} \leq 4t_{\text{hit}}$ and for any vertex transitive graph G we have $t_{\text{hit}}/2 \leq t_{\text{meet}} \leq 2t_{\text{hit}}$.

Proof. We define a pair of chains $((X_t)_{t \geq 0}, (Y_t)_{t \geq 0})$ with arbitrary start vertices $X_0, Y_0 \in V$, called *sequential* random walks, by

$$X_{t+1} = \begin{cases} v \in N(X_t) \text{ w.p. } \frac{1}{2|N(X_t)|} & \text{if } t \text{ is even} \\ X_t & \text{otherwise} \end{cases},$$

and

$$Y_{t+1} = \begin{cases} v \in N(Y_t) \text{ w.p. } \frac{1}{2|N(Y_t)|} & \text{if } t \text{ is odd} \\ Y_t & \text{otherwise} \end{cases}.$$

In particular, for odd t (even t , respectively) the random-walk is lazy meaning $X_{t+1} = X_t$ (and $Y_{t+1} = Y_t$, respectively).

Consider two “non-sequential” random walks $(X'_t)_{t \geq 0}$ and $(Y'_t)_{t \geq 0}$ with $X'_0 = X_0$ and $Y'_0 = Y_0$. We will couple their decisions with the walks $(X_t)_{t \geq 0}$ and $(Y_t)_{t \geq 0}$, by setting $X'_t = X_{2t}$ and $Y'_t = Y_{2t}$. Due to this coupling and since each random walk is lazy w.p. $1/2$,

$$\mathbb{P} [X'_{t+1} = Y'_{t+1} \mid X_{2t} = Y_{2t}] = \mathbb{P} [X_{2t+2} = Y_{2t+2} \mid X_{2t} = Y_{2t}] \geq 1/4,$$

and

$$\mathbb{P} [X'_{t+1} = Y'_{t+1} \mid X_{2t+1} = Y_{2t+1}] = \mathbb{P} [X_{2t+2} = Y_{2t+2} \mid X_{2t+1} = Y_{2t+1}] = 1/2$$

Let $t_{\text{meet}}^{\text{seq}}(u, v)$ be the meeting time of the sequential chains X_t and Y_t , *i.e.*,

$$t_{\text{meet}}^{\text{seq}}(u, v) = \min\{t \geq 0 \mid X_t = Y_t, X_0 = u, Y_0 = v\}$$

and $t_{\text{meet}}^{\text{seq}} = \max_{u,v} t_{\text{meet}}^{\text{seq}}(u, v)$. We seek to relate t_{meet} with $t_{\text{meet}}^{\text{seq}}$. Clearly, $t_{\text{meet}}^{\text{seq}}/2 \leq t_{\text{meet}}$ since a meeting of $X'_t = Y'_t$ implies that $X_{2t} = Y_{2t}$. For an upper bound on t_{meet} recall that X'_t and Y'_t meet, *i.e.*, $X'_t = Y'_t$ w.p. at least $1/4$ whenever $X_{2t-2} = Y_{2t-2}$ or $X_{2t-1} = Y_{2t-1}$. Hence, by independence $t_{\text{meet}} = \max_{u,v} t_{\text{meet}}(u, v) \leq 4(t_{\text{meet}}^{\text{seq}}/2) = 2t_{\text{meet}}^{\text{seq}}$. We conclude,

$$t_{\text{meet}}^{\text{seq}}(u, v)/2 \leq t_{\text{meet}}(u, v) \leq 2t_{\text{meet}}^{\text{seq}}(u, v). \quad (43)$$

We proceed by deriving upper and lower bounds on $t_{\text{meet}}^{\text{seq}}$, which gives us bounds on t_{meet} . We will make use of the following statement that is a weaker version of the original statement Aldous and Fill [AF02, Proposition 3.3]. For all $u, v \in V$ we have

$$\min_{u' \in V} t_{\text{hit}}(\pi, u') \leq t_{\text{meet}}^{\text{seq}}(u, v) - (t_{\text{hit}}(u, v) - t_{\text{hit}}(\pi, v)) \leq \max_{u' \in V} t_{\text{hit}}(\pi, u')$$

Using (43) we derive,

$$t_{\text{meet}}(u, v) \leq 2t_{\text{meet}}^{\text{seq}}(u, v) \leq 2(t_{\text{hit}}(u, v) - t_{\text{hit}}(\pi, v) + \max_{u' \in V} t_{\text{hit}}(\pi, u')),$$

and

$$t_{\text{meet}}(u, v) \geq t_{\text{meet}}^{\text{seq}}(u, v)/2 \geq (\min_{u' \in V} t_{\text{hit}}(\pi, u') + (t_{\text{hit}}(u, v) - t_{\text{hit}}(\pi, v)))/2,$$

which yields the first part of the claim. For vertex transitive chains we get using $t_{\text{hit}}(\pi, u) = t_{\text{hit}}(\pi, u')$ for all $u, u' \in V$ and thus

$$t_{\text{meet}}^{\text{seq}}(u, v) = t_{\text{hit}}(u, v).$$

Thus, putting everything together and fixing $u, v \in V$ to be the nodes maximizing $t_{\text{hit}}(u, v)$, we derive

$$t_{\text{hit}} = t_{\text{hit}}(u, v) = t_{\text{meet}}^{\text{seq}}(u, v) \leq t_{\text{meet}}^{\text{seq}} \leq 2t_{\text{meet}}.$$

Similarly,

$$t_{\text{meet}} \leq 2t_{\text{meet}}^{\text{seq}} = 2 \max_{u,v} t_{\text{meet}}^{\text{seq}}(u, v) \leq 2t_{\text{hit}}.$$

This yields [Proposition B.9](#). □

B.3 Proof of [Theorem 1.4](#)

The proof follows from [Theorem B.4](#) and [Theorem B.6](#).

C Applications to Concrete Topologies

Here we derive t_{meet} and t_{coal} on specific topologies. Note that more general bounds for certain graph classes like regular graphs or vertex-transitive graphs have been stated earlier, see, *e.g.*, [Theorem 1.3](#) or [Appendix B](#).

C.1 2-Dimensional Grids/Tori and Paths/Cycles

Next we apply our machinery to the 2-dimensional grid and the 2-dimensional torus. For the continuous case a manual approach to bound meeting and coalescence times can be found in [[Cox89](#)]. Thanks to our general results, we can not only easily derive the correct bound on t_{meet} , but also on t_{coal} . First, we recall the following well-known fact that for 2-dim. grid and torus: For any integer $t = O(n)$,

$$p_{u,u}^t = \pi(u) + \Omega(t^{-1}), \quad (44)$$

which can be derived, *e.g.*, by using the central limit theorem. Further, $t_{\text{mix}} = \Theta(n)$, and combining these two results, we immediately obtain

$$C_{\min} := \min_{u \in V} \sum_{t=0}^{t_{\text{mix}}} \sum_{v \in V} (p_{u,v}^t)^2 \geq \frac{1}{2} \min_{u \in V} \sum_{t=0}^{t_{\text{mix}}} p_{u,u}^{2t} = \Omega(\log n).$$

Thus, by [Theorem B.1.\(iii\)](#), $t_{\text{meet}} = \Omega(n \log n)$. For the upper bound, we apply [Theorem 1.3](#) together with the well-known bound $t_{\text{hit}} = O(n \log n)$ to derive $t_{\text{coal}} = O(t_{\text{hit}}) = O(n \log n)$.

For cycles or paths, the corresponding formula to (44) is, for $t = O(n^2)$,

$$p_{u,u}^t = \pi(u) + \Omega(t^{-1/2}).$$

Hence $C_{\min} = \Omega(n)$, and therefore the third statement of [Theorem B.1](#) implies $t_{\text{meet}} = \Omega(n^2)$. For the upper bound, we apply [Theorem 1.3](#) together with $t_{\text{hit}} = O(n^2)$ to derive $t_{\text{coal}} = O(t_{\text{hit}}) = O(n^2)$. Alternatively, the upper bound on t_{coal} could be also shown by using $t_{\text{mix}} = O(n^2)$ and applying the third statement [Theorem B.6](#).

C.2 d -Dimensional Grids and Tori, $d \geq 3$

Here the bounds on t_{meet} and t_{coal} follow immediately from our general results. First, for any regular graphs we have $t_{\text{meet}} = \Omega(n)$ ([Theorem B.1.iii](#)). Further, it is well-known that $t_{\text{hit}} = O(n)$ (*e.g.*, [[LPW06](#)]), and the result follows by $t_{\text{coal}} = O(t_{\text{hit}})$ shown in [Theorem 1.3](#). Alternatively, we could also use $t_{\text{meet}} = O(t_{\text{hit}})$ ([Proposition B.9](#)) to deduce $t_{\text{meet}} = O(n)$. Combining this with the fact that $t_{\text{mix}} = O(n^{2/d})$ [[AF02](#)], [Theorem 1.1](#) yields the correct bound $t_{\text{coal}} = O(n)$.

C.3 Hypercubes

Tight bounds for the hypercube can be obtained through different tools we provide.

Firstly, it follows trivially from [Theorem 1.3](#): Since the hypercube is regular (in fact, it is even vertex-transitive) it suffices to consider the hitting time. We have $t_{\text{hit}} = O(n)$ (see *e.g.*, [[Lov93](#)]) and recall that $t_{\text{hit}} = \Omega(n)$ by [Theorem B.1](#). Hence applying [Theorem 1.3](#) yields $t_{\text{coal}} = \Theta(t_{\text{hit}}) = \Theta(n)$.

Alternatively, we could also use the more elementary bound $t_{\text{meet}} \leq 2t_{\text{hit}}$ by [Proposition B.9](#) to conclude $t_{\text{meet}} = O(n)$. Since it is a well-known fact that $t_{\text{mix}} = O(\log n \cdot \log \log n)$ [[LPW06](#)], we obtain by [Theorem 1.1](#) that $t_{\text{coal}} = \Theta(t_{\text{meet}}) = \Theta(n)$.

C.4 (Regular) Expanders

It is not surprising that on regular expander graphs, we have $t_{\text{coal}} = \Theta(n)$ and there is a multitude of approaches to establish this (for instance, the result is a consequence of the main result by [CEOR13]). With regard to our bounds, the easiest route is to follow the arguments for the hypercube: Combine the result $t_{\text{hit}} = O(n)$ (e.g., [BK89]) together with our bound $t_{\text{coal}} = O(t_{\text{hit}})$ (Theorem 1.3). The lower bound $t_{\text{coal}} \geq t_{\text{meet}} = \Omega(n)$ holds for any regular graph.

C.5 Real World Graph Models

There is a variety of different graph models for “real world” networks. In this subsection we demonstrate that random walks coalesce quickly on these graphs by establishing several bounds on t_{coal} which are sublinear in n .

First note that common features of real world graph models are (i) a power law degree distribution with exponent $\beta \in (2, 3)$ and (ii) high expansion, i.e., $1 - \lambda_2$ is not too large, and hence $t_{\text{mix}} = O(\log n)$. Notice that (i) $\beta \in (2, 3)$ implies that w.h.p. we have $\Delta = O(n^{1-\varepsilon})$, and hence $\|\pi\|_2^2 \leq \max_{u \in V} \pi(u) \leq n^{-\varepsilon}$, for $\varepsilon > 0$.

For the sake of concreteness, let us take a specific model by Gkantsidis, Mihail and Saberi [GMS03], which was also analyzed by Cooper *et al.* [CEOR13]. In this model, for some $\alpha \in (2, 3)$ we generate a random graph which has $\Theta(n/d^\alpha)$ vertices of degree d and an eigenvalue gap $1/(1 - \lambda_2) = O(\log^2 n)$. Cooper *et al.* [CEOR13] derived the general bound $t_{\text{coal}} = O(\frac{1}{1-\lambda_2} \cdot (\|\pi\|_2^2 + \log^4 n))$, which implies $t_{\text{coal}} = O(n^{(\alpha-1)/2} \cdot \log^2 n)$ - a sublinear bound on the coalescing time. However, this leaves open how close t_{coal} and t_{meet} are.

Combining Theorem 1.1 with the fact that $t_{\text{mix}} = O(\log n / (1 - \lambda_2)) = O(\log^3 n)$, we immediately obtain $t_{\text{coal}} = \Theta(t_{\text{meet}})$, without having to know the actual value of t_{meet} .¹³

More generally, we have the following result, saying that we have $t_{\text{coal}} = \Theta(t_{\text{meet}})$ whenever t_{mix} is slightly smaller than $1/\|\pi\|_2^2$:

Theorem C.1. *Let $G = (V, E)$ be any graph. Then,*

$$t_{\text{coal}} = O\left(t_{\text{meet}} \cdot \left(1 + \sqrt{t_{\text{mix}} \cdot \|\pi\|_2^2 \cdot \log n}\right)\right).$$

In particular, whenever $t_{\text{mix}} \cdot \log^2 n \leq 1/\|\pi\|_2^2$, we have $t_{\text{coal}} = \Theta(t_{\text{meet}})$.

Proof. First, by the third statement of Theorem B.1, we have $t_{\text{meet}} \geq \frac{1}{9\|\pi\|_2^2}$. Inserting this into Theorem 1.1 yields the upper bound. The lower bound for the setting $t_{\text{mix}} \cdot \log^2 n \leq 1/\|\pi\|_2^2$ trivially holds since $t_{\text{coal}} \geq t_{\text{meet}}$. \square

It is worth comparing this result with the bound derived by Cooper *et al.* [CEOR13]:

$$t_{\text{coal}} = O\left(\frac{1}{1 - \lambda_2} \cdot \left(\frac{1}{\|\pi\|_2^2} + \log^4 n\right)\right). \quad (45)$$

The advantage of (45) is that requires relatively little knowledge about G ; only $\frac{1}{1-\lambda_2}$ and $\frac{1}{\|\pi\|_2^2}$ (which is equivalent to knowing the degree distribution) are needed. One potential drawback of the bound in (45) however, is that it involves the product of two factors $\frac{1}{1-\lambda_2}$ and $\frac{1}{\|\pi\|_2^2}$, each of which is

¹³That being said, deriving the correct bound on t_{meet} is an interesting open problem. So far, it seems rather difficult to use one of our “off-the-shelf” bounds or the results from [CEOR13]. One potential route towards a tight bound may involve stronger bounds on R_{max} , as suggested by the second upper bound on $t_{\text{meet}}(u, v)$ in Theorem B.1.

a lower bound on the meeting time on its own. For instance for regular graphs, by [Theorem B.6](#), we immediately obtain that $t_{\text{coal}} = O(t_{\text{meet}} \log n) = O\left(n \log n \cdot \sqrt{\frac{1}{1-\lambda_2} \cdot \log\left(\frac{1}{1-\lambda_2}\right)}\right)$. As a consequence, for regular graphs, our bound improves over the bound in [\(45\)](#) whenever $\frac{1}{1-\lambda_2} \geq \log^{2+\varepsilon} n$ for an arbitrarily small constant $\varepsilon > 0$.

It is also interesting to consider an alternative graph model for real world networks, proposed by Mihail et al. [\[MPS06\]](#). Also in this model, the degree distribution has the same Power law with exponent $\alpha \in (2, 3)$, but there is a stronger bound on the spectral gap, $\frac{1}{1-\lambda_2} = O(1)$ [\[MPS06\]](#). Hence [Theorem C.1](#) implies $t_{\text{coal}} = \Theta(t_{\text{meet}})$. Further, thanks to [Proposition B.2](#) (or alternatively, the bound by Cooper *et al.* [\(45\)](#)) we get the explicit bound $t_{\text{coal}} = O(1/\|\pi\|_2^2)$, which is asymptotically tight due to the trivial lower bound $t_{\text{meet}} = \Omega(1/\|\pi\|_2^2)$ ([Theorem B.1](#)).

C.6 Binary Trees

In this subsection, we derive a lower bound $t_{\text{meet}} = \Omega(n \log n)$ for complete binary trees. Unfortunately, this bound does not follow directly from our general results and a manual analysis is required. To some extent, this is due to the structural difference between nodes close to the leaves and nodes close to the root. While a collision close to the leaves triggers $\Theta(n \log n)$ expected additional collisions, a collision near the root triggers only $\Theta(n)$ additional collisions.

Our proof consists of the following two steps. In [Section C.6.1](#), we first provide a lower bound on the probability that a random walk starting from any nodes $u \in V$ is on a leaf after $O(\log n)$ steps. We also show that any t -step probability $p_{u,v}^t$ is at the most return probability for a leaf. Both results shown in [Lemma C.3](#) are derived by projecting the random walk on the tree to a random walk on a weighted path of length $\log_2 n - 1$.

In [Section C.6.2](#), we proceed to analyzing the expected number of collisions between two random walks in n steps. The main component is [Lemma C.5](#), establishing that this number is at least $\Omega(\log n)$ provided the walks start from the same vertex not too far from the root. This result is complemented by a union-bound type argument in [Lemma C.4](#), showing that it is unlikely that two random walks collide on a vertex close to the root. Combining the two results and applying them to [\(1\)](#) establishes the desired lower bound $t_{\text{meet}} = \Omega(n \log n)$.

For the other bounds on t_{meet} and t_{coal} , we combine $t_{\text{coal}} = O(t_{\text{hit}})$ ([Theorem 1.3](#)) with the well-known fact $t_{\text{hit}} = O(n \log n)$ (cf. [\[AF02\]](#)) to obtain $t_{\text{coal}} = O(n \log n)$. Together with the established lower bound, this shows that t_{meet} and t_{coal} are both of order $\Theta(n \log n)$.

C.6.1 Bounds on the t -step probabilities

We assume that the complete binary tree has $\log_2 n - 1$ levels, *i.e.*, there are $n/2$ leaves and the total number of nodes is $n - 1$. We define $\mathcal{L} \subseteq V$ to be the set of leaves. For the analysis, it will be helpful to relate a random walk on the binary tree to a corresponding random walk on a weighted path \tilde{G} of length $\log_2 n - 1$ with nodes $\tilde{V} = \{1, 2, \dots, \log_2 n - 1\}$, where each vertex on the path corresponds to all vertices in the binary tree on the same level. Let Q denote the $(\log_2 n - 1) \times (\log_2 n - 1)$ transition matrix of the corresponding weighted random walk. For $i, j \in \tilde{V}$ we have

$$q_{ij} = \begin{cases} 1/2 & \text{if } i = j \\ 1/3 & \text{if } j > i, i \neq 1 \\ 1/2 & \text{if } j > i, i = 1 \\ 1/6 & \text{if } j < i, i \neq \log_2 n - 1 \\ 1/2 & \text{if } j < i, i = \log_2 n - 1. \end{cases}$$

Let $\tilde{\pi}$ denote the stationary distribution of this process. Since the random walk on \tilde{G} is also lazy, [Lemma A.1](#) implies that $q_{\log_2 n-1, \log_2 n-1}^t \geq \frac{1}{4}$ for all $t \geq 0$. Define $\tau_{u, \mathcal{L}} := \min_{v \in \mathcal{L}} \tau_{u, v}$ to be the first time-step a leaf is visited, where the walk starts at u ; $\tau_{u, \mathcal{L}} = 0$ if u is a leaf. We will frequently use the following two simple facts about random walks on binary trees:

Lemma C.2 ([\[Moo73\]](#)). *Let G be any tree, and u and v be two adjacent nodes. Then $t_{\text{hit}}(u, v) = 2 \cdot n_{u, v} - 1$, where $n_{u, v}$ is the number of vertices in the subtree containing u obtained by deleting the edge $\{u, v\}$.*

Lemma C.3. *Let G be a complete binary tree, and let u be an arbitrary node. Then the following statements hold: there is a constant $c_1 > 0$, so that*

1. *for any $t \geq c_1 \log_2 n$, $p_{u, \mathcal{L}}^t \geq \frac{1}{5}$. Moreover, if u is a leaf, then the same inequality holds for all $t \geq 0$.*
2. *For any vertex $v \in V$, $p_{u, v}^t \leq 7 \cdot p_{w, w}^{t-2c_1 \log_2 n}$, where w is any leaf.*

Proof. To derive a lower bound on $p_{u, \mathcal{L}}^t$, we consider the contracted binary tree \tilde{G} . Recall that $\tau_{u, \mathcal{L}}$ is the random variable of the first time-step at which the random walk visits a leaf where the random walk starts at u ; $\tau_{u, \mathcal{L}} = 0$ if u is a leaf. By conditioning on the first visit to a leaf,

$$p_{u, \mathcal{L}}^t \geq \sum_{s=0}^{2c_1 \log_2 n} \mathbb{P}[\tau_{u, \mathcal{L}} = s] \cdot q_{\mathcal{L}, \mathcal{L}}^{t-s} \geq \frac{1}{4} \cdot (1 - \mathbb{P}[\tau_{u, \mathcal{L}} > 2c_1 \log_2 n]),$$

since the chain Q can be seen as a projection of P to the line. Our next claim is that $\mathbb{P}[\tau_{u, \mathcal{L}} > c_1 \log_2 n] \leq n^{-2}$, provided that the constant $c_1 > 0$ is sufficiently large. This can be derived by coupling the random walk on a binary tree, starting from the root, with a random walk on the integers, starting from zero and waiting until the random walk reaches the vertex $\log_2 n - 1$. Using a Chernoff bound for $X := \sum_{i=1}^{c_1 \log_2 n} X_i$, with $\mathbb{P}[X_i = +1] = 1/3$, $\mathbb{P}[X_i = -1] = 1/6$ and $\mathbb{P}[X_i = 0] = 1/2$, we conclude that $\mathbb{P}[\tau_{u, \mathcal{L}} > c_1 \log_2 n] \leq n^{-2}$, which implies the first statement.

To prove the second statement, consider the first $c_1 \log_2 n$ steps of a random walk starting at u . Similarly as before,

$$p_{u, v}^t \leq \sum_{s=0}^{c_1 \log_2 n} \mathbb{P}[\tau_{u, \mathcal{L}} = s] \cdot \max_{w \in \mathcal{L}} p_{w, v}^{t-s} + \mathbb{P}[\tau_{u, \mathcal{L}} > 4 \log_2 n].$$

As seen above, $\mathbb{P}[\tau_{u, \mathcal{L}} > c_1 \log_2 n] \leq n^{-2}$ and therefore

$$p_{u, v}^t \leq \max_{t-c_1 \log_2 n \leq s \leq t} \max_{w \in \mathcal{L}} p_{w, v}^s + n^{-2}. \tag{46}$$

Let us now compare $p_{w, v}^t$ to $p_{v, w}^t$. Since the random walk is time-reversible, we have

$$p_{w, v}^t \cdot \pi(w) = p_{v, w}^t \cdot \pi(v),$$

and hence

$$p_{u, v}^t \leq \max_{t-c_1 \log_2 n \leq s \leq t} \max_{w \in \mathcal{L}} 3 \cdot p_{v, w}^s + n^{-2}.$$

Applying [\(46\)](#) to each $p_{v, w}^s$, we conclude that

$$p_{u, v}^t \leq 6 \cdot \max_{t-2c_1 \log_2 n \leq s \leq t} \max_{w, w' \in \mathcal{L}} p_{w, w'}^s + 4n^{-2}.$$

Further, by symmetry $p_{w,w'}^s$ is maximized if $w = w'$, so that

$$p_{u,v}^t \leq 6 \cdot \max_{t-2c_1 \log_2 n \leq s \leq t} p_{w,w}^s + 4n^{-2},$$

where w is any leaf. Applying [Lemma A.1](#), it follows that the maximum is attained for $s = t - 2c_1 \log_2 n$ and $p_{w,w}^s \geq \frac{1}{2n-2}$, which implies the second statement. \square

C.6.2 Establishing the Lower Bound on the Meeting Time

We now prove that the meeting time on binary trees is $\Omega(n \log n)$. The intuition for this is as follows. While two random walks of length $\Theta(n \log n)$ will lead to $\Theta(\log n)$ expected collisions, it turns out that the distribution of collisions is poorly concentrated. In fact we will prove that, conditional on the existence of at least one collision, the expected number of total collisions is $\Omega(\log n)$. This will imply the desired lower bound on the meeting time. A slight complication is that the collision could occur on different nodes, which is why we will first bound the probability for a collision to occur close to the root.

Let us define U to be the set of all nodes that have distance at least $\frac{1}{2} \log_2 n$ from the root. Note that $|V \setminus U| \leq 2 \cdot \sqrt{n}$. Further, let \mathcal{E} denote the event that two random walks starting from the stationary distribution of length $n \log_2 n$ meet on a vertex in $V \setminus U$.

Lemma C.4. *We have $\mathbb{P}[\mathcal{E}] \leq n^{-1/3}$.*

Proof. By the Union Bound, $\mathbb{P}[\mathcal{E}] \leq \sum_{t=1}^{n \log_2 n} \sum_{u \in V \setminus U} \pi(u)^2 = n \log_2 n \cdot 2\sqrt{n} \cdot \left(\frac{2}{n}\right)^2 \leq \frac{8 \log_2 n}{\sqrt{n}}$. \square

Lemma C.5. *For any node $u \in U$ and any $2\sqrt{n} \leq t \leq n/200$, we have $\sum_{v \in \mathcal{L}} (p_{u,v}^t)^2 = \Omega(1/t)$.*

Proof. Recall that $\sum_{v \in \mathcal{L}} (p_{u,v}^t)^2$ is the probability of two non-interacting, independent random walks starting from u to meet at the same leaf at time t .

Our first claim is that with probability at least $c_1 > 0$, both random walks reach a leaf before returning to u within $4 \log_2 n$ steps. To prove this claim, recall that with probability at least $1 - 2n^{-2}$, both random walks reach a leaf before step $4 \log_2 n$. Secondly, by [[Lov93](#), Proposition 2.3], applied to the collapsed binary tree \tilde{G} , where node u is at level ℓ and the leafs are in level $\log_2 n - 1$, it follows that the probability that a random walk starting at u visits a leaf before returning to u is

$$p := \frac{1}{(t_{\text{hit}}^{\tilde{G}}(\ell, \log_2 n - 1) + t_{\text{hit}}^{\tilde{G}}(\log_2 n - 1, \ell)) \cdot \pi_{\ell}^{\tilde{G}}}.$$

Further, $t_{\text{hit}}^{\tilde{G}}(\ell, \log_2 n - 1) = O(\log n)$ and $t_{\text{hit}}^{\tilde{G}}(\log_2 n - 1, \ell) \leq \min_{w \in \mathcal{L}} t_{\text{hit}}(w, u) \leq 2^{\log_2 n - \ell + 2}$ by [Lemma C.2](#) and $\pi_{\ell}^{\tilde{G}} \leq 2^{-\log_2 n + \ell + 2}$, where $\text{dist}(\mathcal{L}, u) = \min_{w \in \mathcal{L}} \text{dist}(w, u)$. Hence p is at least some constant > 0 . Hence with probability at least $p^2 - 2n^{-1}$, both random walks reach a leaf before time $4 \log_2 n$ without returning to u .

Consider now the original binary tree, and one of the two random walks starting from a leaf w at some time $\in [1, 4 \log_2 n]$ up until time step $t - 4 \log_2 n$. Consider the shortest path from w to the root, and let z be a node that is on this shortest path and has distance $\log_2(100t)$ from w . Applying [Lemma C.2](#), it follows that $t_{\text{hit}}(w, z) = \sum_{i=1}^{\log_2(100t)} 2^i - 1 \geq 49t$. By Markov's inequality,

$$\mathbb{P}[T_{\text{hit}}(w, z) \geq 2t_{\text{hit}}(w, z)] \leq 1/2. \tag{47}$$

Now divide the random walk into consecutive epochs of length $2t_{\text{hit}}(w, z) + 4 \log_2 n$. Combining (47) and [Lemma C.3](#) it follows that the random walk will visit the vertex z in each epoch with

probability at least $\frac{1}{2}$, conditional on having not visited the vertex z in any of the previous epochs. Therefore for any integer $\lambda \geq 1$,

$$\mathbb{P}[T_{\text{hit}}(w, z) \geq \lambda \cdot 3 \cdot t_{\text{hit}}(w, z)] \leq \mathbb{P}[T_{\text{hit}}(w, z) \geq \lambda \cdot (2 \cdot t_{\text{hit}}(w, z) + 4 \log_2 n)] \leq 2^{-\lambda},$$

where the first inequality holds since $t_{\text{hit}}(w, z) = \Omega(t) = \omega(\log n)$. Hence,

$$t_{\text{hit}}(w, z) \leq \frac{1}{20} t_{\text{hit}}(w, z) + \mathbb{P}\left[T_{\text{hit}}(w, z) \geq \frac{1}{20} t_{\text{hit}}(w, z)\right] \cdot c_1 \cdot 3t_{\text{hit}}(w, z) + \sum_{\lambda=c_1}^{\infty} 2^{-\lambda} \cdot \lambda \cdot 3t_{\text{hit}}(w, z),$$

and it follows that by choosing the constant $c_1 > 0$ large enough, there is a constant $c_2 = c_2(c_1) > 0$ so that

$$\mathbb{P}\left[T_{\text{hit}}(w, z) \geq \frac{1}{10} t_{\text{hit}}(w, z)\right] \geq c_2.$$

Hence with probability at least $c_2 > 0$, the random walk does not reach the node z before time t . Further, with probability at least $1 - n^{-2}$, the random walk visits a leaf, say, w' , before step t , say at step s , and therefore by [Lemma C.3](#), the random walk is at a leaf at step t with probability at least $(1 - n^{-2}) \cdot \frac{1}{5}$.

Now define $\mathcal{F}_1 := \{T_{\text{hit}}(w, z) \geq t\}$ and $\mathcal{F}_2 := \{X_t \in \mathcal{L}\}$. Clearly, the events \mathcal{F}_1 and \mathcal{F}_2 are positively correlated so that

$$\mathbb{P}[\mathcal{F}_1 \cap \mathcal{F}_2] \geq \mathbb{P}[\mathcal{F}_1] \cdot \mathbb{P}[\mathcal{F}_2] \geq c_2 \cdot (1 - n^{-2}) \cdot \frac{1}{5}.$$

Combining all the events, we conclude that with constant probability $c_3 > 0$ both random walks are on a leaf at step t and have never left the subtree with root z . For one walk, the distribution will be uniform over all the leaves within a subtree whose root is the vertex closest to the root ever visited. Hence let \mathcal{L}_1 be all the leaves that have a non-zero probability to be visited at step t by the first random walk, and \mathcal{L}_2 similarly. W.l.o.g. let $|\mathcal{L}_1| \leq |\mathcal{L}_2|$ and observe that $\mathcal{L}_1 \subseteq \mathcal{L}_2$ since both walks start at the same node. Therefore,

$$\sum_{v \in \mathcal{L}} (p_{u,v}^t)^2 \geq (p^2 - 2n^{-1}) \cdot c_3 \cdot \sum_{v \in \mathcal{L}_1 \cap \mathcal{L}_2} \frac{1}{|\mathcal{L}_1|} \cdot \frac{1}{|\mathcal{L}_2|} \geq (p^2 - 2n^{-1}) \cdot c_3 \cdot \frac{1}{|\mathcal{L}_2|} \geq (p^2 - 2n^{-1}) \cdot c_3 \cdot \frac{1}{10t},$$

where the last inequality holds since we are conditioning on the event that none of the two random walks reaches the vertex z . \square

Theorem C.6. *For the binary tree it holds that $t_{\text{meet}} = \Omega(n \log n)$.*

Proof. We first only consider collisions on nodes in U by two random walks $(X_t)_{t \geq 0}$, $(Y_t)_{t \geq 0}$ starting from the stationarity distribution. More formally, we are interested in the random variable

$$Z := \sum_{t=1}^{cn \log_2 n} \sum_{v \in V \setminus U} \mathbf{1}_{X_t = Y_t = v}.$$

By linearity of expectations,

$$\mathbb{E}[Z] \leq \sum_{t=1}^{cn \log_2 n} \sum_{v \in V} \mathbb{P}[X_t = Y_t = v] \leq cn \log_2 n \cdot \sum_{v \in V} \pi(v)^2 \leq n \log_2 n \cdot n \cdot \left(\frac{2}{n}\right)^2 = 4c \log_2 n,$$

and clearly,

$$\mathbb{P}[Z \geq 1] = \frac{\mathbb{E}[Z]}{\mathbb{E}[Z \mid Z \geq 1]}.$$

Hence to derive an upper bound on $\mathbb{P}[Z \geq 1]$, we will derive a lower bound on $\mathbb{E}[Z \mid Z \geq 1]$. In order to this, it will be helpful if we can work under the assumption that the first collision occurs in the first half of the walk. To this end, let Z_1 be the indicator random variable that is 1 if a collision appears before time $\frac{1}{2}n \log_2 n$ and Z_2 be the indicator random variable if a collision appears after time $\frac{1}{2}n \log_2 n$. Applying [Observation A.8](#), we get $\mathbb{P}[Z_1 \geq 1 \mid Z \geq 1] \geq \frac{1}{2}$. Thus,

$$\begin{aligned} \mathbb{E}[Z \mid Z \geq 1] &\geq \mathbb{P}[Z_1 \geq 1 \mid Z \geq 1] \cdot \mathbb{E}[Z \mid Z_1 \geq 1] \\ &\geq \frac{1}{2} \cdot \sum_{t=0}^{t_{\text{mix}}} \sum_{v \in V} (p_{u,v}^t)^2 \\ &\geq \frac{1}{2} \cdot \sum_{t=2\sqrt{n}}^{t_{\text{mix}}} \sum_{v \in \mathcal{L}} (p_{u,v}^t)^2 \geq \sum_{t=2\sqrt{n}}^{t_{\text{mix}}} \Omega(1/t) = c_4 \log_2 n, \end{aligned}$$

where $c_4 > 0$ is a constant (the penultimate inequality is due to [Lemma C.5](#)).

Consequently, for the modified process where collisions are only allowed on nodes in U

$$\mathbb{P}[Z \geq 1] \leq \frac{4c \log_2 n}{c_4 \log_2 n}.$$

Hence by choosing $c = \min\{c_4/8, 1\}$, it follows that $\mathbb{P}[Z \geq 1] \leq 1/2$.

This implies that with probability at least $1/2$, no collision occurs on nodes in U before time $cn \log_2 n$. Furthermore, by [Lemma C.4](#), with probability at least $1 - n^{-1/3}$ there is no collision on nodes in $V \setminus U$. Therefore, with probability at least $1/2 - n^{-1/3}$, there is no collision among two random walks before time $cn \log_2 n$, and we have shown that $t_{\text{meet}} = \Omega(n \log n)$. \square

C.7 Star

Clearly, the coalescing time is $\Theta(\log n)$ which could be easily shown by a direct analysis. For the sake of completeness, we point out that the upper bound also follows from [Proposition 3.4](#) and the fact that $t_{\text{meet}} = O(1)$. The matching lower bound follows from the general bound $t_{\text{coal}} = \Omega(\log n)$, holding for *any graph* (see [Lemma 5.2](#)).