On coalescence time in graphs–When is coalescing as fast as meeting?

VARUN KANADE, University of Oxford and The Alan Turing Institute, UK
FREDERIK MALLMANN-TRENN, King's College London, UK
THOMAS SAUERWALD, University of Cambridge, UK

Coalescing random walks is a fundamental distributed 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 et al. [14] and Cooper et al. [19], 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. Finally, we prove a tight worst case bound for the coalescence time of $O(n^3)$. By duality, our results yield identical bounds on the voter model.

Our techniques also yield a new bound on the hitting time and cover time of regular graphs, improving and tightening previous results by Broder and Karlin [12], as well as those by Aldous and Fill [2].

CCS Concepts: • **Mathematics of computing → Markov processes**.

Additional Key Words and Phrases: Coalescing Random Walks, Population Protocols, Voter Model

ACM Reference Format:
Varun Kanade, Frederik Mallmann-Trenn, and Thomas Sauerwald. 2022. On coalescence time in graphs–When is coalescing as fast as meeting?. 1, 1 (December 2022), 47 pages. https://doi.org/10.1145/nnnnnnn.nnnnnnn

Contents

An extended abstract based on this work appeared in SODA 2019 [32].

Authors' addresses: Varun Kanade, University of Oxford and The Alan Turing Institute, UK; Frederik Mallmann-Trenn, King's College London, UK; Thomas Sauerwald, University of Cambridge, UK.

Manuscript submitted to ACM

1 INTRODUCTION

Coalescing random walks is a fundamental distributed 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 and communication networks [4, 6, 24, 40, 42]). Other applications of the coalescence process appear in robotics [28]; 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 has annihilating particles [17].

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 (this duality assumes that the random walks are non-lazy). Despite recent progress by Cooper et al. [14, 19] and Berenbrink

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

et al. [11], 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 [8–10, 15, 16, 20, 26, 29]. However, the behavior of these processes is fundamentally different and despite their efficiency in reaching consensus on expanders and 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. [14] and Hassin and Peleg [30] 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 $t_{\mathrm{meet}} \leqslant t_{\mathrm{coal}}$; as for an upper bound, it can be shown that $t_{\mathrm{coal}} = O(t_{\mathrm{meet}} \log n)$, 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 [30].

Aldous [1] showed in *continuous-time* that the meeting time is bounded by the maximum hitting time, $t_{\mathrm{hit}} := \max_{u,v} t_{\mathrm{hit}}(u,v)$, where $t_{\mathrm{hit}}(u,v)$ 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 $O(t_{\mathrm{hit}} \log n)$ for the coalescing time; however, in general $O(t_{\mathrm{hit}})$ may be a loose upper bound on t_{meet}. In recent work, Cooper et al. [14] provide results that are better than $O(t_{\mathrm{meet}} \log n)$ for several interesting graph classes, notably expanders and power-law graphs. They show that $t_{\mathrm{coal}} = O((\log^4 n + \|\pi\|_2^{-2}) \cdot (1-\lambda_2)^{-1})$, 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. [11] show that $t_{\mathrm{coal}} = O(m/(d_{\min} \cdot \Phi))$, 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. [14] and Berenbrink et al. [11], 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 $t_{\mathrm{coal}}/t_{\mathrm{meet}}$ and $t_{\mathrm{meet}}/t_{\mathrm{mix}}$, where $t_{\mathrm{mix}} = t_{\mathrm{mix}}(1/e)$ denotes the mixing time.[2] In particular, the result shows that if $t_{\mathrm{meet}}/t_{\mathrm{mix}} = \Omega(\log^2 n)$, then $t_{\mathrm{coal}} = O(t_{\mathrm{meet}})$; 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 [2] 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 [30] 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.

[2]The *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.
[3]We call a graph *almost-regular* if $\deg(u) = \Theta(\deg(v))$ for all $u, v \in V$.
[4]Cooper et al. [14] 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	$\Theta(n)$	$\Theta(n \log n)$	[31] & [39]	$\Theta(n \log n)$	[31] & [39]	$\Theta(n \log n)$
Clique	$\Theta(1)$	$\Theta(n)$	[14], [11] & [31]	$\Theta(n)$	[14], [11] & *Thm.* 1.1	$\Theta(n)$
Cycle	$\Theta(n^2)$	$\Theta(n^2)$	[11] & *Prop.* 5.7	$\Theta(n^2)$	[11]	$\Theta(n^2)$
Rand. r-reg.	$\Theta(\log n)$	$\Theta(n)$	[19], [14], [11] & *Thm.* 1.1	$\Theta(n)$	[19], [14], [11] & *Thm.* 1.1	$\Theta(n)$
Hypercube	$\Theta(\log n \log \log n)$	$\Theta(n)$	*Thm.* 1.1	$\Theta(n)$	*Thm.* 1.1	$\Theta(n)$
Path	$\Theta(n^2)$	$\Theta(n^2)$	[11] & [31]	$\Theta(n^2)$	[11] & [31]	$\Theta(n^2)$
Star	$\Theta(1)$	$\Theta(1)$	folklore	$\Theta(\log n)$	[30], *Prop.* 3.4 & *Thm.* 1.4	$\Theta(n)$
Torus ($d = 2$)	$\Theta(n)$	$\Theta(n \log n)$	*Prop.* 5.7	$\Theta(n \log n)$	[39]	$\Theta(n \log n)$
Torus ($d > 2$)	$\Theta(n^{2/d})$	$\Theta(n)$	*Thm.* 1.1	$\Theta(n)$	*Thm.* 1.1	$\Theta(n)$

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 [2]. Note that for regular graphs, $t_{\mathrm{meet}} = \Omega(n)$ (e.g., Theorem 5.1.iii). The results for the hypercube and the torus with $d > 2$ are new (see Section 6). Note that $t_{\mathrm{coal}} = \Omega(t_{\mathrm{meet}})$.

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. In the continuous time setting, Cox [22] show that the coalescence time is bounded by $\Theta(t_{\mathrm{hit}})$ for tori. Oliveira [37] showed that the coalescence time is $O(t_{\mathrm{hit}})$ in general. In a different work, Oliveira [38] 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 [38, Theorem 1.2], implies that $t_{\mathrm{coal}} = O(t_{\mathrm{meet}})$ whenever $t_{\mathrm{mix}} \cdot \pi_{\mathrm{max}} = O(1/\log^4 n)$. One of our main results, Theorem 1.1, implies $t_{\mathrm{coal}} = O(t_{\mathrm{meet}})$ whenever $t_{\mathrm{mix}}/t_{\mathrm{meet}} = O(1/\log^2 n)$. Notice that since $t_{\mathrm{meet}} = \Omega(1/\|\pi\|_2^2) \geqslant \Omega(1/\pi_{\mathrm{max}})$, our condition is considerably more general—however, the results in [38] 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 $t_{\mathrm{coal}} \gg t_{\mathrm{meet}}$ such as the star graph, and together with Theorem 1.2, demonstrate that the trade-off between meeting and mixing time is the best possible.

1.1 Main results

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 $[t_{\mathrm{meet}}, O(t_{\mathrm{meet}} \cdot \log n)]$. 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.

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 $t_{\mathrm{coal}} = O(t_{\mathrm{meet}} \log n)$ is well-known. However, this bound can be improved substantially, as demonstrated by our result below (recall that for any graph G, $t_{\mathrm{mix}} = O(t_{\mathrm{meet}})$ [2]).

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

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

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

The proof of Theorem 1.1 appears in Section 3. One interesting aspect about this bound is that it can be used to establish $t_{\mathrm{coal}} = \Theta(t_{\mathrm{meet}})$ even without explicit bounds on the quantities t_{meet} or t_{mix}. Another interesting feature of our theorem is that the main result of Cooper et al. [14, Theorem 1] can be reproven by combining [14, Theorem 2] with Theorem 1.1 (see Proposition 5.2).

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 $t_{\mathrm{hit}} \gg t_{\mathrm{meet}}$, 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* $(\alpha_n)_{n \geqslant 0}$*,* $\alpha_n \in [1, \log^2 n]$ *there exists a family of almost-regular graphs* (G_n)*, with* G_n *having* $\Theta(n)$ *nodes and satisfying* $\frac{t_{\mathrm{meet}}}{t_{\mathrm{mix}}} = \Theta(\alpha_n)$ *such that*

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

The above two results show that the condition $t_{\mathrm{meet}} \geqslant t_{\mathrm{mix}} \log^2 n$ in Theorem 1.1 is best possible up to constant factors. A natural question is therefore whether in the case of structured sub-classes such as vertex-transitive graphs, or special graphs such as grids, tori, binary trees, cycles, real-world (power-law) graphs, *etc.*, better bounds can be obtained. 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.

In an earlier version [31], we also proved that

$$t_{\mathrm{coal}} = O\left(t_{\mathrm{hit}} + t_{\mathrm{meet}} \cdot \log(\Delta/d)\right) \tag{1}$$

for any graph G with maximum degree Δ and average degree d. We omit the proof of this result here, since very recently [39] proved that $t_{\mathrm{coal}} = O(t_{\mathrm{hit}})$ 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, which is the expected time until a random walk has visit all vertices in G (starting from a worst-case vertex). The detailed results are given in Section 5, and here we state only one of these results:

Theorem 1.3. *For any regular graph, it holds that*

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

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,* $t_{\mathrm{cov}} = O(n \log n/(1-\lambda_2)) = O(n \log n/\Phi)$*. Furthermore,*

$$t_{\mathrm{meet}} \leqslant t_{\mathrm{coal}} = O(n/\sqrt{1-\lambda_2}) = O(n/\Phi).$$

We point out that so far the best possible bound on t_{coal} for regular graphs has been $t_{\mathrm{coal}} = O(n/(1-\lambda_2))$ from [14].[6] The best possible bound on t_{hit} (and t_{cov}) in terms of $1-\lambda_2$, was $t_{\mathrm{hit}} = O(n/(1-\lambda_2))$ and $t_{\mathrm{cov}} = O(n \log n/(1-\lambda_2))$ due to Broder and Karlin [12] from 1989. In all four cases, t_{meet}, t_{coal}, t_{hit}, and t_{cov}, Theorem 1.3 improves the dependency

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

[6] Alternatively, the same bound as the known bound can also be derived from the bound on the conductance in [11] together with Cheeger's inequality.

on $1/(1-\lambda_2)$ (or, equivalently t_{mix}), by almost a square-root (we refer the reader to Theorem 5.4 and Theorem 5.6 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 [2, Corollary 6.2.1] by a factor of $1/\Phi$. After an extended abstract of our work was published in SODA'19 ([32]), Oliveira and Peres [39] improved versions of these bounds for non-regular graphs, by showing that the maximum degree Δ in the above statement can be replaced by the average degree d.

Finally, we also derive a general lower bound on t_{meet} that combines the lower trivial bound, $1/\|\pi\|_2^2$, with the minimum number of collisions (see Theorem 5.1.(iii)).

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 $O(n^2)$. These two new upper bounds for general and regular graphs complete the picture of worst-case bounds:

Theorem 1.4. *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 graph G we have $t_{\text{meet}}, t_{\text{coal}} \in [\Omega(n), O(n^2)]$.

The proof of Theorem 1.4 appears in Section 4.

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 [5, 25], where Efremenko and Reingold [25, Section 6] highlight the difficulty in analyzing the hitting time distribution before its expectation. In the context of concentration inequalities for Markov chains, Lezaud [34, 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 their expansion (see Czumaj and Sohler [23]). 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. our approaches can be leveraged to derive new bounds on other random walk quantities such as hitting times or cover times (see Section 5).

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 (expected) 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 $\mathcal{G}_1$ cannot be eliminated. The domination of the real process by the group-restricted one is established by introducing a

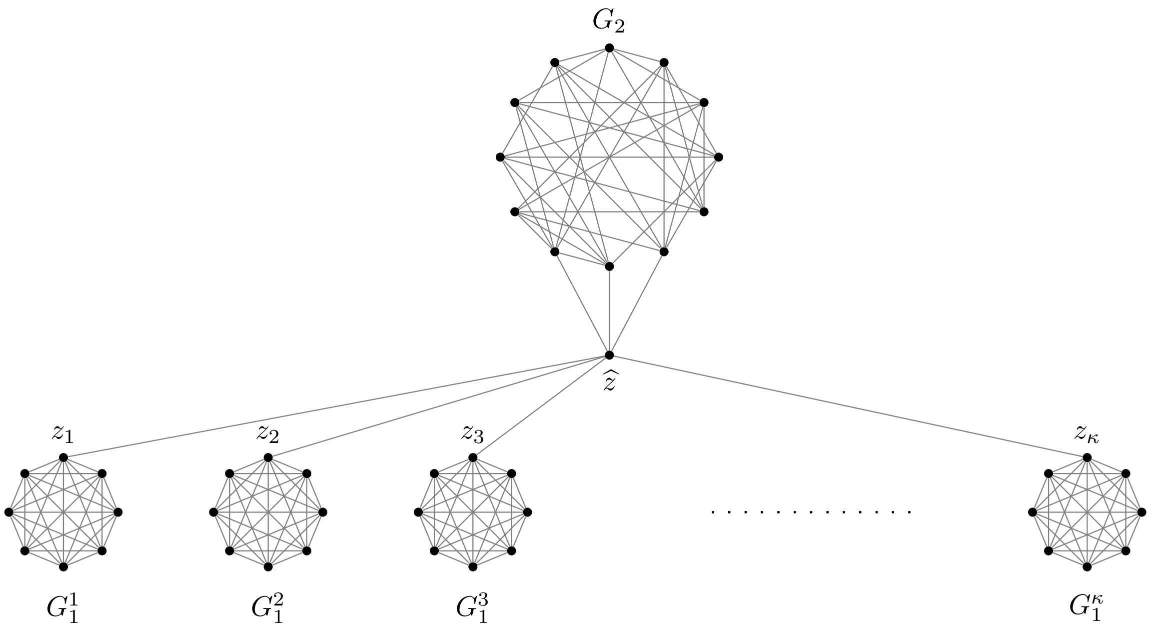

Fig. 1. The graph described in Section 3.4 with $t_{\text{coal}} = \Omega(t_{\text{meet}} + \sqrt{t_{\text{meet}}/t_{\text{mix}}} \cdot \log n \cdot t_{\text{mix}})$.

formal concept called immortal process at the beginning of Section 3.1. We believe that this domination might be of further use to analyse other interacting particle processes.

In the group-restricted 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 11). 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 (G_2 in Figure 4). 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 $t_{\text{meet}}/t_{\text{mix}} \geqslant \log^2 n$. Therefore, the problem of bounding t_{coal} reduces to bounding t_{meet}.

For some of the other results, including Theorem 1.2, 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 $(X_t)_{t\geqslant 0}$ and $(Y_t)_{t\geqslant 0}$, 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

$$\mathbb{P}[\, Z \geqslant 1\,] = \frac{\mathbb{E}[\, Z\,]}{\mathbb{E}[\, Z \mid Z \geqslant 1\,]}. \tag{2}$$

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

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

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.,* [18]) to derive accurate bounds on the hitting (and cover time) on various classes of random graphs, and in Barlow et al. [7] to bound the collisions of random walks on infinite graphs. Using (2), 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 (2) is that in order to lower (or upper) bound the probability that the two walks meet, we need to derive a corresponding bound on $\mathbb{E}[\, Z \mid Z \geqslant 1\,]$, *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:

$$\mathbb{E}[\, Z \mid Z \geqslant 1\,] \leqslant \max_{u\in V} \sum_{t=0}^{\tau-1} \sum_{v\in V} \left(p_{u,v}^t\right)^2. \tag{3}$$

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 $Z \geqslant 1$) at some vertex u before time-step τ yields the bound in (3). 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 Section 5, or the last paragraph in this section for a summary).

Worst-Case Bounds. We proceed by establishing that $t_{\mathrm{coal}} = O(n^3)$ on all graphs. The proof of $t_{\mathrm{coal}} = O(n^3)$ (Theorem 1.4) follows by reducing the number of walks from n to $(\Delta/d)^{O(1)} \leqslant (n^2/|E|)^{O(1)}$ in $O(t_{\mathrm{hit}})$. We have, by Proposition 5.7, $t_{\mathrm{meet}} \leqslant 2t_{\mathrm{hit}} + 1 = O(n \cdot |E|)$, where this last bound follows from [3].

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

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

steps the coalescing terminates. The fact that this is tight can be easily verified by considering the Barbell graph.[7]

For regular graphs, the same argument as before shows that $t_{\mathrm{coal}} = O(n^2)$, and this is matched by the cycle, for instance.

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

Bounds on t_{meet} *and Other Results.* In Section 5, we derive several bounds on t_{meet}. These bounds are derived more directly by (2) and/or (3), 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. [14] 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 [12].

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 $N(u) = \{v\colon (u, v) \in E\}$ 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 $p^t_{u,v}$ 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 $p^t_{u,\cdot}$ 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 $d(t) = \max_u \|p^t_{u,\cdot} - \pi\|_{\text{TV}}$ and $\bar{d}(t) = \max_{u,v}\|p^t_{u,\cdot} - p^t_{v,\cdot}\|_{\text{TV}}$, where $\|\cdot\|_{\text{TV}}$ denotes the total variation distance, which can be defined for any two discrete distributions μ and ν as $\|\mu - \nu\|_{\text{TV}} = \frac{1}{2}\sum_{\omega\in\Omega}|\mu(\omega) - \nu(\omega)|$. Following Aldous and Fill [2], we define the *mixing time* to be $t_{\text{mix}}(\varepsilon) = \min\{t \geqslant 0 : \bar{d}(t) \leqslant \varepsilon\}$ and for convenience we will write $t_{\text{mix}} = t_{\text{mix}}(1/e)$. We define separation from stationarity to be $s(t) = \min\{\varepsilon\colon p^t_{u,v} \geqslant (1-\varepsilon)\pi(v) \text{ for all } u, v \in V\}$. Then $s(\cdot)$ is submultiplicative, i.e., $s(t_1 + t_2) \leqslant s(t_1)\cdot s(t_2)$ [33, Exercise 6.4], so in particular, non-increasing [33, Exercise 6.4]. Hence we can define the *separation threshold time* $t_{\text{sep}} = \min\{t \geqslant 0 : s(t) \leqslant 1/e\}$ and, by [2, Lemma 4.11] $t_{\text{sep}} \leqslant 4t_{\text{mix}}$. We write $T_{\text{hit}}(u, v)$ to denote the first time-step $t \geqslant 0$ 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 $t_{\text{avg-hit}} = \sum_{u,v\in V}\pi(u)\cdot\pi(v)\cdot t_{\text{hit}}(u, v)$.

Finally, let $T_{\text{cov}}(u)$ be the first time until a random walk starting at u has visited all vertices in G. Then the cover time of G, denoted by $t_{\text{cov}}(G)$, is defined as $t_{\text{cov}}(G) := \max_{u\in V}\mathbb{E}[\,T_{\text{cov}}(u)\,]$.

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^{π}_{meet} 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.,*

$$\mathbb{P}[\, S_{t+1} \mid \mathcal{F}_t \,] = \mathbb{P}[\, S_{t+1} \mid S_t \,], \tag{4}$$

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 $T_{\text{coal}}(S_0) = \min\{t \geqslant 0 \mid |S_t| = 1\}$. 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 Auxiliary 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 [37, 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 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 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 id $i \in \mathbb{N}$ which is on node $v \in V$ to be

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

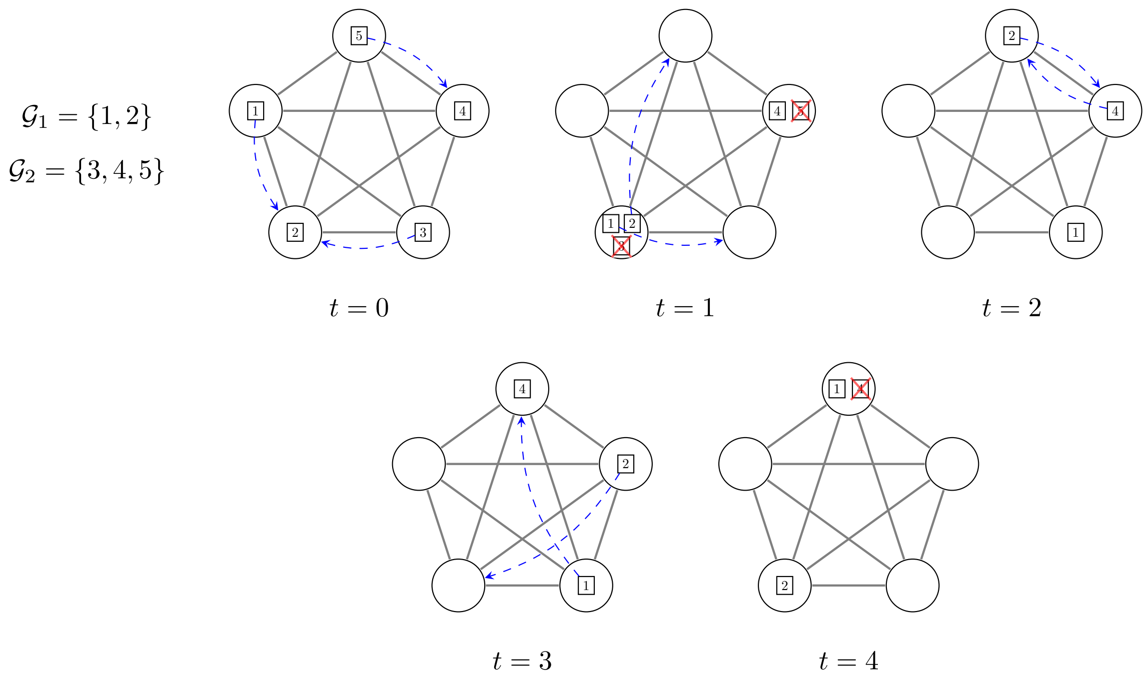

Fig. 2. Illustration of the process P_{imm}.

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.

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^k_{imm} 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

$$T^k \leqslant T^k_{\text{imm}}.$$

Note that this is not obvious since it could happen that a random walk of $\mathcal{G}_1$, that dies in the original process, is able to eliminate many of the other random walks hence leading to a small T^k_{imm} (in comparison to T^k). In particular, the standard coupling, in which the random choices P and P_{imm} are exactly the same, fails. Instead, we use two non-trivial couplings.

Proof of Proposition 3.1. 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^{Q}_t denote the set of tuples of alive walks in process $Q \in \{P, P_{\text{int}}, P_{\text{imm}}\}$ at time t. Let $\bar{S}^{Q}_t = \{v \mid (v,i) \in S^{Q}_t\}$ 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 $t \geqslant 1$). In order to prove the proposition, we show that there exists a coupling, such that for any $t \in \mathbb{N}$

$$\bar{S}^{\mathrm{P}}_t \subseteq \bar{S}^{P_{\text{int}}}_t \tag{5}$$

$$\bar{S}^{P_{\text{int}}}_t \subseteq \bar{S}^{P_{\text{imm}}}_t \tag{6}$$

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

$$T^k = \min\{t \geqslant 0 \colon |\bar{S}^{P}_t| \leqslant k\} \leqslant \min\{t \geqslant 0 \colon |\bar{S}^{P_{\text{imm}}}_t| \leqslant k\} = T^k_{\text{imm}}.$$

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 $\mathrm{id}(r) \in \mathcal{G}_i$ for $i \in \{1, 2\}$. Whenever at least one walk arrives[8] 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

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

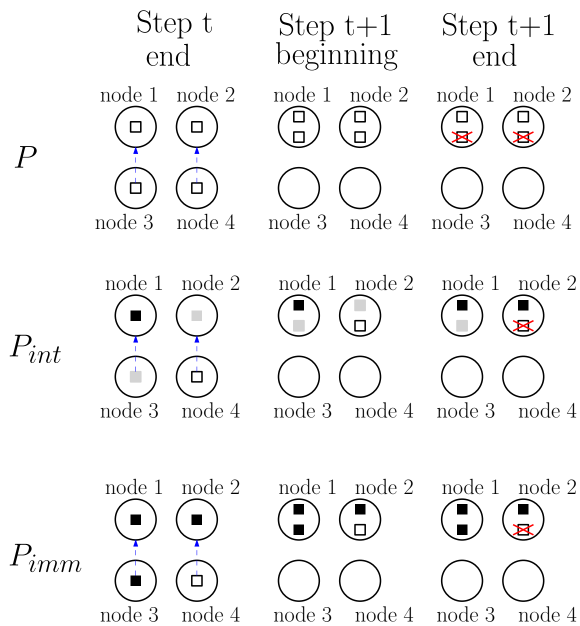

Fig. 3. An illustration of couplings between the processes. The squares depict the random walks. Walks of $\mathcal{G}_1$ are colored black and grey (phantom) and the nodes of $\mathcal{G}_2$ are white. The blue arrows denote the moving decisions. Observe that in P_{int} a phantom becomes alive (and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

   (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.

Couplings. We couple all decisions by the random walks in P_{int} with the decisions of the random walks P_{imm}. The coupling between the random walks' decisions in P with those of P_{int} in the following way. For the (unique) random walk at $v \in \bar{S}_t^{\mathrm{P}}$ 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 $\bar{S}_t^{P_{\mathrm{int}}}$ (there might be several walks of P_{int}, however only one is alive and we couple with this alive walk).[9]

Proof of (5) . We prove (5) by induction on t starting from the same initial configuration: if $v \in \bar{S}_t^{\mathrm{P}}$, then $v \in \bar{S}_t^{P_{\mathrm{int}}}$. 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 $v \in \bar{S}_t^{\mathrm{P}}$ under process P, we coupled its transition to node $Y_v(t+1)$ (where we possibly have $Y_v(t+1) = v$) with the corresponding alive walk of $\bar{S}_t^{P_{\mathrm{int}}}$ (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.,* $S = \{Y_v(t+1) \colon v \in \bar{S}_t^{\mathrm{P}}\}$. 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 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 least[10] one alive walk remains. Thus, our coupling ensures that if $v \in \bar{S}_{t+1}^{\mathrm{P}}$, then $v \in \bar{S}_{t+1}^{P_{\mathrm{int}}}$. 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 (5) and we turn to proving (6).

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 $t+1$.

We now prove the second invariant. Note that whenever a walk r of type $\mathcal{G}_2$ in $P_{\mathrm{imm}}(P_{\mathrm{int}})$ 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 $P_{\mathrm{int}}(P_{\mathrm{imm}})$. We have that either r' is of type $\mathcal{G}_1$ or r' is of type $\mathcal{G}_2$ and that $\mathrm{id}(r') < \mathrm{id}(r)$. In either case, r is also killed (labeled dead) in $P_{\mathrm{int}}(P_{\mathrm{imm}})$. 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 (6). This finishes the proof. □

3.2 Meeting Time Distribution Prior to t_{meet}

Let $(X_t)_{t\geqslant 0}$ and $(Y_t)_{t\geqslant 0}$ 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.

[9] We will show that such a walk always exists in P_{int}, nonetheless, for completeness we can define a suitable random decision should no such walk in P_{int} be available.
[10] By definition, there is actually exactly one alive walk.

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

$$\mathbb{P}\left[\,\mathrm{intersect}(X_t, Y_t, 5t_{\mathrm{mix}})\,\right] \geqslant \frac{1}{32\alpha},$$

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

$$\mathbb{P}\left[\,\mathrm{intersect}(X_t, Y_t, c\, b\, t_{\mathrm{mix}})\,\right] \geqslant \frac{b}{\alpha},$$

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

$$\mathbb{P}\left[\,\mathrm{intersect}(\widetilde{X}_t, \widetilde{Y}_t, \ell)\,\right] \geqslant \frac{1}{2}. \tag{7}$$

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

$$\mathbb{P}\left[\,\mathrm{intersect}(\widetilde{X}_t, \widetilde{Y}_t, \ell)\,\right] \leqslant 2\lceil\alpha\rceil \cdot \mathbb{P}\left[\,\mathrm{intersect}(\widetilde{X}_t, \widetilde{Y}_t, t_{\mathrm{mix}})\,\right], \tag{8}$$

and hence combining equation (7) and (8) yields

$$\mathbb{P}\left[\,\mathrm{intersect}(\widetilde{X}_t, \widetilde{Y}_t, t_{\mathrm{mix}})\,\right] \geqslant \frac{1}{4\lceil\alpha\rceil}.$$

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

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

Observing that for any $\alpha \geqslant 1$, 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 $5t_{\mathrm{mix}}$. 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 $(1-x)^k \leqslant e^{-xk} \leqslant 1-(1-e^{-1})xk$ for $0 \leqslant x < 1$, $k \geqslant 0$ and $xk \leqslant 1$. We derive that the probability of intersecting after k blocks is at least $1-(1-\frac{1}{32\alpha})^k \geqslant (1-e^{-1})k/(32\alpha) = b/\alpha$.

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

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 two independent random walks starting from the stationary distribution* [11]. *Let $p = \mathbb{P}\left[\,\mathrm{intersect}(X_t, Y_t, \tau)\,\right]$ and let $\mathcal{T}_\tau$ denote the set all possible trajectories of a walk of*

[11] The lemma also holds for arbitrary distributions, but for the sake of simplicity we only state this slightly less general version.

length τ in G (including possible self-loops). We define the following two categories C_1 and C_2 with $C_1 \subseteq C_2$

$$C_1 := \{(z_0, \ldots, z_\tau) \in \mathcal{T}_\tau : \mathbb{P}\left[\, \exists 0 \leqslant s \leqslant \tau, Y_s = z_s \,\right] \geqslant \sqrt{p}\}$$
$$C_2 := \{(z_0, \ldots, z_\tau) \in \mathcal{T}_\tau : \mathbb{P}\left[\, \exists 0 \leqslant s \leqslant \tau, Y_s = z_s \,\right] \geqslant p/3\}.$$

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

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

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

a contradiction. This completes the proof. □

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 [30]. For the sake of completeness, we give a self-contained proof.

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 $T := \min\{t \geqslant t_0 : |S_t| \leqslant 99/100 \cdot |S_{t_0}|, |S_{t_0}| \geqslant 100\}$ denoting the first time-step the number of coalescing random walks reduces by a factor of $99/100$, we will prove that $\mathbb{E}\left[\, T \,\right] = O(t_{\text{meet}})$. 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.,*

$$\mathbb{P}\left[\text{intersect}(X_t, Y_t, c\frac{e-1}{e} t_{\text{meet}})\right] \geqslant \frac{e-1}{e}.$$

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

We condition on the successful coupling of $X_{4t_{\text{mix}}}$ with a node drawn from π and that $(X_t)_{t \geqslant 4t_{\text{mix}}} \in C_2$, 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 $(Y_t)_{t \geqslant 0} b$ 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, $(Y_t)_{t \geqslant 4t_{\text{mix}}}$ meets $(X_t)_{t \geqslant 4t_{\text{mix}}}$ 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.

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

The above holds as long as $Z_{\ell-1} \geqslant 100$. 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. □

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 $1, \ldots, k'$ and starting nodes $v_1, \ldots, v_{k'}$. Thus, $S_0 = \{(v_i, i) \mid 1 \leqslant i \leqslant k'\}$.
- Let $\mathcal{G}_1 = \{1, \ldots, k\}$ and $\mathcal{G}_2 = \{k+1, \ldots, k'\}$. 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 in 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

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

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 $g \geqslant k$: 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 $p_\tau \leqslant \mathbb{P}\left[\text{intersect}(X_t, Y_t, \tau)\right]$ 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 $k \geqslant \frac{3}{(1-e^{-1})\cdot p_\tau}$. For some $1 \leqslant g \leqslant |S_0| - k$, define the stopping condition $T_g = \min\{t \geqslant 0 \mid |\text{IDs}(S_t) \cap \mathcal{G}_2| \leqslant g\}$. Then the expected stopping time satisfies*

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

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 $4t_{\text{mix}} \geqslant t_{\text{sep}}$ w.c.p. with a node drawn from π. Thus, conditioning on the success of this coupling we have, by Lemma 3.3, $\mathbb{P}\left[(Z_t)_{t=4t_{\text{mix}}}^{4t_{\text{mix}}+\tau} \in C_1\right] \geqslant \frac{p_\tau}{3}$ or $\mathbb{P}\left[(Z_t)_{t=4t_{\text{mix}}}^{4t_{\text{mix}}+\tau} \in C_2\right] \geqslant \frac{\sqrt{p_\tau}}{3}$. 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 walks 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. a 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 $Z_j = |\text{IDs}(S_{j\cdot(4t_{\text{mix}}+\tau)}) \cap \mathcal{G}_2|$, *i.e.*, the number of walks remaining in $\mathcal{G}_2$ after j blocks of time have passed. For any $j \geqslant 1$, we will show that there exists a constant $c > 0$ such that,

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

By using Lemma A.9, we get $\mathbb{E}\left[T_g\right] = O\left((4t_{\text{mix}} + \tau) \cdot \frac{1}{\sqrt{p_\tau}} \cdot (\log|\mathcal{G}_2| - \log g)\right)$ (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 jth block. In the remainder we show that we have indeed $\mathbb{E}\left[Z_j \mid \mathcal{F}_{j-1}\right] \leqslant Z_{j-1} \cdot \left(1 - c\sqrt{p_\tau}\right)$.

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

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 $k \geqslant \frac{3}{p_\tau\cdot(1-e^{-1})}$, 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 $P_{\mathrm{imm}}(S_0, k)$; 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 $4t_{\mathrm{mix}}$ 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 $4t_{\mathrm{mix}}$ and $4t_{\mathrm{mix}} + \tau$ 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

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

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

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 $4t_{\mathrm{mix}}$ 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 $4t_{\mathrm{mix}}$ to $4t_{\mathrm{mix}} + \tau$ 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 $4t_{\mathrm{mix}}$ and $4t_{\mathrm{mix}} + \tau$ 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 $4t_{\mathrm{mix}}$ with a node drawn from π. Conditioned on this between time-steps $4t_{\mathrm{mix}}$ to $4t_{\mathrm{mix}} + \tau$, this random walk hits any trajectory whose portion between time-steps $4t_{\mathrm{mix}}$ to $4t_{\mathrm{mix}} + \tau$ lies in C_2 with probability at least $p_\tau/3$ (by definition of C_2 in Lemma 3.3). Since $k = |\mathcal{G}_1| \geqslant \frac{3}{(1-e^{-1})\cdot p_\tau}$, 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 $4t_{\mathrm{mix}}$ to $4t_{\mathrm{mix}} + \tau$ 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

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

□

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 processP_{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 $\alpha = t_{\mathrm{meet}}/t_{\mathrm{mix}}$. Let $T_1 = \min\{t \geqslant 0 \mid |S_t| \leqslant 2\lceil\alpha\rceil\}$. Then $\mathbb{E}\left[\, T_1 \,\right] = O(t_{\mathrm{mix}} \cdot \sqrt{\alpha} \cdot \log|S_0|)$.*

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 $\{1, 2, \ldots, |S_0|\}$. Let $S_0 = \{(v_1, 1), \ldots, (v_{|S_0|}, |S_0|)\}$ and $\bar{S}_0 = \{i\colon (v, i) \in S_0\}$. We consider the process $P_{\mathrm{imm}}(S_0, k)$ and $k = \lceil\alpha\rceil$. Let T_1^* be the stopping time defined by $|\mathrm{IDs}(\bar{S}_t) \cap \mathcal{G}_2| \leqslant \alpha$ for the process $P_{\mathrm{imm}}(S_0, k)$. By definition of P_{imm} and Proposition 3.1, it follows that T_{imm} stochastically dominates T. Thus, it suffices to bound $\mathbb{E}\left[\, T_{\mathrm{imm}} \,\right]$. W.l.o.g. we assume that $\alpha \geqslant 6\frac{e-1}{e}$, 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,

$$p = \mathbb{P}\left[\, \mathrm{intersect}(X_{t \geqslant 0}, Y_{t \geqslant 0}, 6ct_{\mathrm{mix}}) \,\right] \geqslant \frac{6}{\alpha},$$

Thus, we have

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

Applying Claim 3.5 with $g = \alpha$, $\tau = 6ct_{\mathrm{mix}}$ (where c is a constant as given by Lemma 3.2), $p_\tau = 6/\alpha$, and observing that $k \geqslant \frac{3}{(1-e^{-1}) \cdot p_\tau}$, 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 $O(t_{\mathrm{meet}})$.

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

Proof. We will consider the coalescence process in phases. Let ℓ be the largest integer such that $|S_0| \geqslant \left(\frac{4}{3}\right)^{\ell}$. For $j \geqslant 1$, the j^{th} phase ends when $|S_t| < \left(\frac{4}{3}\right)^{\ell - j + 1}$. The $(j+1)^{th}$ 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 $\ell - j + 1 \geqslant 32$.

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

$$\left(\frac{4}{3}\right)^{\ell - j + 1} \leqslant |S_{t_j}| < \left(\frac{4}{3}\right)^{\ell - j + 2} \tag{9}$$

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 $S_0' = \{(v_1, 1), \ldots, (v_{n_j}, n_j)\}$ and $\bar{S}_0' = \{v_1, \ldots, v_{n_j}\}$. 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

$$k_j := \left\lceil \frac{|S_0'|}{2} \right\rceil$$

be the size of $\mathcal{G}_1$ and consider the process $P_{\mathrm{imm}}(S_0', k_j)$ as defined in Section 3.3. Let

$$g_j := \left\lfloor \frac{|S_0'| - k_j}{3} \right\rfloor$$

and

$$T_2^*(j) := \min\{t \mid |\mathrm{IDs}(S_t') \cap \mathcal{G}_2| \leqslant g_j\}.$$

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

$$|S_t'| \leqslant g_j + k_j \leqslant \frac{|S_0'| - k_j}{3} + k_j = \frac{|S_0'|}{3} + \frac{2k_j}{3} \leqslant \frac{|S_0'|}{3} + \frac{|S_0'|}{3} + \frac{2}{3} < \frac{3}{4} \cdot |S_0'|.$$

By Proposition 3.1, $T_2^*(j)$ stochastically dominates $T_2(j)$ and hence it suffices to bound $\mathbb{E}\left[\, T_2^*(j)\,\right]$. In order to bound $\mathbb{E}\left[\, T_2^*(j)\,\right]$, 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 $\ell-j+1 \geqslant 32$ we have $b_j \leqslant b_{\ell-31} \leqslant ((e-1)/e)\alpha$. Furthermore, we have $b_j \geqslant b_0 \geqslant 4\alpha\log\alpha(3/4)^{\ell} \geqslant 1$, where the last inequality follows from $(4/3)^{\ell} \leqslant |S_0| \leqslant 4\alpha\log\alpha$, 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

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

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

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

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

$$\frac{8}{p_j} = \frac{8}{32\log(4/3)(\ell-j+1)}(4/3)^{\ell-j+1} \leqslant \frac{1}{4}\cdot\left(\frac{4}{3}\right)^{\ell-j+1} \leqslant \frac{1}{2}\cdot|S_0'| \leqslant k_j,$$

where we used (9) and $|S_0'| = |S_{t_j}|$ in the second-last inequality. Thus we can apply Claim 3.5 and derive

$$\mathbb{E}\left[\, T_2^*(j)\,\right] \leqslant (\tau_j + 4t_{\mathrm{mix}})\cdot\frac{1}{\sqrt{p_j}}\cdot\left(\log|\mathrm{IDs}(S_0')\cap\mathcal{G}_2| - \log g_j\right)$$

and we continue by dissecting that bound. Since $b_j \geqslant 1$, there exists a suitably large constant c_1, so that $\tau_j + 4t_{\mathrm{mix}} \leqslant c_1 b_j t_{\mathrm{mix}}$. 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, $|\mathrm{IDs}(S_0')\cap\mathcal{G}_2|/g_j \leqslant 3$, hence $\log|\mathrm{IDs}(S_0')\cap\mathcal{G}_2| - \log g_j \leqslant \log(3)$. Putting everything together, we get that there is a constant c_2 such that,

$$\mathbb{E}\left[\, T_2^*(j)\,\right] \leqslant c_2\cdot t_{\mathrm{mix}}\cdot\alpha\cdot\sqrt{\ell-j+1}\left(\frac{3}{4}\right)^{(\ell-j+1)/2} \tag{10}$$

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

$$\begin{aligned}
\mathbb{E}\left[\, T_2\,\right] &= \mathbb{E}\left[\,\widetilde{T}\,\right] + \mathbb{E}\left[\,\mathbb{E}\left[\, T_2 - \widetilde{T} \mid \widetilde{T}\,\right]\,\right] \\
&\leqslant c_2\cdot t_{\mathrm{mix}}\cdot\alpha\sum_{j=0}^{\ell-30}\sqrt{\ell-j+1}\left(\frac{3}{4}\right)^{(\ell-j+1)/2} + c_3 t_{\mathrm{meet}} \qquad (11)\\
&\leqslant c_2\cdot t_{\mathrm{mix}}\cdot\alpha + c_3 t_{\mathrm{meet}} = O(t_{\mathrm{meet}}) \qquad (12)
\end{aligned}$$

Above, in (11) we used (10) and the fact that $\mathbb{E}\left[\, T_2 - \widetilde{T} \mid \widetilde{T}\,\right] \leqslant c_3 t_{\mathrm{meet}}$ for some constant $c_3 > 0$ and in step (12), we used the fact that $\sum_{j=32}^{\infty} jc^j < 1$ for $c \leqslant \sqrt{3/4}$. □

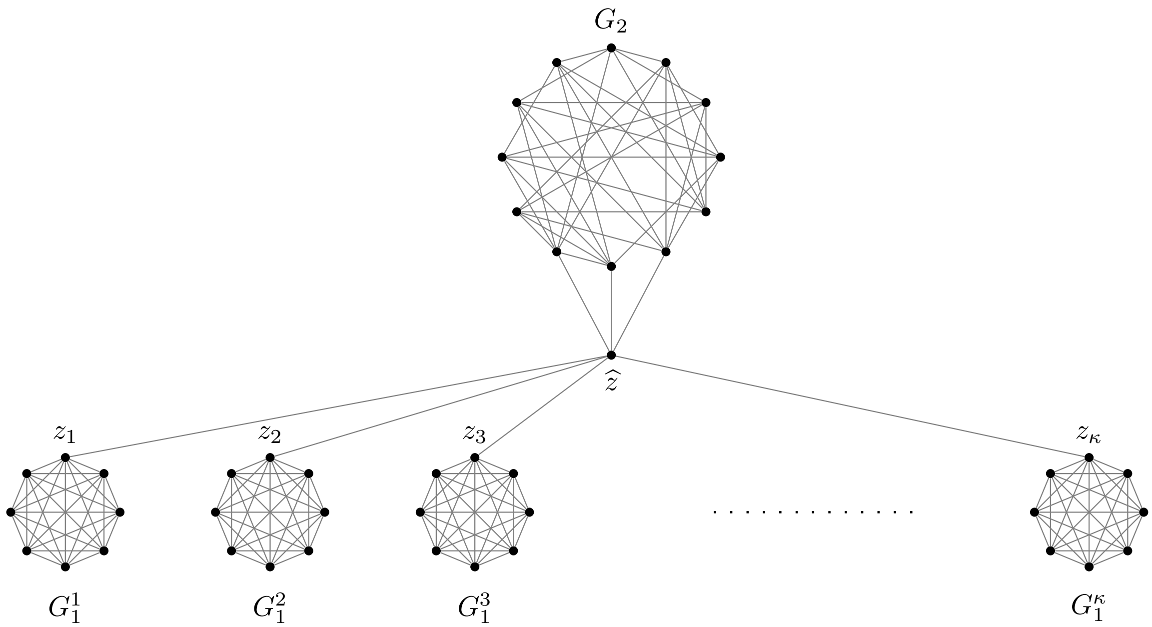

Fig. 4. The graph with $t_{\text{coal}} = \Omega(t_{\text{meet}} + \sqrt{t_{\text{meet}}/t_{\text{mix}}} \cdot \log n \cdot t_{\text{mix}})$.

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 [2, 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 an Erdős-Rényi random graph $G(n/\sqrt{\alpha'}, \sqrt{\alpha'}/\sqrt{n})$ (the nodes having expected degree $\sqrt{n}$), 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 $G_1^1, G_1^2, \ldots, G_1^\kappa$), and a node $\widehat{z}$, 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 – in fact, the maximum and minimum degree differ by a factor of at most $(1 + o(1))$, since by applying a Chernoff bound, it follows that with high probability that all degrees of the Erdős-Rényi random graph are within $\sqrt{n} \pm O(n^{1/4}\sqrt{\log n})$.

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 auxiliary lemma which shows that the walk restricted to V_2 behaves similarly to the walk restricted to $V_2 \cup \{\widehat{z}\}$, 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 $\widehat{Q}$ be the transition matrix of the random walk on the subgraph of G induced by $V_2 \cup \{\widehat{z}\}$. Let $p^t_{u,v}, q^t_{u,v}, \widehat{q}^t_{u,v}$ denote the corresponding transition probabilities for a walk starting at u to end up at node v after t steps. Let $S^* = \{u \in V_2 \cap N(\widehat{z})\}$. Then the following statements hold:*

(i) For any $u, v \in V_2$ we have $\|p^t_{u,\cdot} - q^t_{u,\cdot}\|_{\text{TV}} \leqslant (1 + o(1)) \sum_{i=1}^{t-1} p^i_{u,S^*}/(\sqrt{n}) \leqslant (1 + o(1))t/(2\sqrt{n})$.
(ii) For any $u, v \in V_2$ we have $\|\widehat{q}^t_{u,\cdot} - q^t_{u,\cdot}\|_{\text{TV}} \leqslant (1 + o(1)) \sum_{i=1}^{t-1} p^i_{u,S^*}/(2\sqrt{n}) \leqslant (1 + o(1))t/(2\sqrt{n})$.
(iii) For any $u, v \in V_2$ we have that after $t = t_{\text{mix}}(G_2)$ time steps $\|p^t_{u,\cdot} - p^t_{v,\cdot}\|_{\text{TV}} \leqslant o(1) + 2/e$.

Proof. Let $(X_t)_{t \geqslant 0}$ be the Markov chain with transition matrix P and let $(Y_t)_{t \geqslant 0}$ 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

$$\begin{aligned}\mathbb{P}\left[X_t \neq Y_t \mid X_{t-1} = Y_{t-1} \right] &= \mathbb{P}\left[X_t \neq Y_t \mid X_{t-1} = Y_{t-1}, X_{t-1} \in S^* \right] \cdot \mathbb{P}\left[X_{t-1} \in S^* \right] \\ &\quad + \mathbb{P}\left[X_t \neq Y_t \mid X_{t-1} = Y_{t-1}, X_{t-1} \in V_2 \setminus S^* \right] \cdot \mathbb{P}\left[X_{t-1} \in V_2 \setminus S^* \right] \\ &\leqslant \mathbb{P}\left[X_t = \widehat{z} \mid X_{t-1} \in S^* \right] p^{t-1}_{u,S^*} + 0 \leqslant (1 + o(1)) p^{t-1}_{u,S^*}/(2\sqrt{n}),\end{aligned}$$

where we used the laziness of the process and the fact that the degrees are in $(1 \pm o(1))\sqrt{n}$ for all nodes in S_2 by a Chernoff bound. We have, by [33, Proposition 4.7],

$$\|p^t_{u,\cdot} - p^t_{v,\cdot}\|_{\text{TV}} = \inf\{\mathbb{P}\left[X \neq Y \right] \mid (X, Y) \text{ is a coupling of } p^t_{u,\cdot} \text{ and } p^t_{v,\cdot}\}.$$

Hence, by a union bound over t steps,

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

To prove the second part we redefine $(X_t)_{t \geqslant 0}$ 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 Corollary A.3, we get

$$\begin{aligned}\|p^t_{u,\cdot} - \pi^{G_2}(\cdot)\|_{\text{TV}} &\leqslant \|p^t_{u,\cdot} - q^t_{u,\cdot}\|_{\text{TV}} + \|q^t_{u,\cdot} - \pi^{G_2}(\cdot)\|_{\text{TV}} \\ &\leqslant (1 + o(1)) \frac{t_{\text{mix}}(G_2)}{\sqrt{n}} + \|q^t_{u,\cdot} - \pi^{G_2}(\cdot)\|_{\text{TV}} \\ &\leqslant o(1) + \|q^t_{u,\cdot} - \pi^{G_2}(\cdot)\|_{\text{TV}} \leqslant o(1) + 1/e,\end{aligned}$$

where the last inequality follows form the definition of mixing time. Again, by the triangle inequality, $\|p^t_{u,\cdot} - p^t_{v,\cdot}\|_{\text{TV}} \leqslant o(1) + 2/e$. ◻

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 $c \geqslant 1$ such that $\mathbb{P}[\, T_{\text{hit}}(u, \widehat{z}) \geqslant n/c \,] \geqslant 1/2$.
(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 Corollary A.3.
Fix an arbitrary $t \in [2t_{\text{mix}}(G_2), 2t_{\text{mix}}(G_2) + 7]$. In the following we show $\|\widehat{q}^t_{u,\cdot} - \pi^{\widehat{G}}(\cdot)\|_{\text{TV}} \leqslant 1/e$. 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.*,

$$\mathcal{D} = \{D' \colon \text{ for } u \sim D' \text{ we have } \mathbb{P}[\, u = \widehat{z} \,] = 0\}. \tag{13}$$

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

$$\|\widehat{q}^t_{u\sim D',\cdot} - \pi^{\widehat{G}}(\cdot)\|_{\text{TV}} = 0 + \frac{1}{2}\sum_{v\in V_2}\left|\widehat{q}^t_{u\sim D',v} - \pi^{\widehat{G}}(v)\right| + \frac{1}{2}\left|\widehat{q}^t_{u\sim D',\widehat{z}} - \pi^{\widehat{G}}(\widehat{z})\right|.$$

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 [33, 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

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

Thus, for $\delta_v := |\widehat{q}^t_{u,v} - q^t_{u,v}|$, we get by using triangle inequality,

$$\begin{aligned}
\frac{1}{2}\sum_{v\in V_2}\left|\widehat{q}^t_{u\sim D',v} - \pi^{\widehat{G}}(v)\right| &\leqslant \frac{1}{2}\sum_{v\in V_2}\left|\widehat{q}^t_{u\sim D',v} - \pi^{G_2}(v)\right| + \frac{1}{2}\sum_{v\in V_2}|\pi^{G_2}(v) - \pi^{\widehat{G}}(v)| \\
&\leqslant \frac{1}{2}\sum_{v\in V_2}\left|\widehat{q}^t_{u\sim D',v} - \pi^{G_2}(v)\right| + \frac{1}{2}\sum_{v\in V_2}\pi^{G_2}(v)\zeta \\
&\leqslant \frac{1}{2}\sum_{v\in V_2}\left|q^t_{u\sim D',v} - \pi^{G_2}(v)\right| + \frac{1}{2}\sum_{v\in V_2}|\delta_v| + \frac{1}{2}\sum_{v\in V_2}\pi^{G_2}(v)|\zeta| \\
&\leqslant d^Q(t) + 1/32 + \frac{\zeta}{2}, \\
&\leqslant d^Q(t) + 1/32 + 1/32,
\end{aligned} \tag{14}$$

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

The above equation (14) only consider the variation distance w.r.t. V_2. For $\widehat{z}$ we have

$$\frac{1}{2}|\widehat{q}^t_{u\sim D',\widehat{z}} - \pi^{\widehat{G}}(\widehat{z})| \leqslant (2t_{\text{mix}}(G_2)+7)2(1+o(1))/\sqrt{n} \leqslant 1/32.$$

Putting everything together we get we get

$$\|\widehat{q}^t_{u\sim D',\cdot} - \pi^{\widehat{G}}(\cdot)\|_{\text{TV}} = \frac{1}{2}\sum_{v\in V_2}\left|\widehat{q}^t_{u\sim D',v} - \pi^{\widehat{G}}(v)\right| + \frac{1}{2}\left|\widehat{q}^t_{u\sim D',\widehat{z}} - \pi^{\widehat{G}}(\widehat{z})\right|$$
$$\leqslant d^Q(t) + 1/32 + 1/32 + 1/32 \leqslant 1/e^2 + 3/32 \tag{15}$$

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

$$\widehat{q}^7_{\widehat{z},\widehat{z}} \leqslant \frac{1}{2^7} + \sum_{i\leqslant 7}\sum_{v\in N(\widehat{z})}\widehat{q}^{i-1}_{\widehat{z},v}\cdot\frac{4}{\sqrt{n}} \leqslant \frac{1}{2^7} + \frac{1}{\sqrt{\alpha'}} \leqslant 1/32.$$

The set of distribution for the position of the random walk at time 7 conditioning on $X_7 \neq \widehat{z}$ gives the same distribution $\mathcal{D}$ as defined in (13). Let $D_{\widehat{z}} \in \mathcal{D}$ be distribution of the random at time 7 starting at $\widehat{z}$. Hence, by (15), we get

$$\|\widehat{q}^{2t_{\text{mix}}(G_2)+7}_{\widehat{z},\cdot} - \pi^{\widehat{G}}(\cdot)\|_{\text{TV}} \leqslant \widehat{q}_{\widehat{z},V(\widehat{G})\setminus\{\widehat{z}\}}\cdot\|\widehat{q}^{2t_{\text{mix}}(G_2)}_{u\sim D_{\widehat{z}},\cdot} - \pi^{\widehat{G}}(\cdot)\|_{\text{TV}} + \widehat{q}_{\widehat{z},\widehat{z}}\cdot 1 \tag{16}$$
$$\leqslant 1\cdot(1/e^2 + 3/32) + 1/32 \leqslant 1/e. \tag{17}$$

Thus, for $t' = 2t_{\text{mix}}(G_2) + 7$ we have $\|\widehat{q}^{t'}_{\widehat{z},\cdot} - \pi^{\widehat{G}}(\cdot)\|_{\text{TV}} \leqslant 1/e$. Together with (15), we conclude that for all $u \in V'$, $\|\widehat{q}^{t'}_{u,\cdot} - \pi^{\widehat{G}}(\cdot)\|_{\text{TV}} \leqslant 1/e$ and by definition of t_{mix} and we get $t_{\text{mix}}(\widehat{G}) \leqslant 2t_{\text{mix}} + 7 = O(1)$.

(ii) To prove $\mathbb{P}[\, T_{\text{hit}}(u,\widehat{z}) \geqslant n/c\,] \geqslant 1/2$ for $u \in V_2$ we show that the random walk restricted to $\widehat{G}$ does not hit $\widehat{z}$ 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

$$\mathbb{P}\left[\, T^G_{\text{hit}}(u,\widehat{z}) \leqslant n/c_1 \,\right] = \mathbb{P}\left[\, T^{\widehat{G}}_{\text{hit}}(u,\widehat{z}) \leqslant n/c_1 \,\right] \leqslant \sum_{t=1}^{n/c_1}\widehat{q}^t_{u,\widehat{z}}$$
$$\leqslant \sum_{t=1}^{c_2\log n} 2/\sqrt{n} + \sum_{t=c_2\log n}^{n/c_1}\widehat{q}^t_{u,\widehat{z}}$$
$$\leqslant o(1) + n/c_1\cdot(\pi^{\widehat{G}}(\widehat{z}) + 1/n^2) \leqslant 1/2,$$

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

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

$$\mathbb{P}\left[\,\widetilde{Z} \geqslant \gamma\,\right] \leqslant 1/4.$$

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

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

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

$$\mathbb{P}\left[\, T^G_{\text{hit}}(u,\widehat{z}) \geqslant n/c_1 \,\right] \geqslant \mathbb{P}\left[\, \cup_{t=0}^{n/c_1} X_t = \widetilde{X}_t \,\right] \geqslant 1 - \mathbb{P}\left[\, \widetilde{Z} \geqslant \gamma \,\right] - (1-p) \geqslant 1/2$$

and the proof is complete.

(iii) We proceed by showing $t_{\text{hit}}(u,\widehat{z}) = O(n)$ 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 $S^* = N(\widehat{z})$ be the neighbors of $\widehat{z}$ in G_2. For every $v \in S^*$ we have $\pi^{G_2}(v) = \frac{(1-o(1))\sqrt{n}}{\frac{n}{\sqrt{\alpha'}}(1+o(1))\sqrt{n}+\frac{\sqrt{n}}{\sqrt{\alpha'}}} \geqslant \frac{\sqrt{\alpha'}}{1.2n}$. Hence, after $t = t_{\text{sep}}(G_2)$ we have that

$$q^t_{u,S^*} := \sum_{v \in S^*} q^t_{u,v} \geqslant \sum_{v \in S^*} \pi^{G_2}(v)(1-e^{-1}) \geqslant \frac{\sqrt{n}}{\sqrt{\alpha'}} \cdot \frac{\sqrt{\alpha'}}{1.2n}(1-e^{-1}) = \frac{1-e^{-1}}{1.2\sqrt{n}}.$$

By Lemma 3.8, we have for any $u \in V_2$ that $\|p^t_{u,\cdot} - q^t_{u,\cdot}\|_{\text{TV}} \leqslant (1+o(1))t_{\text{sep}}(G_2)/(2\sqrt{n})$. To bound $T^G_{\text{hit}}(u,\widehat{z})$ we show that after $t_{\text{sep}} + 1 = O(1)$ steps the random walk hits $\widehat{z}$ w.p. $\Omega(1/n)$.

We distinguish between two cases.

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

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

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

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

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

We now show a similar statement if $u \in V_1$. Let $(X_t)_{t \geqslant 0}$ be a random walk on G starting on u. Observe that X_t (the walk on G) hits $\widehat{z}$ with probability $p^1_{u,z^1} \cdot p^1_{z^1,\widehat{z}} = \Omega(1/n)$ in 2 time steps. Hence, for any $u \in V$ we $\mathbb{P}\left[\, T_{\text{hit}}(u,\widehat{z}) = O(1) \,\right] = \Omega(1/n)$. Thus, repeating this iteratively and using independence yields $t_{\text{hit}}(u,\widehat{z}) = O(n)$ 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 ([41]). *For any* $\beta < 1/2$, *let* $t_{\mathrm{hit}}(\beta) = \max_{u,A:\pi(A)\geqslant\beta} t_{\mathrm{hit}}(u,A)$. *Then there exist positive constants* c_β *and* c'_β *such that*

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

In the following we show for any β close enough to $1/2$, that any $A \subseteq V$ satisfying $\pi(A) \geqslant \beta$ 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* $\pi(A) \geqslant \beta$, *define* $H(A) = \{i \mid |G_1^i \cap A| \geqslant |V_1|/(2e)\}$. *Then,* $|H(A)| \geqslant \kappa/(2e)$.

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

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

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

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* $t_{\mathrm{mix}}(G) = \Theta(n)$.

Proof. First we show $t_{\mathrm{mix}} = \Omega(n)$. The *conductance* of $G = (V, E)$ is defined by

$$\Phi(G) = \min_{\substack{U \subseteq V, \\ 0 < \mathrm{vol}(U) \leqslant \mathrm{vol}(V)/2}} \frac{|E(U, V \setminus U)|}{\mathrm{vol}(U)}.$$

In particular, for $U = V_1$ we get that $\Phi(G) \leqslant \frac{4}{n}$. Hence, by Cheeger's inequality and

$$\left(\frac{1}{1 - \lambda_2(G)} - 1\right) \cdot \log(\frac{e}{2}) \leqslant t_{\mathrm{mix}}(1/e)$$

(see, *e.g.*, [33, Chapter 12]),

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

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

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 $\pi(A) \geqslant \beta$. First, we apply Claim 3.11 to conclude that $|H(A)| \geqslant \kappa/(2e)$. This immediately implies that with $Z := \{z^i : i \in H(A)\}$, $|Z| \geqslant \kappa/(2e)$. 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 $2\max_{u \in V} t_{\mathrm{hit}}(u, \widehat{z}) = O(n)$ 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 $t_{\mathrm{hit}}(u, A) = O(n)$, and hence by Theorem 3.10, $t_{\mathrm{mix}} = O(n)$ as needed.

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

Part (ii). If the random walk is on z^*, then since $\deg(z^*) = \kappa + \sqrt{n/\alpha'}$, $|Z| \geqslant \kappa/(2e)$, 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. □

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 (2). 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 says 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 $V_1^1 \cup V_1^2 \cup \cdots \cup V_1^{\kappa}$. 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 $\widehat{z}$.

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

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

We have $\mathbb{E}[Z_1] \leqslant \ell n \left(\frac{2}{n}\right)^2$, $\mathbb{E}[Z_2] \leqslant \ell \frac{n}{\sqrt{\alpha'}} \left(\frac{2}{n}\right)^2$, and $\mathbb{E}[Z_*] \leqslant \ell \left(\frac{2}{n}\right)^2$, since $\max_u \pi(u) \leqslant 2/n$. Conditioning on $Z_1 \geqslant 1$ 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 $\mathbb{E}[Z_1 \mid Z_1 \geqslant 1]$. 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 $T_{\text{hit}}(u, \widehat{z}) \geqslant n/c$ for both walks, where $c > 0$ is a large enough constant. By Lemma 3.9.(ii), we have that $\mathbb{P}[\mathcal{E}_1] \geqslant (1/2)^2 = 1/4$ due to independence of the walks. For any $t < n/c$ let $\widehat{p}_{u,\cdot}^t$ 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 $\sum_{v \in V_1} \widehat{p}_{u,v}^t = 1$ implying that $\sum_{v \in V_1} (\widehat{p}_{u,v}^t)^2 \geqslant \sum_{v \in V_1} \left(\frac{1}{|V_1|}\right)^2 = 1/|V_1|$. Hence, we get

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

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

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

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

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

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 $(X_t)_{t \geqslant 0}$ and $(Y_t)_{t \geqslant 0}$ on G, both starting from arbitrary nodes. Note $t_{\text{sep}} = t_{\text{sep}}(G) \leqslant 4t_{\text{mix}} = O(n)$ by Lemma 3.12, and

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

We assume in the following that $\left\{X_{t_{\text{sep}}} \in V_2\right\} \cap \left\{Y_{t_{\text{sep}}} \in V_2\right\}$. We have $t_{\text{mix}}(G_2) = O(1)$, by Corollary A.3. Consider a random walk $(\widetilde{X}_t)_{t \geqslant t_{\text{sep}}}$ restricted to G_2 that starts at vertex $X_{t_{\text{sep}}} \in V_2$ and let $\widetilde{P}$ denote the transition matrix. Furthermore, in order to couple the random walk $\widetilde{X}_t$ restricted to G_2 with a random walk in G, we will consider the random variable

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

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

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

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

$$\mathbb{P}\left[\widetilde{Z} \geqslant \gamma\right] \leqslant 1/8.$$

Consider now the straightforward coupling between a random walk $(X_t)_{t \geqslant t_{\text{sep}}}$ in G that starts at vertex $\widetilde{X}_{t_{\text{sep}}} \in V_2$ and the random walk $(\widetilde{X}_t)_{t \geqslant t_{\text{sep}}}$ restricted to G_2 that starts at the same vertex. Whenever the random walk $\widetilde{X}_t$ is at a vertex in $V_2 \setminus \{N(\widehat{z})\}$, then the random walk X_t makes the same transition. If the random walk $\widetilde{X}_t$ is at vertex $z' \in N(\widehat{z})$, then there is a coupling so that the random walk X_t makes the same transition as $\widetilde{X}_t$ with probability $\frac{(1-\nu)\sqrt{n}}{2(1-\nu)\sqrt{n}+2}$ for some $\nu = o(1)$. Conditional on the event $\{\widetilde{Z} \leqslant \gamma\}$ occurring, the random walk $\widetilde{X}_t$ follows the random walk X_t up until step n/c with probability at least

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

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

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

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

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

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

Therefore, by the Union bound,

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

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

□

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 $t_{\mathrm{coal}}(G) = \Omega(\sqrt{\alpha'} \cdot n \log n)$.*

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 two of the n^{ε} random walks meet on a node in $V \setminus V_2$ 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 in $V \setminus (V_2 \cup \widehat{z})$.
(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 $\mathbb{P}\left[\mathcal{E}_3 \mid \overline{\mathcal{E}_2} \right] < 1/e$, which implies, by union bound,

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

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 $t_{\mathrm{coal}}(G) = \Omega(\sqrt{\alpha'} \cdot n \log n)$ 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 $\widehat{z}$ in a fixed step $t \geqslant 0$ is at most $O(1/n^2)$.

Hence, by the Union bound over $\binom{n^{\varepsilon}}{2}$ pairs of random walks and $\sqrt{\alpha'} \cdot n \log^2 n \leqslant n \log^3 n$ steps, the probability of any two random walks meeting on $\widehat{z}$ 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 walks 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 $G_1^1, G_1^2, \ldots, G_1^{\kappa}$. 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

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

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

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

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

(iii) We now prove $\mathbb{P}\left[\mathcal{E}_3 \mid \overline{\mathcal{E}_2}\right] < 1/4$. 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 $\widehat{z}$ required before the random walk hits G_2 instead of returning to G_1 is w.c.p. at least $(1 - o(1))\sqrt{\alpha'}/2$; this is because the fraction of edges from $\widehat{z}$ to G_2 is $\sqrt{n/\alpha'}/(\sqrt{n/\alpha'} + (1 \pm o(1))\sqrt{n})$. 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 $\lambda \geqslant 1$. 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

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

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 $\{z^1, z^2, \ldots, z^\kappa\}$. 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 $t_{\text{coal}} \in [\Omega(\log n), O(n^3)]$

Given that worst-case upper and lower bounds have long been known for t_{mix}, t_{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.

4.1 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 4.1. *For any graph G we have $t_{\text{coal}} = O(n \cdot |E| \cdot \log(|E|/n))$, so in particular, $t_{\text{coal}} = O(n^3)$.*

This result was first derived in a self-contained fashion in an earlier version of this work, which gave the first tight bound on the worst-case coalescing time. Due to the subsequent work [39], Theorem 4.1 follows immediately.

Proof of Theorem 4.1. This follows directly from $t_{\text{coal}} = O(t_{\text{hit}})$ [39, Theorem 5.1], and the standard bound $t_{\text{hit}} \leqslant n \cdot 2|E|$ (cf. [3]). □

4.2 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 $\mathbb{P}\left[Z_{u,t} = 1 \right] = 1/2$ and $\mathbb{P}\left[Z_{u,t} = 0 \right] = 1/2$. If $Z_{u,t} = 1$, then the random walk on u at time t (if there is any), moves to a neighboring 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 4.2. *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 $t_u \leqslant \tau$ such $Z_{u,t_u} = 1$. In symbols, let T be the first point in time where all walks coalesced, then $T \geqslant T'$, with $T' := \min\{t' \in \mathbb{N} \colon |\{u \colon \exists t_u \leqslant t' \text{ s.t. } Z_{u,t_u} = 1\}| \geqslant n-1\}$. Let Y_u be the indicator variable which is 1 if $Z_{u,t} = 0$ for all $t \leqslant \tau := \log n/2$. The process ensures independence of the Y_u. Due to the laziness of the random walk, $\mathbb{P}\left[Y_u = 1 \right] = 1/2^\tau = 1/\sqrt{n}$. Thus, using the independence of the Y_u,

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

where 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. □

4.3 Proof of Theorem 1.4

We are now ready to put all the pieces together. The upper bound on general graphs follows directly from $t_{\text{meet}} \leqslant t_{\text{coal}} = O(n^3)$, by Theorem 4.1. The lower bound on the meeting time holds by definition and the lower bound on the coalescing time follows from Lemma 4.2. For the upper bound on regular graphs we have $t_{\text{meet}} \leqslant t_{\text{coal}} = O(t_{\text{hit}}) = O(n^2)$ having used the standard bound $t_{\text{hit}} = O(n^2)$ for regular graphs (see [2]). The lower bound follows from $t_{\text{coal}} \geqslant t_{\text{meet}} \geqslant t^{\pi}_{\text{meet}} = \Omega(n)$, by Theorem 5.1.

5 BASIC BOUNDS

Although the focus of this work is on understanding the coalescence time, in order to apply our general results, we also devise some tools to obtain lower and upper bounds on t_{meet}. In Theorem 5.1 (Section 5.1) we establish upper and lower bounds on the meeting time in terms of $\|\pi\|_2^2 = \sum_{u \in V} \pi(u)^2$. Section 5.1 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 5.7 (Section 5.2) we establish a discrete-time counterpart of [2, 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.

5.1 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 (2) and its counterpart for the hitting times (cf. Cooper and Frieze [18]). In these bounds, we will use the following notation:

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

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 [14]), using $\sum_{t=0}^{\infty} \lambda_2^t = \frac{1}{1-\lambda_2}$.

Before stating the next result, we recall that G is Γ-approximative regular if $\Delta/\delta = \Gamma$.

Theorem 5.1. *For any graph $G = (V, E)$, the following statements hold:*

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

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

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

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

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

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

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

(iii) For any graph G it holds that,

$$t_{\text{meet}} \geqslant t_{\text{meet}}^{\pi} \geqslant \frac{3}{16\|\pi\|_2^2}.$$

In particular, if the graph G is Γ-approximate regular, then $t_{\text{meet}} \geqslant t_{\text{meet}}^{\pi} = \frac{3}{16} \cdot \frac{n}{\Gamma}$. Furthermore, if $\frac{C_{\min}}{16\|\pi\|_2^2 \cdot t_{\text{mix}}} \geqslant 1$, then we have

$$t_{\text{meet}} \geqslant t_{\text{meet}}^{\pi} \geqslant \frac{C_{\min}}{16\|\pi\|_2^2}.$$

Let us make two remarks. First, note that the second upper bound on $t_{\text{meet}}(u,v)$ depends only logarithmically on Γ. Secondly, the final lower bound involving $C_{\min}$ could be tightened further by having a different range of the sum in $C_{\min}$, but for simplicity we only state this version.

Proof. We begin by proving the first part. Consider one random walk $(X_t)_{t \geqslant 0}$, 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, $\mathbb{E}[\,Z\,] \geqslant t_{\text{mix}} \cdot \frac{1}{e} \cdot \pi(v)$, and (2) yields

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

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

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

This means that in every epoch of length $t_{\text{sep}} + t_{\text{mix}} \leqslant 5t_{\text{mix}}$, 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

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

The claim for Γ-approximate regular graphs follows from the observation that $\min_{u \in V} \pi(u) \geqslant 1/(\Gamma n)$. We continue with the second part. Consider two independent random walks, $(X_t)_{t \geqslant 0}$, $(Y_t)_{t \geqslant 0}$ 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 $t_{\mathrm{sep}} + t_{\mathrm{mix}} - 1$, *i.e.,*

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

By linearity of expectation,

$$\mathbb{E}[\,Z\,] = \sum_{t=t_{\mathrm{sep}}}^{t_{\mathrm{sep}}+t_{\mathrm{mix}}-1} \sum_{u \in V} \mathbb{P}[\,X_t = u\,] \cdot \mathbb{P}[\,Y_t = u\,] \tag{20}$$

$$\geqslant \sum_{t=t_{\mathrm{sep}}}^{t_{\mathrm{sep}}+t_{\mathrm{mix}}-1} (1-1/e) \cdot \pi_u \cdot (1-1/e) \cdot \pi_u \tag{21}$$

$$\geqslant t_{\mathrm{mix}} \cdot \frac{1}{e} \cdot \|\pi\|_2^2, \tag{22}$$

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

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

Plugging this into (2) and using (22) we finally arrive at

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

Hence,

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

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

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

and note that $S_0, \ldots, S_{\log_2 n}$ 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

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

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

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

Then,

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

Furthermore,

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

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

$$\begin{aligned}\mathbb{E}\left[Z_j \mid Z_j \geqslant 1 \right] &\leqslant 2 \cdot \max_{u \in S_i} \sum_{t=0}^{t_{\text{mix}}-1} \sum_{v \in V} p_{u,v}^t \cdot p_{v,u}^t \\ &\leqslant 2 \cdot \max_{u \in V} \sum_{t=0}^{t_{\text{mix}}-1} p_{u,u}^{2t} \leqslant 2 \cdot \max_{u \in V} \sum_{t=0}^{t_{\text{mix}}-1} p_{u,u}^t = 2 \cdot R_{\max},\end{aligned}$$

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

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

Finally, for the third statement, let $(X_t)_{t \geqslant 0}$, $(Y_t)_{t \geqslant 0}$ be two random walks starting from stationarity. Define $\widetilde{Z}$ as the number of collisions between steps 0 and $\ell := \frac{1}{4\|\pi\|_2^2}$, *i.e.*,

$$\widetilde{Z} := \sum_{t=0}^{\ell} \mathbf{1}_{X_t = Y_t}.$$

Then, $\mathbb{E}\left[\widetilde{Z} \right] = \frac{1}{4\|\pi\|_2^2} \cdot \|\pi\|_2^2 = \frac{1}{4}$. By Markov's inequality, $\mathbb{P}\left[\widetilde{Z} \geqslant 1 \right] \leqslant \frac{1}{4}$, and hence $t_{\text{meet}}^{\pi} \geqslant \frac{3}{4} \cdot \frac{1}{4\|\pi\|_2^2} = \frac{3}{16} \cdot \frac{1}{\|\pi\|_2^2}$. If the graph is Γ-approximative regular, then the maximum degree Δ satisfies $\Delta \leqslant \Gamma \cdot d$, where d is the average degree, and so

$$\|\pi\|_2^2 = \sum_{u \in V} \left(\frac{\deg(u)}{2|E|} \right)^2 \leqslant \frac{nd}{\Gamma d} \cdot \left(\frac{\Gamma d}{nd} \right)^2 = \frac{\Gamma}{n}.$$

For the next part of the third statement, let $\widetilde{Z}$ be the random variable counting the number of collisions between steps 0 and $2t_{\text{mix}}$, *i.e.*,

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

Then,

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

In order to lower bound $\mathbb{E}\left[\widetilde{Z} \mid \widetilde{Z} \geqslant 1 \right]$, let us write $\widetilde{Z} = \widetilde{Z}_1 + \widetilde{Z}_2$ with $\widetilde{Z}_1 := \sum_{t=0}^{t_{\text{mix}}} \mathbf{1}_{X_t = Y_t}$ and $\widetilde{Z}_2 := \sum_{t=t_{\text{mix}}+1}^{2t_{\text{mix}}-1} \mathbf{1}_{X_t = Y_t}$. By Observation A.8,

$$\mathbb{P}\left[\widetilde{Z}_1 \geqslant 1 \mid \widetilde{Z} \geqslant 1 \right] \geqslant \frac{1}{2}.$$

Therefore, by law of total expectation,

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

Hence,

$$\mathbb{P}\left[\widetilde{Z} \geqslant 1\right] \leqslant \frac{2t_{\mathrm{mix}} \cdot \|\pi\|_2^2}{\frac{1}{2} \cdot C_{\min}} =: p.$$

By the precondition, we have $\frac{C_{\min}}{16\|\pi\|_2^2 \cdot t_{\mathrm{mix}}} \geqslant 1$, so consider now $\lceil \frac{C_{\min}}{16\|\pi\|_2^2 \cdot t_{\mathrm{mix}}} \rceil \geqslant 1$ consecutive time-intervals of length $2t_{\mathrm{mix}}$ each. We conclude that, if B denotes the total number of collisions between the walks across all the intervals,

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

Hence by Markov's inequality, $\mathbb{P}\left[\, B \geqslant 1 \,\right] \leqslant \frac{1}{2}$ and thus

$$t_{\mathrm{meet}}^{\pi} \geqslant \frac{1}{2} \cdot \left\lceil \frac{C_{\min}}{16\|\pi\|_2^2 \cdot t_{\mathrm{mix}}} \right\rceil \cdot 2t_{\mathrm{mix}} \geqslant \frac{C_{\min}}{16\|\pi\|_2^2}.$$

□

It is interesting to compare the upper bound on t_{meet} in Theorem 5.1 with the bound $t_{\mathrm{meet}} = O(\frac{1}{1-\lambda_2} \cdot (\frac{1}{\|\pi\|_2^2} + \log n))$ from Cooper et al. [14, Theorem 2]. Using the trivial bound $C_{\max} \leqslant t_{\mathrm{mix}}$ and $t_{\mathrm{mix}} = O(\frac{\log n}{1-\lambda_2})$, we obtain $t_{\mathrm{meet}} = O(\frac{1}{1-\lambda_2} \cdot \frac{\log n}{\|\pi\|_2^2})$, 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 $C_{\max} \ll t_{\mathrm{mix}}$ could be established.

Proposition 5.2. *For any graph $G = (V, E)$,*

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

which is at least as good as the bound of [14, Theorem 1] and equally good if one uses the trivial bound $t_{\mathrm{mix}} = O\left(\frac{\log n}{1-\lambda_2}\right)$.

Proof. We first derive the claimed bound (by a case distinction), and then compare it to the bound of [14, Theorem 1]. The first case is $t_{\mathrm{meet}}/t_{\mathrm{mix}} \geqslant \log^2 n$. In this case,

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

where the first inequality is by Theorem 1.1 and the second inequality by [14, Thereom 2]. The second case is $t_{\mathrm{meet}}/t_{\mathrm{mix}} \leqslant \log^2 n$, so $t_{\mathrm{meet}} \leqslant t_{\mathrm{mix}} \cdot \log^2 n$. By Proposition 3.4, $t_{\mathrm{coal}} = O(t_{\mathrm{meet}} \log n) = O\left(t_{\mathrm{mix}} \log^3 n\right)$. Using $t_{\mathrm{mix}} = \Omega(\frac{1}{1-\lambda_2})$, we can state the following (unconditional) upper bound:

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

Using $t_{\mathrm{mix}} = O\left(\frac{\log n}{1-\lambda_2}\right)$ we derive indeed $O\left(\frac{1}{1-\lambda_2} \cdot (\frac{1}{\|\pi\|_2^2} + \log^4 n)\right)$, which is the same bound as [14, Theorem 1]. □

We will make use of the following standard result about return probabilities (related bounds were proven in [2, Proposition 6.16 (iii)], [36, Theorem 4.9].

Lemma 5.3 ([39, Theorem 1.2]). *Let G be any regular graph. Then for any $t \leqslant 10n^2$,*

$$p_{u,u}^t \leqslant \frac{20}{\sqrt{t+1}}.$$

Theorem 5.4. *For any regular graph we have*

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

and by Cheeger's inequality we obtain $t_{\mathrm{hit}} = O(\frac{n}{\Phi})$*, where* Φ *is the conductance of* G*.*

For non-regular graphs, the best bound on the hitting time was shown in [39, Theorem 1.1], which states that for any graph G with average degree d and minimum degree δ,

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

Proof. By applying first [33, Lemma 10.2] and then [33, Proposition 10.26],

$$\begin{aligned} t_{\mathrm{hit}} &\leqslant 2 \max_{u \in V} \sum_{v \in V} t_{\mathrm{hit}}(v,u)\pi(v) \\ &= \max_{u \in V} \frac{2}{\pi(u)} \sum_{t=0}^{\infty} \left(p_{u,u}^t - \pi(u)\right) =: \max_{u \in V} f(u). \end{aligned}$$

In the following we will upper bound $f(u)$ for an arbitrary vertex $u \in V$. We write $t_{\mathrm{rel}} = 1/(1-\lambda_2)$.

For lazy random walks, it holds that for any $u \in V$ and $t \geqslant 1$, $p_{u,u}^t \geqslant \pi(u)$ (Lemma A.1, third statement). Hence, for any $u \in V$,

$$\begin{aligned} f(u) &\leqslant \frac{2}{\pi(u)} \sum_{t=0}^{2t_{\mathrm{rel}}} \left(p_{u,u}^t\right) + \frac{2}{\pi(u)} \sum_{t=2t_{\mathrm{rel}}+1}^{\infty} \left(p_{u,u}^t - \pi(u)\right) \\ &\leqslant \frac{O(1)}{\pi(u)} \sqrt{t_{\mathrm{rel}}} + \frac{2}{\pi(u)} 2t_{\mathrm{rel}} \sum_{k=0}^{\infty} \left(p_{u,u}^{2(t_{\mathrm{rel}}+kt_{\mathrm{rel}})} - \pi(u)\right), \end{aligned} \tag{23}$$

where the bound on the first term of the last inequality follows from $p_{u,u}^{\sqrt{t_{\mathrm{rel}}}} \leqslant O\left(\frac{1}{\sqrt{t_{\mathrm{rel}}}}\right)$ (Lemma 5.3).

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

$$\begin{aligned} \sum_{k=0}^{\infty} \left(p_{u,u}^{2(t_{\mathrm{rel}}+kt_{\mathrm{rel}})} - \frac{1}{n}\right) &= \sum_{k=0}^{\infty} \left(\sum_{v \in V} \left(p_{u,v}^{t_{\mathrm{rel}}+kt_{\mathrm{rel}}}\right)^2 - \frac{1}{n}\right) \\ &\overset{(*)}{=} \sum_{k=0}^{\infty} \left(\sum_{v \in V} \left(p_{u,v}^{t_{\mathrm{rel}}+kt_{\mathrm{rel}}}\right)^2 - 2 \sum_{v \in V} p_{u,v}^{t_{\mathrm{rel}}+kt_{\mathrm{rel}}} \cdot \frac{1}{n} + \sum_{v \in V} \frac{1}{n^2}\right) \\ &\overset{(**)}{=} \sum_{k=0}^{\infty} \left\| p_{u,\cdot}^{t_{\mathrm{rel}}+kt_{\mathrm{rel}}} - \frac{1}{n} \right\|^2 \\ &= \sum_{k=0}^{\infty} \left\| p_{u,\cdot}^{t_{\mathrm{rel}}} P^{k \cdot t_{\mathrm{rel}}} - \frac{1}{n} \right\|^2 \\ &= \sum_{k=0}^{\infty} \left\| (P^{k \cdot t_{\mathrm{rel}}})^T p_{u,\cdot}^{t_{\mathrm{rel}}} - \frac{1}{n} \right\|^2 \\ &= \sum_{k=0}^{\infty} \left\| P^{k \cdot t_{\mathrm{rel}}} p_{u,\cdot}^{t_{\mathrm{rel}}} - \frac{1}{n} \right\|^2 \end{aligned} \tag{24}$$

where we have used that for regular graphs, P is symmetric, i.e., $P = P^T$; moreover, (*) follows from the the fact that $\sum_{v\in V} p_{u,v}^t = 1$ and (**) is the Binomial theorem. Now using the spectral decomposition of $P^{k\cdot t_{\mathrm{rel}}}$ ([33, Exercise 12.4]), we have

$$\left\| P^{k\cdot t_{\mathrm{rel}}} p_{u,\cdot}^{t_{\mathrm{rel}}} - \frac{1}{n} \right\|^2 \leqslant \left(\lambda_2(P^{k\cdot t_{\mathrm{rel}}})\right)^2 \cdot \left\| p_{u,\cdot}^{t_{\mathrm{rel}}} - \frac{1}{n} \right\|^2$$

since due to laziness, all eigenvalues are non-negative, and thus λ_2 is the second largest eigenvalue of P (and $P^{k\cdot t_{\mathrm{rel}}}$) in absolute value. Hence, using that $\sum_{v\in V} p_{u,v}^{t_{\mathrm{rel}}} = 1$, we get

$$\begin{aligned}
\sum_{k=0}^{\infty} \left(p_{u,u}^{2(t_{\mathrm{rel}}+kt_{\mathrm{rel}})} - \frac{1}{n} \right) &\leqslant \sum_{k=0}^{\infty} \lambda_2^{2k\cdot t_{\mathrm{rel}}} \left\| p_{u,\cdot}^{t_{\mathrm{rel}}} - \frac{1}{n} \right\|^2 \\
&= \sum_{k=0}^{\infty} \lambda_2^{2k\cdot t_{\mathrm{rel}}} \left(\sum_{v\in V} (p_{u,v}^{t_{\mathrm{rel}}})^2 - 2\frac{1}{n} \sum_{v\in V} p_{u,v}^{t_{\mathrm{rel}}} + \sum_{v\in V} \frac{1}{n^2} \right) \\
&\leqslant \sum_{k=0}^{\infty} \lambda_2^{2k\cdot t_{\mathrm{rel}}} p_{u,u}^{2t_{\mathrm{rel}}} \\
&\overset{(*)}{=} O\left(\sum_{k=0}^{\infty} \lambda_2^{2k\cdot t_{\mathrm{rel}}} \frac{1}{\sqrt{t_{\mathrm{rel}}}} \right) \\
&\overset{(**)}{=} O\left(\frac{1}{\sqrt{t_{\mathrm{rel}}}} \right), \qquad (25)
\end{aligned}$$

where $(*)$ follows from $p_{u,u}^{2t_{\mathrm{rel}}} \leqslant O\left(\frac{1}{\sqrt{t_{\mathrm{rel}}}}\right)$ (Lemma 5.3) and $(**)$ follows since $g(y) = y^{1/(1-y)}$ is bounded from above by $1/e$ for any $y \in (0,1)$ and hence the sum is a geometric series. Combining (23), (24) and (25) yields the claim.

□

It turns out that the hitting time bound of Theorem 5.4 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 5.5 ([11]). *For every n, $d \geqslant 3$, and constant Φ, there exists a d-regular graph G with n nodes and a constant conductance such that the expected coalescing 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 time on G is $\Omega(n/\Phi)$. Therefore, the coalescence time and hitting time are of order $\Omega(n/\Phi)$.*

We will now apply hitting time bound of Theorem 5.4 to deduce results for three other quantities.

Theorem 5.6. *For any regular graph G, the following three bounds hold:*

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

and

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

and

$$t_{\mathrm{cov}} = O\left(\frac{n}{\sqrt{1-\lambda_2}} \cdot \log n \right).$$

The upper bound on t_{meet} and t_{coal} gives $t_{\mathrm{meet}} = O(n^2)$ for cycles and paths, and $t_{\mathrm{meet}} = O(n)$ on regular expanders (since $1/(1-\lambda_2) = O(1)$). It thus improves the bound by Cooper et al. [14, Theorem 1], which states that for any regular graph, $t_{\mathrm{meet}} = O(n/(1-\lambda_2))$. For any regular graph, we also improve the best-known bound on the cover time t_{cov} in terms of the eigenvalue gap, which is $t_{\mathrm{cov}} = O\left(n \log n/(1-\lambda_2)\right)$ established by Broder and Karlin [12].

Proof. The proof of the first part follows from Theorem 5.4 and $t_{\mathrm{meet}} \leqslant 2t_{\mathrm{hit}} + 1$ for regular graphs (Proposition 5.7). The second statement follows by combining the Theorem 5.4 with $t_{\mathrm{coal}} = O(t_{\mathrm{hit}})$ [39]. The third statement follows also from Theorem 5.4 and Matthew's bound $t_{\mathrm{cov}} = O(t_{\mathrm{hit}} \cdot \log n)$, which holds for any graph G. □

5.2 Relating Meeting Time to t_{hit}

In this section we state the following proposition that relates the meeting time to the hitting time. It can be seen as an analogous version of [2, Proposition 14.5] in discrete time. The upper bound is very similar to the bound given by [21, Theorem 2].

Proposition 5.7. *For any graph $G = (V, E)$ and $u, v \in V$ we have*

$$(\min_{u' \in V} t_{\mathrm{hit}}(\pi, u') + (t_{\mathrm{hit}}(u,v) - t_{\mathrm{hit}}(\pi, v)))/2 \leqslant t_{\mathrm{meet}}(u,v) \leqslant \max_{u,v \in V}(t_{\mathrm{hit}}(u,v) - t_{\mathrm{hit}}(\pi,v)) + \max_{u' \in V} t_{\mathrm{hit}}(\pi, u') + 1.$$

Consequently, for any graph we have $t_{\mathrm{meet}} \leqslant 2t_{\mathrm{hit}} + 1$ and for any vertex transitive graph G we have $t_{\mathrm{hit}}/2 \leqslant t_{\mathrm{meet}} \leqslant t_{\mathrm{hit}} + 1$.

Proof. We define a pair of chains $((X_t)_{t\geqslant 0}, (Y_t)_{t \geqslant 0})$ 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 $(X'_t)_{t\geqslant 0}$ and $(Y'_t)_{t\geqslant 0}$ with $X'_0 = X_0$ and $Y'_0 = Y_0$. We will couple their decisions with the walks $(X_t)_{t\geqslant 0}$ and $(Y_t)_{t\geqslant 0}$, by setting $X'_t = X_{2t}$ and $Y'_t = Y_{2t}$ for all integers $t \geqslant 0$. Due to this coupling and since each random walk is lazy w.p. 1/2,

$$\mathbb{P}\left[\, X'_t = Y'_t \mid X_{2t} = Y_{2t} \,\right] = 1,$$

and

$$\mathbb{P}\left[\, X'_{t+1} = Y'_{t+1} \mid X_{2t+1} = Y_{2t+1} \,\right] = \mathbb{P}\left[\, X_{2t+2} = Y_{2t+2} \mid X_{2t+1} = Y_{2t+1} \,\right] = 1/2$$

Let $t_{\mathrm{meet}}^{\mathrm{seq}}(u,v)$ be the meeting time of the sequential chains X_t and Y_t, *i.e.,*

$$t_{\mathrm{meet}}^{\mathrm{seq}}(u,v) = \min\{t \geqslant 0 \mid X_t = Y_t, X_0 = u, Y_0 = v\}$$

and $t_{\mathrm{meet}}^{\mathrm{seq}} = \max_{u,v} t_{\mathrm{meet}}^{\mathrm{seq}}(u,v)$. We seek to relate t_{meet} and $t_{\mathrm{meet}}^{\mathrm{seq}}$. Clearly, for any pair $u, v \in V$, $t_{\mathrm{meet}}^{\mathrm{seq}}(u,v)/2 \leqslant t_{\mathrm{meet}}(u,v)$ 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 1/2 whenever $X_{2t} = Y_{2t}$ or $X_{2t-1} = Y_{2t-1}$. Hence, we can consider intervals of length $t_{\mathrm{meet}}^{\mathrm{seq}}$. In each of them, we have, independent of what happened in previous intervals, that the random walks meet w.p. at least 1/2. Thus, $t_{\mathrm{meet}}(u,v) \leqslant 2(t_{\mathrm{meet}}^{\mathrm{seq}} + 1)/2 = t_{\mathrm{meet}}^{\mathrm{seq}} + 1$. We conclude,

$$t_{\text{meet}}^{\text{seq}}(u,v)/2 \leqslant t_{\text{meet}}(u,v) \leqslant t_{\text{meet}}^{\text{seq}} + 1. \tag{26}$$

We proceed by deriving a lower bounds on $t_{\text{meet}}^{\text{seq}}(u,v)$ and an upper bound on $t_{\text{meet}}^{\text{seq}}(u,v)$, which gives us the desired bounds on t_{meet}. We will make use of the following statement that is a weaker version of the original statement Aldous and Fill [2, Proposition 3.3]. For all $u,v \in V$ we have

$$\min_{u' \in V} t_{\text{hit}}(\pi,u') \leqslant t_{\text{meet}}^{\text{seq}}(u,v) - (t_{\text{hit}}(u,v) - t_{\text{hit}}(\pi,v)) \leqslant \max_{u' \in V} t_{\text{hit}}(\pi,u')$$

Using (26) we derive,

$$t_{\text{meet}}(u,v) \leqslant t_{\text{meet}}^{\text{seq}} + 1 \leqslant \max_{u,v \in V} (t_{\text{hit}}(u,v) - t_{\text{hit}}(\pi,v)) + \max_{u' \in V} t_{\text{hit}}(\pi,u') + 1,$$

and

$$t_{\text{meet}}(u,v) \geqslant t_{\text{meet}}^{\text{seq}}(u,v)/2 \geqslant (\min_{u' \in V} t_{\text{hit}}(\pi,u') + (t_{\text{hit}}(u,v) - t_{\text{hit}}(\pi,v)))/2,$$

which yields the first part of the claim. This also implies that $t_{\text{meet}} \leqslant 2t_{\text{hit}} + 1$. In the remainder of the proof we focus on vertex transitive chains. By vertex transitivity, we have $t_{\text{hit}}(\pi,u) = t_{\text{hit}}(\pi,u')$ for all $u,u' \in V$ and thus using (26)

$$t_{\text{meet}}^{\text{seq}}(u,v) = t_{\text{hit}}(u,v).$$

Thus, putting everything together and fixing $u,v \in V$ to be the nodes maximizing $t_{\text{hit}}(u,v)$, we derive

$$t_{\text{hit}} = t_{\text{hit}}(u,v) = t_{\text{meet}}^{\text{seq}}(u,v) \leqslant t_{\text{meet}}^{\text{seq}} \leqslant 2t_{\text{meet}}.$$

Similarly, let $u,v \in V$ be the nodes maximising $t_{\text{meet}}^{\text{seq}}(u,v)$, we derive

$$t_{\text{meet}} \leqslant t_{\text{meet}}^{\text{seq}} + 1 = t_{\text{meet}}^{\text{seq}}(u,v) + 1 = t_{\text{hit}}(u,v) + 1 \leqslant t_{\text{hit}} + 1.$$

This yields Proposition 5.7. □

5.3 Proof of Theorem 1.3

The proof follows from Theorem 5.4 and Theorem 5.6.

6 SPECIFIC GRAPHS

Here we present bounds on the coalescence time for specific graphs.

6.1 d-Dimensional Grids and Tori, $d > 2$

Here the bounds on t_{meet} and t_{coal} follow immediately from our general results. First, recall that for any regular (or almost-regular) graph we have $t_{\text{meet}} = \Omega(n)$ (Theorem 5.1.iii). Combining this with the fact that $t_{\text{mix}} = O(n^{2/d})$ [2], our Theorem 1.1 yields the bound $t_{\text{coal}} = \Theta(t_{\text{meet}})$. Further, it is a well-known result that $t_{\text{hit}} = O(n)$ (*e.g.*, [33]), and using $t_{\text{meet}} \leqslant 2t_{\text{hit}} + 1$ (Proposition 5.7) gives $t_{\text{meet}} = O(n)$, and hence also $t_{\text{coal}} = O(t_{\text{meet}}) = O(n)$. Let us point out that one could also use the more recent result $t_{\text{coal}} = O(t_{\text{hit}})$ ([39]) to deduce that $t_{\text{coal}} = O(n)$ without appealing to the meeting time.

6.2 Hypercubes

First, for any regular graph we again have $t_{\text{meet}} = \Omega(n)$ (Theorem 5.1.iii). Furthermore, since $t_{\text{meet}} \leqslant 2t_{\text{hit}} + 1$ and using that $t_{\text{hit}} = O(n)$ (see *e.g.,* [35]), we get $t_{\text{meet}} = O(n)$. It is a well-known fact that $t_{\text{mix}} = O(\log n \cdot \log\log n)$ [33], so applying Theorem 1.1 yields $t_{\text{coal}} = \Theta(t_{\text{meet}}) = \Theta(n)$.

Acknowledgments. F. Mallmann-Trenn was in part supported by the EPSRC grant EP/W005573/1. The authors would like to thank Petra Berenbrink, Robert Elsässer, Nikolaos Fountoulakis, Peter Kling, Roberto Oliveira, Perla Sousi and He Sun for helpful discussions and in particular Yuval Peres for pointing out how to further improve the bound on the hitting time (Theorem 1.3). Moreover, the authors would like to thank the anonymous reviewers of this article for their thorough reading and patience.

REFERENCES

[1] David Aldous. 1991. Meeting times for independent Markov chains. *Stochastic Processes and their Applications* 38, 2 (1991), 185–193.

[2] D. Aldous and J. Fill. 2002. Reversible Markov Chains and Random Walks on Graphs. Unpublished. http://www.stat.berkeley.edu/~aldous/RWG/book.html.

[3] R. Aleliunas, R.M. Karp, R.J. Lipton, L. Lovász, and C. Rackoff. 1979. Random Walks, Universal Traversal Sequences, and the Complexity of Maze Problems. In *FOCS*. 218–223.

[4] Dan Alistarh. 2020. Distributed Computing Column 77 Consensus Dynamics: An Overview. *ACM SIGACT News* 51, 1 (2020), 57–57.

[5] Noga Alon, Chen Avin, Michal Koucký, Gady Kozma, Zvi Lotker, and Mark R. Tuttle. 2011. Many Random Walks Are Faster Than One. *Combinatorics, Probability & Computing* (2011).

[6] Hagit Attiya and Keren Censor. 2008. Tight bounds for asynchronous randomized consensus. *J. ACM* 55, 5 (2008), 20:1–20:26. https://doi.org/10.1145/1411509.1411510

[7] Martin Barlow, Yuval Peres, and Perla Sousi. 2012. Collisions of Random Walks. *Annales de l' Institut Henri Poincare* 48, 4 (2012), 922–946.

[8] Luca Becchetti, Andrea Clementi, Emanuele Natale, Francesco Pasquale, and Riccardo Silvestri. 2015. Plurality Consensus in the Gossip Model. In *SODA*. 371–390.

[9] Luca Becchetti, Andrea E. F. Clementi, Emanuele Natale, Francesco Pasquale, and Luca Trevisan. 2016. Stabilizing Consensus with Many Opinions. In *SODA*. 620–635.

[10] Petra Berenbrink, Andrea E. F. Clementi, Robert Elsässer, Peter Kling, Frederik Mallmann-Trenn, and Emanuele Natale. 2017. Ignore or Comply?: On Breaking Symmetry in Consensus. In *PODC*. 335–344. https://doi.org/10.1145/3087801.3087817

[11] Petra Berenbrink, George Giakkoupis, Anne-Marie Kermarrec, and Frederik Mallmann-Trenn. 2016. Bounds on the Voter Model in Dynamic Networks. In *43rd International Colloquium on Automata, Languages, and Programming, ICALP 2016, July 11-15, 2016, Rome, Italy*. 146:1–146:15. https://doi.org/10.4230/LIPIcs.ICALP.2016.146

[12] Andrei Z. Broder and Anna R. Karlin. 1989. Bounds on the cover time. *Journal of Theoretical Probability* 2, 1 (1989), 101–120.

[13] Amin Coja-Oghlan. 2007. On the Laplacian Eigenvalues of Gn,p. *Combinatorics, Probability and Computing* 16, 6 (2007), 923–946. https://doi.org/10.1017/S0963548307008693

[14] Colin Cooper, Robert Elsässer, Hirotaka Ono, and Tomasz Radzik. 2013. Coalescing Random Walks and Voting on Connected Graphs. *SIAM Journal on Discrete Mathematics* 27, 4 (2013), 1748–1758.

[15] Colin Cooper, Robert Elsässer, and Tomasz Radzik. 2014. The Power of Two Choices in Distributed Voting. In *Proc. ICALP*. 435–446.

[16] Colin Cooper, Robert Elsässer, Tomasz Radzik, Nicolaás Rivera, and Takeharu Shiraga. 2015. Fast Consensus for Voting on General Expander Graphs. In *DISC*. 248–262.

[17] Colin Cooper, Alan Frieze, and Tomasz Radzik. 2009. Multiple Random Walks and Interacting Particle Systems. In *Automata, Languages and Programming*, Susanne Albers, Alberto Marchetti-Spaccamela, Yossi Matias, Sotiris Nikoletseas, and Wolfgang Thomas (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 399–410.

[18] Colin Cooper and Alan M. Frieze. 2005. The Cover Time of Random Regular Graphs. *SIAM J. Discrete Math.* 18, 4 (2005), 728–740.

[19] Colin Cooper, Alan M. Frieze, and Tomek. Radzik. 2009. Multiple Random Walks in Random Regular Graphs. *SIAM J. Discret. Math.* 23, 4 (2009), 1738–1761.

[20] Colin Cooper, Tomasz Radzik, Nicolás Rivera, and Takeharu Shiraga. 2016. Fast plurality consensus in regular expanders. *CoRR* abs/1605.08403 (2016). http://arxiv.org/abs/1605.08403

[21] Don Coppersmith, Prasad Tetali, and Peter Winkler. 1993. Collisions Among Random Walks on a Graph. *SIAM Journal on Discrete Mathematics* 6, 3 (1993), 363–374. https://doi.org/10.1137/0406029 arXiv:https://doi.org/10.1137/0406029

[22] J. Cox. 1989. Coalescing Random Walks and Voter Model Consensus Times on the Torus in $\mathbb{Z}^d$. *The Annals of Probability* 17, 4 (1989), 1333–1366.

[23] Artur Czumaj and Christian Sohler. 2010. Testing Expansion in Bounded-Degree Graphs. *Combinatorics, Probability & Computing* 19, 5-6 (2010), 693–709.

[24] Benjamin Doerr, Leslie Ann Goldberg, Lorenz Minder, Thomas Sauerwald, and Christian Scheideler. 2011. Stabilizing Consensus With the Power of Two Choices. In *Proc. SPAA*. 149–158.

[25] Klim Efremenko and Omer Reingold. 2009. How Well Do Random Walks Parallelize?. In *RANDOM*. 476–489.

[26] Robert Elsässer, Tom Friedetzky, Dominik Kaaser, Frederik Mallmann-Trenn, and Horst Trinker. 2016. Efficient k-Party Voting with Two Choices. *CoRR* abs/1602.04667 (2016). http://arxiv.org/abs/1602.04667

[27] Robert Elsässer and Thomas Sauerwald. 2011. Tight bounds for the cover time of multiple random walks. *Theor. Comput. Sci.* 412, 24 (2011), 2623–2641.

[28] Melvin Gauci, Monica E. Ortiz, Michael Rubenstein, and Radhika Nagpal. 2017. 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*. 1404–1412. http://dl.acm.org/citation.cfm?id=3091319

[29] Mohsen Ghaffari and Johannes Lengler. 2017. Tight Analysis for the 3-Majority Consensus Dynamics. *CoRR* abs/1705.05583 (2017). http://arxiv.org/abs/1705.05583

[30] Y. Hassin and D. Peleg. 2001. Distributed Probabilistic Polling and Applications to Proportionate Agreement. *Information and Computation* 171, 2 (2001), 248–268.

[31] Varun Kanade, Frederik Mallmann-Trenn, and Thomas Sauerwald. 2018. On coalescence time in graphs–When is coalescing as fast as meeting? arXiv:1611.02460 [cs.DM]

[32] Varun Kanade, Frederik Mallmann-Trenn, and Thomas Sauerwald. 2019. On coalescence time in graphs-When is coalescing as fast as meeting?. In *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA'19*. http://arxiv.org/abs/1611.02460 A full version with all proofs is available at http://arxiv.org/abs/1611.02460.

[33] D. Levin, Y. Peres, and E. Wilmer. 2006. *Markov chains and mixing times*. American Mathematical Society.

[34] Pascal Lezaud. 1989. Chernoff-type bound for finite Markov chains. *Ann. Appl. Probab.* 8, 3 (1989), 849–867.

[35] L. Lovász. 1993. Random Walks on Graphs: A Survey. *Combinatorics, Paul Erdős is Eighty* 2 (1993), 1–46.

[36] Russell Lyons and Shayan Oveis Gharan. 2017. Sharp Bounds on Random Walk Eigenvalues via Spectral Embedding. *International Mathematics Research Notices* 2018, 24 (05 2017), 7555–7605. https://doi.org/10.1093/imrn/rnx082 arXiv:http://oup.prod.sis.lan/imrn/article-pdf/2018/24/7555/27176676/rnx082.pdf

[37] Roberto Imbuzeiro Oliveira. 2012. On the coalescence time of reversible random walks. *Trans. Am. Math. Soc.* 364, 4 (2012), 2109–2128.

[38] Roberto Imbuzeiro Oliveira. 2013. Mean field conditions for coalescing random walks. *Ann. Probab.* 41, 5 (2013), 3420–3461.

[39] Roberto I. Oliveira and Yuval Peres. 2019. Random walks on graphs: new bounds on hitting, meeting, coalescing and returning. In *Proceedings of the Sixteenth Workshop on Analytic Algorithmics and Combinatorics, ANALCO 2019, San Diego, CA, USA, January 6, 2019*. 119–126.

[40] David Peleg. 2002. Local majorities, coalitions and monopolies in graphs: a review. *Theoretical Computer Science* 282, 2 (2002), 231–257.

[41] Yuval Peres and Perla Sousi. 2015. Mixing Times are Hitting Times of Large Sets. *Journal of Theoretical Probability* 28, 2 (2015), 488–519.

[42] Etienne Perron, Dinkar Vasudevan, and Milan Vojnović. 2009. Using Three States for Binary Consensus on Complete Graphs. In *Proc. INFOCOM*. 2527–2535.

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. Here, we recall that a non-lazy random walk moves at each step to a neighbour chosen uniformly at random. A lazy random walk stays at the current vertex with probability $1/2$, and otherwise moves to a neighbour chosen uniformly at random.

Lemma A.1 (cf. [33, Equation 12.13]). *Let P be the transition matrix of a reversible Markov chain with state space Ω. Then the following statements hold: If P is irreducible, then for any two states $x, y \in \Omega$,*

$$p^t_{x,y} \leqslant \pi(y) + \sqrt{\frac{\pi(y)}{\pi(x)}} \cdot \lambda^t,$$

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

Note that the t-step distribution q^t of a non-lazy random walk can be expressed by $q^t = q^{t-1} \cdot Q$, with transition matrix $Q := D^{-1} \cdot A$.

Next define the normalised Laplacian matrix by

$$L = I - D^{-1/2}AD^{-1/2} = I - \widetilde{A},$$

where $\widetilde{A} = D^{1/2}QD^{-1/2}$. Hence,

$$D^{-1/2}\widetilde{A} = QD^{-1/2}$$

It is easy to verify that the eigenvalues of $\widetilde{A}$ and Q are the identical. Indeed, if v satisfies $\widetilde{A}v = \lambda v$, then

$$D^{-1/2}\widetilde{A}v = D^{-1/2}\lambda v = \lambda(D^{-1/2}v),$$

but also

$$(D^{-1/2}\widetilde{A})v = (QD^{-1/2})v = Q(D^{-1/2}v).$$

Hence $\widetilde{A}$ and Q have the set of same eigenvalues $1 = \lambda_1 \geqslant \lambda_2 \geqslant \cdots \geqslant \lambda_n \geqslant -1$. Further, let $0 = \mu_1 \leqslant \mu_2 \leqslant \cdots \leqslant \mu_n$ be the eigenvalues of L. Let $\mu := \max\{|1-\mu_2|, |1-\mu_n|\}$.

Lemma A.2 ([13, Theorem 1.2, (2)]). *Let G be a graph sampled from $G(n,p)$ with $p = \omega(\log n/n)$ and $\hat{d} = (n-1)\cdot p$. Then, with probability $1-o(1)$,*

$$\mu = O\left(\frac{1}{\sqrt{\hat{d}}}\right).$$

By definition of the normalised Laplacian, we have the relation

$$\lambda_i = 1 - \mu_i,$$

and hence the conclusion of the lemma above implies

$$\max\{\lambda_2, |\lambda_n|\} = O\left(\frac{1}{\sqrt{\hat{d}}}\right).$$

Corollary A.3. *Let n_0 be a sufficiently large constant. Let H_{n^*} be a $G(n^*, p)$ graph with $p \geqslant (n^*)^{-c}$ for some constant $0 < c < 1$. Then, there exists a constant $C = C(c) > 0$ such that with probability $1 - o(1)$, the lazy random walk on $G(n^*, p)$ satisfies $\max_{n^* \geqslant n_0}\{t_{\mathrm{sep}}(H_{n^*}), t_{\mathrm{mix}}(H_{n^*})\} \leqslant C$.*

Proof. Using a Chernoff bound, it follows that for the random graph H_{n^*}, the following two conditions hold: $\Delta \leqslant (1+o(1))p \cdot n^*$, $\delta \geqslant (1-o(1))p \cdot n^*$. Further, by Lemma A.2,

$$\max\{\lambda_2, |\lambda_{n^*}|\} = O(1/\hat{d}).$$

Hence by Lemma A.1, it is easy to see that for $t \geqslant c_1$ being a sufficiently large constant $c_1 > 0$, for all pairs of vertices $x, y \in V$,

$$\left|q^t_{x,y} - \pi(y)\right| \leqslant 1/(n^*)^2.$$

Recall $p^t_{x,\cdot}$ is t-step distribution of a lazy-walk, and $d(t) = \max_u \|p^t_{u,\cdot} - \pi\|_{\mathrm{TV}}$., as well as, $\bar{d}(t) = \max_{u,v} \|p^t_{u,\cdot} - p^t_{v,\cdot}\|_{\mathrm{TV}}$. By conditioning on the number of loops taken by a t-step random walk,

$$p^t_{x,\cdot} = \sum_{s=0}^{t} \underbrace{\binom{t}{s} 2^{-t}}_{:=a_s} q^s_{x,\cdot}.$$

Further, by the triangle inequality,

$$\|p^t_{x,\cdot} - \pi\|_{\mathrm{TV}} \leqslant \sum_{s=0}^{t} a_s \cdot \|q^s_{x,\cdot} - \pi\|_{\mathrm{TV}}.$$

By choosing $t = c_2$ for another constant $c_2 > c_1$, by using a Chernoff bound for a Binomial random variable, it follows that

$$\sum_{s=0}^{c_1-1} a_s = \mathbb{P}\left[\, \mathsf{Bin}(c_2, 1/2) \leqslant c_1 \,\right] \leqslant 1/e^2.$$

Applying this to the previous inequality yields,

$$\begin{aligned} d(t) = \|p^t_{x,\cdot} - \pi\|_{\mathrm{TV}} &\leqslant \sum_{s=0}^{c_1-1} a_s \cdot 1 + \max_{s \geqslant c_1} \|p^s_{x,\cdot} - \pi\|_{\mathrm{TV}} \\ &\leqslant 1/e^2 + 2n^* \cdot 1/(n^*)^2 < 1/(2e), \end{aligned}$$

for sufficiently large n^*. By [33, Lemma 4.10], we have $\bar{d}(t) \leqslant 2d(t)$, and hence $t_{\mathrm{mix}}(H_{n^*}) \leqslant c_2$ and $t_{\mathrm{sep}}(H_{n^*}) \leqslant 4t_{\mathrm{mix}}(H_{n^*}) \leqslant 4c_2$, which completes the proof. □

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 [27, 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 $\{1, \ldots, n\}$ so that $\mathbb{P}[\, Z_1 = i \,] \geqslant \varepsilon \cdot \mathbb{P}[\, Z_2 = i \,]$ for every $1 \leqslant i \leqslant n$. Then, there is a coupling $(\widetilde{Z}_1, \widetilde{Z}_2)$ of (Z_1, Z_2) and an event $\mathcal{E}$ with $\mathbb{P}[\, \mathcal{E} \,] \geqslant \varepsilon$ so that*

$$\mathbb{P}\left[\, \widetilde{Z}_1 = i \,\middle|\, \mathcal{E} \,\right] = \mathbb{P}\left[\, \widetilde{Z}_2 = i \,\right] \qquad \text{for every } 1 \leqslant i \leqslant n.$$

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

$$\widetilde{Z}_1 = \widetilde{Z}_2 = i, \qquad \text{if } i \text{ satisfies } \varepsilon \textstyle\sum_{k=1}^{i-1} \mathbb{P}[\, Z_2 = k \,] \leqslant U < \varepsilon \sum_{k=1}^{i} \mathbb{P}[\, Z_2 = k \,].$$

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

The following lemma is an immediate consequence of Lemma A.4.

Lemma A.5. *Consider a random walk $(X_t)_{t \geqslant 0}$, starting from an arbitrary but fixed vertex x_0. Then there is an event $\mathcal{E}$ with $\mathbb{P}[\, \mathcal{E} \,] \geqslant 1 - 1/e$ such that for every $v \in V$,*

$$\mathbb{P}\left[\, X_{4t_{\mathrm{mix}}} = v \,\middle|\, \mathcal{E} \,\right] = \pi(v)$$

Proof. Consider the random walk $(X_t)_{t \geqslant 0}$ after step $s := t_{\mathrm{sep}} \leqslant 4t_{\mathrm{mix}}$. By definition of t_{sep}, $p^s_{u,v} \geqslant (1 - 1/e)\pi(v)$. Applying Lemma A.4, where Z_1 is the distribution given by $p^t_{u,v}$ 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 $X_{4t_{\mathrm{mix}}}$ as well. □

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,*

$$\max\{(1/e)t_{\text{mix}}, t^{\pi}_{\text{meet}}\} \leqslant t_{\text{meet}} \leqslant \frac{2}{(1-1/e)^2} \cdot \left(4t_{\text{mix}} + 2t^{\pi}_{\text{meet}}\right),$$

and similarly, $t_{\text{coal}} \leqslant 4 \cdot (4t_{\text{mix}} + 2t^{\pi}_{\text{coal}})$., where t^{π}_{coal} is the expected coalescence time when all n walks start at positions drawn independently from the stationary distribution π.

Proof. We begin by proving the lower bound on t_{meet}. First, consider a coupling of two random walks, $(X_t)_{t\geqslant 0}$ and $(Y_t)_{t\geqslant 0}$ that are run for $t = e \cdot t_{\text{meet}}$ time-steps and start from arbitrary vertices x_0 and y_0. As soon as the two random walks are at the same vertex for some time s, all future transitions are identical, i.e., $X_t = Y_t$ for all $t \geqslant s$. Then, we have

$$\bar{d}(t) = \max_{u,v} \|p^t_{u,\cdot} - p^t_{v,\cdot}\|_{\text{TV}} \leqslant \mathbb{P}\left[\, \cap_{s=0}^{t} \{X_s \neq Y_s\} \,\right] \leqslant \frac{1}{e},$$

where the first inequality is due to the coupling method [33, Equation 5.5] and the second inequality follows by Markov's inequality. The above inequality implies $t_{\text{mix}} \leqslant e \cdot t_{\text{meet}}$. Furthermore, $t^{\pi}_{\text{meet}} \leqslant t_{\text{meet}}$ holds by definition, and the lower bound follows.

For the upper bound, we divide the two random walks into consecutive epochs of length $\ell := 4t_{\text{mix}} + 2t^{\pi}_{\text{meet}}$. 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 $(X_t)_{t\geqslant 0}$ 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 $s + 2t^{\pi}_{\text{meet}}$ 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^{π}_{coal} is shown in exactly the same way. ◻

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

Note that $t_{\text{hit}} \geqslant \frac{2}{\pi_{\min}} - 2$ is tight in the sense that the hitting time of the clique is indeed $2(n-1) = \frac{2}{\pi_{\min}} - 2$ 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 $(X_t)_{t\geqslant 0}$ starting at u. Then it is well-known (cf. [2]) that for the first return $\tau^+(u,u) := \min\{t > 0 : X_t = u, X_0 = u\}$, we have $\mathbb{E}\left[\, \tau^+(u,u) \,\right] = 1/\pi(u) = 1/\pi_{\min}$. By conditioning on the first step of the random walk, we obtain

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

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 $t_{\text{hit}}(v,u) \geqslant \frac{2}{\pi_{\min}} - 2$, and the first claim follows. The second part follows from observing that if $n \geqslant 2$ we have $\pi_{\min} \leqslant 1/2$ and thus $t_{\text{hit}} \geqslant \frac{2}{\pi_{\min}} - 2 \geqslant \frac{1}{\pi_{\min}} \geqslant n$, where the last inequality follows from the simple pigeon hole principle. □

Observation A.8. *Consider two random walks $(X_t)_{t\geqslant 0}$ and $(Y_t)_{t\geqslant 0}$ 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 $\mathbb{P}[\, Z_1 \geqslant 1 \mid Z \geqslant 1\,] \geqslant \frac{1}{2}$.*

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

Lemma A.9. *Let $\mathcal{F}_{t-1}$ be the filtration at time $t-1$.[12] Let $(X_t)_{t\geqslant 0}$ be a stochastic process satisfying (i) $\mathbb{E}[\, X_t | \mathcal{F}_{t-1}\,] \leqslant \beta \cdot X_{t-1}$, for some $\beta < 1$, and (ii) $X_t \geqslant 0$ for all $t \geqslant 0$. Let $\tau(g) := \min\{t \geqslant 0 | X_t \leqslant g\}$ for $g \in (0, |X_0|)$, then*

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

Proof. By the iterative law of expectation, we have

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

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

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

Therefore,

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

□

[12] Informally speaking, is the history of all random decisions up to time t.