

Tight Bounds for Coalescing-Branching Random Walks on Regular Graphs

Petra Berenbrink, George Giakkoupis, Peter Kling

► **To cite this version:**

Petra Berenbrink, George Giakkoupis, Peter Kling. Tight Bounds for Coalescing-Branching Random Walks on Regular Graphs. SODA 2018 - Proceedings of the 29th ACM-SIAM Symposium on Discrete Algorithms, Jan 2018, New Orleans, United States. pp.1715-1733. <hal-01635757>

HAL Id: hal-01635757

<https://hal.inria.fr/hal-01635757>

Submitted on 15 Nov 2017

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Tight Bounds for Coalescing-Branching Random Walks on Regular Graphs

Petra Berenbrink^{*†}

George Giakkoupis[‡]

Peter Kling^{*}

Abstract

A *Coalescing-Branching Random Walk* (CoBRA) is a natural extension to the standard random walk on a graph. The process starts with one pebble at an arbitrary node. In each round of the process every pebble splits into k pebbles, which are sent to k random neighbors. At the end of the round all pebbles at the same node coalesce into a single pebble. The process is also similar to randomized rumor spreading, with each informed node pushing the rumor to k random neighbors each time it receives a copy of the rumor. Besides its mathematical interest, this process is relevant as an information dissemination primitive and a basic model for the spread of epidemics.

We study the *cover time* of CoBRA walks, which is the time until each node has seen at least one pebble. Our main result is a bound of $O(\phi^{-1} \log n)$ rounds with high probability on the cover time of a CoBRA walk with $k = 2$ on any regular graph with n nodes and conductance ϕ . This bound improves upon all previous bounds in terms of graph expansion parameters (Dutta et al. [13], Mitzenmacher et al. [27], Cooper et al. [8, 9]). Moreover, we show that for any connected regular graph the cover time is $O(n \log n)$ with high probability, independently of the expansion. Both bounds are asymptotically tight.

Since our bounds coincide with the worst-case time bounds for PUSH rumor spreading on regular graphs until all nodes are informed, this raises the question whether CoBRA walks and PUSH rumor spreading perform similarly in general. We answer this negatively by separating the cover time of CoBRA walks and the rumor spreading time of PUSH by a super-polylogarithmic factor on a family of tree-like regular graphs.

1 Introduction

A *Coalescing-Branching Random Walk* with branching factor k , also called k -CoBRA walk, is a generalization of the simple random walk on a graph and was introduced by Dutta et al. [12, 13]. The process starts with one particle on an arbitrary node of a connected graph.

At the beginning of each round, every particle splits into k parts which move independently to a randomly chosen neighbor (branching). Particles that arrive at a node during the same round merge into one particle, such that there is at most one particle on each node after that round (coalescing). Note that a branching factor of $k = 1$ gives the standard random walk on graphs.

Random walks and their variants are a natural choice for processes that spread information in networks in a simple and robust manner. In particular, random walk particles are stateless and spread information without knowledge about the global network topology. A standard random walk takes polynomial time to disseminate information from one to all nodes (*cover time*). In particular, on any connected graph of n nodes it needs at least $\Omega(n \log n)$ rounds in expectation. In the worst case, this number can become quadratic on regular graphs [24] and cubic on non-regular graphs [17]. One way to speed up the cover time is the use of *parallel random walks*, where multiple particles start from the same source and perform *independent* random walks. On many graphs, this leads to a speedup that equals the number of particles [15]. CoBRA walks represent another natural idea to speed up the cover time. In contrast to parallel random walks, the number of active particles in CoBRA walks changes during the process. This allows CoBRA walks to “adapt” to the graph topology and limits the number of neighbors contacted by any one node during a round. However, complex dependencies arise between the particles’ positions, which complicates their analysis.

Another process that is closely related to CoBRA walks is *randomized rumor spreading* [19, 25]. Here, a rumor is planted on an arbitrary node of the graph. In the simplest protocol, the so-called PUSH rumor spreading, every node that has learned the rumor pushes the rumor to a randomly chosen neighbor in each round. Notice the similarity to k -CoBRA walks, which can be seen as pushing the rumor to k randomly chosen neighbors instead of one. However, unlike in rumor spreading, where a node pushes the rumor in *every* round after it gets informed, in CoBRA walks a node pushes the rumor only if it received (a copy of) the rumor by one of its neighbors in the previous round. It is immediate that the *rumor spreading time*

^{*}University of Hamburg.

[†]Simon Fraser University.

[‡]INRIA, Rennes. The author was partially supported by ANR Project NDFusion (ANR-16-TERC-0007), and ANR Project PAMELA (ANR-16-CE23-0016-01).

of PUSH, i.e., the number of rounds before all nodes are informed, is at most k times larger than the cover time of COBRA walks. In particular, for constant k , the rumor spreading time of PUSH is asymptotically a lower bound on the cover time of COBRA walks. However, COBRA walks require potentially fewer messages.

COBRA walks are also relevant to modeling the spread of epidemic diseases. In particular, the dual process of a k -COBRA walk, called the *biased infection with persistent source* process, or k -BIPS, was introduced by Cooper et al. [8]. It is a discrete epidemic process of the *susceptible infected susceptible* (SIS) type. In the k -BIPS process, there is a persistent source of the virus, which remains infected at all times. In each round, every node (infected or not) samples k random neighbors. If some of these neighbors are infected, u becomes infected at the end of the round. Any node (except for the persistent source) which did not sample an infected node will become uninfected at the end of the round. As noted in [8], a similar process has been used to model the spread of the BVDV (Bovine Viral Diarrhea Virus) in animals [23]. It was shown in [8] that the time to infect the whole graph in the k -BIPS process is of the same order as the cover time of the k -COBRA walk.

Several recent papers have been devoted to the analysis of the cover time of 2-COBRA walks. Dutta et al. [12, 13] showed an $O(\log^2 n)$ bound that holds w.h.p. (with high probability¹) for regular constant-degree graphs with sufficiently high expansion parameters. They also provided tight bounds for the complete graph, trees, and the d -dimensional grid. Mitzenmacher et al. [27] gave a w.h.p. bound of $O(d^4 \cdot \phi^{-2} \cdot \log^2 n)$ for d -regular graphs with conductance ϕ , and general bounds of $O(n^{2-1/d} \cdot \log n)$ for d -regular graphs and $O(n^{11/4} \cdot \log n)$ for non-regular graphs. The last two bounds show that COBRA walks beat the corresponding $\Theta(n^2)$ and $\Theta(n^3)$ worst-case cover time of random walks [17, 24]. Cooper et al. [8] showed that the cover time on d -regular graphs is bounded by $O((1-\lambda)^{-3} \cdot \log n)$ w.h.p., where $1-\lambda$ is the *absolute spectral gap* of the graph's random walk matrix. Very recently, Cooper et al. [9] complemented this result by a bound of $O((d^2 + d/(1-\lambda)) \cdot \log n)$. They also proved a bound of $O(m + \Delta^2 \cdot \log n)$ w.h.p. for arbitrary graphs with m edges and maximum degree Δ . Both [8, 9] base their analysis on the duality between the k -COBRA process and the k -BIPS process, while earlier papers use techniques from the analysis of parallel random walks.

Despite the intense research on the problem, no general tight bounds are known so far for the cover

time of COBRA walks. This is in contrast to the rumor spreading process, for which tight upper bounds are known in terms of various expansion measures [20, 21], and a tight upper bound of $O(n \log n)$ is known to hold for all graphs [19].

Contribution. In this paper, we provide the first tight bounds for the cover time of COBRA walks on general regular graphs. Our first main result is the following:

THEOREM 1.1. *For any connected regular graph G with n nodes and conductance ϕ , the cover time of a 2-COBRA walk on G is $O(\phi^{-1} \cdot \log n)$ w.h.p.*

This result improves all previous bounds on regular graphs in terms of expansion: the $O(d^4 \cdot \phi^{-2} \cdot \log^2 n)$ bound in [27], the $O((1-\lambda)^{-3} \cdot \log n)$ bound in [8], and the $O((d^2 + d/(1-\lambda)) \cdot \log n)$ bound in [9]. Note that ϕ is an upper bound on $1-\lambda$, but the latter can be much smaller than ϕ . It is $\lambda := \max\{\lambda_2, |\lambda_d|\}$, where λ_i is the i th largest eigenvalue of the random walk transition matrix on G . The well-known Cheeger's inequality states that $(1-\lambda_2)/2 \leq \phi \leq \sqrt{2(1-\lambda_2)}$, thus $1-\lambda_2$ can be up to quadratic in ϕ . However, $1-|\lambda_d|$ can be even smaller than that. For example, in regular bipartite expanders, ϕ is lower bounded by a positive constant, while $\lambda_d = -1$ and, thus, $1-\lambda = 0$.

The bound of [Theorem 1.1](#) is tight since there are d -regular graphs with diameter $\Omega(\phi^{-1} \cdot \log n)$ for a wide range of parameters d and ϕ [6]. We note that it is not possible to show an $O(\phi^{-1} \cdot \log n)$ upper bound for 2-COBRA walks for *non-regular* graphs; the n -node star graph has conductance $\phi = 1$, but the expected cover time is $\Theta(n \log n)$.

Our second result is a general upper bound that applies to all n -node regular graphs, independently of their expansion.

THEOREM 1.2. *For any connected regular graph G with n nodes, the cover time of a 2-COBRA walk on G is $O(n \log n)$ w.h.p.*

This result improves the previous bounds of $O(n^{2-1/d} \log n)$ for d -regular graphs [27], and the recent $O(n^2 \log n)$ bound which however applies also to non-regular graphs [9]. Our bound is tight, as demonstrated by the example of an n -node graph obtained by two cliques of size $n/2$, by removing one of the edges in each clique, and adding two new edges between the cliques such that the graph remains regular. It is an intriguing open problem whether the $O(n \log n)$ bound holds also for non-regular graphs.

Our bounds in [Theorems 1.1](#) and [1.2](#) coincide with the worst-case time bounds for PUSH rumor spreading

¹The expression *with high probability* refers to a probability of $1 - n^{-\Omega(1)}$.

on regular graphs until all nodes are informed. This raises the question whether COBRA walks and PUSH rumor spreading perform similarly in general. We answer this question negatively, by showing that the rumor spreading time of PUSH is smaller by a super-polylogarithmic factor than the cover time of a COBRA walk on a family of tree-like regular graphs.

THEOREM 1.3. *There is a d -regular graph G with n nodes and $d := 2^{\sqrt{\log n \cdot \log \log n}}$, such that the rumor spreading time of PUSH on G is $O(d \log n)$ w.h.p., while the expected cover time of a 2-COBRA walk on G is $d \cdot 2^{\Omega(\sqrt{\log n / \log \log n})}$.*

Techniques. We summarize our main techniques for the proof of [Theorems 1.1](#) and [1.2](#). Our analysis exploits the duality between COBRA walks and the BIPS process [\[8\]](#), which allows us to analyze the infection time of BIPS instead of the cover time of COBRA walks.

The proof of [Theorem 1.1](#) is conceptually similar to the analysis for the $O(\phi^{-1} \log n)$ bound for rumor spreading [\[20\]](#), but with some important differences. Both the analysis of rumor spreading and the analysis of BIPS look at the expected per round change in the number of informed nodes and infected nodes, respectively. This expectation is then bounded in terms of the conductance. However, for rumor spreading it is straightforward to see that the expected number of informed nodes increases by a factor of $1 + \phi(S)$ (until informing half of the graph), where $\phi(S)$ is the conductance of the set S of currently informed nodes. For BIPS this is not the case: We show that the expected number of informed nodes increases by a factor of $1 + \phi$, where ϕ is the *minimum* graph conductance. To show this bound we relate the conductance $\phi(I)$ of the current set I of infected nodes to that of a different set, I' . The set I' is obtained by an iterative process of adding to I nodes from $V \setminus I$ with many neighbors in I . A second difference is that the above mentioned increase by a factor of $1 + \phi$ is not maintained until half of the graph is infected; the expected increase starts dropping from around the point when $n/3$ nodes are infected until $n/2$ are infected (see [Figure 1](#)). Another difference compared to the rumor spreading analysis is that BIPS is a non-progressive process, such that it is naturally more difficult to bound its evolution. To cope with this we prove our own set of Chernoff-like bounds for stochastic processes.

To prove [Theorem 1.2](#), roughly speaking, we show the following: If the size of the cut between infected and uninfected nodes is $\Omega(d)$, then the expected increase in the number of infected nodes in a round of BIPS is $\Omega(1)$ (under certain assumptions). If there are fewer, say ad edges (for a small, possibly sub-constant $\alpha > 0$),

we show that in expectation it takes at least $1/\alpha^2$ rounds before some of the nodes in I become uninfected, while within $1/\alpha$ rounds the number of infected nodes increases by $\Omega(1/\alpha)$. It follows that we have an increase of $\Omega(1)$ on the number of infected nodes *on average per round*. [Theorem 1.2](#) then follows by a version of Wald's theorem. In the formal proof, we need to address also the case in which the cut between infected and uninfected nodes is large but the expected increase is $o(1)$. This happens if the endpoints of the edges in the cut form an almost bipartite graph. We show that this case can be reduced to the case of a small cut.

2 Other Related Work

In the following we survey some related work from the random walk and rumor spreading literature.

Random walks provide a fundamental mathematical model for many processes in networks. For example, they are used to analyze diffusion load balancing in networks [\[30\]](#), to analyze rumor spreading [\[19\]](#), and as a model to spread infections in social networks [\[29\]](#). It is known that the cover time of a random walk on any graph lies between $\Omega(n \log n)$ and $O(n^3)$ [\[17, 18\]](#). Other works have studied how to speed up random walks by using information about node degrees [\[10, 22\]](#), by using neighborhood exploration [\[3\]](#), or by avoiding already visited edges [\[4\]](#). Parallel random walks were first studied in [\[1\]](#). Tight bounds on the cover time for parallel random walks on several graph classes were given in [\[7, 15\]](#). Biased random walks were considered in [\[2\]](#), where the authors studied the effect of a non-uniform choice between neighbors.

Randomized rumor spreading protocols are simple and efficient protocols for information dissemination in networks. The three standard rumor spreading strategies are PUSH, PULL, and PUSH-PULL [\[11\]](#). As already mentioned, COBRA walks share some similarities with the PUSH protocol. In [\[19\]](#), it was proven that the worst-case PUSH rumor spreading time is $\Theta(n \log n)$ w.h.p. In [\[20\]](#), it was shown that the worst-case PUSH-PULL rumor spreading time is $\Theta(\phi^{-1} \log n)$ w.h.p. The same bound holds for both PUSH-only and PULL-only rumor spreading on *regular* graphs. In [\[21\]](#), a bound of $O(\log n \cdot \log(\Delta)/\alpha)$ was proven for PUSH-PULL rumor spreading in arbitrary graphs with vertex expansion α and maximum degree Δ . In another line of research, [\[5\]](#) and [\[14\]](#) tried to minimize the number of message transmissions for rumor spreading. These protocols tend to restrict communication: nodes are only allowed to communicate a couple of times after receiving the rumor for the first time. In [\[16\]](#), a close relation was established between the cover time of random walks and the broadcast time in the phone-call model. It was shown that,

for any graph with m edges and maximum degree Δ , the cover time of random walks (and, hence, the cover time of CoBRA walks) is at most $O(m/\Delta \cdot \log n)$ times the broadcast time. The authors also showed that this factor is a lower bounded by $\Omega(d^2/(n \cdot \log n))$ for d -regular graphs.

3 Preliminaries

Consider an undirected graph $G = (V, E)$. For a node $u \in V$, let $N(u)$ denote the set of u 's neighbors in G , and let $\deg(u) := |N(u)|$ be u 's degree. Given a set $A \subseteq V$, we define $N_A(u) := N(u) \cap A$, $\deg_A(u) := |N_A(u)|$, and $\gamma_{u,A} := \deg_A(u)/d$. For a pair of sets $A, B \subseteq V$, let $E(A, B) := \sum_{u \in A} \deg_B(u)$. Note that $E(A, B) = E(B, A)$ and that, if $A \cap B = \emptyset$, then $E(A, B)$ equals the number of edges connecting A and B . Define the *volume* of $A \subseteq V$ as $\text{vol}(A) := E(A, V) = \sum_{u \in A} \deg(u)$. The *conductance* ϕ of graph G is

$$\phi := \min_{\substack{A \subseteq V \\ 0 < \text{vol}(A) \leq \text{vol}(V)/2}} \frac{E(A, V \setminus A)}{\text{vol}(A)}.$$

k -CoBRA Walk: Assume we are given a connected graph $G = (V, E)$, a number $k \in \mathbb{N}$, and a set $A_0 \subseteq V$ of initially *active* nodes. The k -CoBRA walk is a random process $(A_t)_{t \in \mathbb{N}_0}$ over the state space $\mathcal{P}(V)$. The set A_t is the set of active nodes at the end of round t . In round $t + 1$, each node $u \in A_t$ samples k nodes independently and uniformly at random from $N(u)$ (with replacement). The set of all chosen nodes forms the set A_{t+1} of active nodes at the end of round $t + 1$.

The *cover time* of the k -CoBRA walk is the minimum number of rounds after which each node has been active at least once (i.e., $\min \{t : \bigcup_{i=0}^t A_i = V\}$).

k -BIPS Process: Assume we are given a connected graph $G = (V, E)$, a number $k \in \mathbb{N}$, and a set $I_0 \subseteq V$ of initially (and persistently) *infected* nodes. The k -BIPS process is a random process $(I_t)_{t \in \mathbb{N}_0}$ over the state space $\mathcal{P}(V)$. The set I_t is the set of infected nodes at the end of round t . In round $t + 1$, each node $u \in V \setminus I_0$ samples k nodes independently and uniformly at random from $N(u)$ (with replacement). The union of I_0 and the set of all nodes that sampled at least one node from I_t forms the set I_{t+1} of infected nodes at the end of round $t + 1$.

The *infection time* of the k -BIPS process is the minimum number of rounds after which all nodes are infected, i.e., $\min \{t : I_t = V\}$.

The following result from [8] allows us to analyze the infection time of the BIPS process instead of analyzing the cover time of a CoBRA walk directly.

THEOREM 3.1. ([8]) *For any graph $G = (V, E)$ and any nodes $u, v \in V$, the time for a k -CoBRA walk starting from u until v is active for the first time is distributed identically to the time for the k -BIPS process with source v until node u is infected for the first time.*

From [Theorem 3.1](#), and a union bound over all n possible nodes v , it follows that a high probability bound on the infection time of k -BIPS implies a high probability bound on the cover time of k -CoBRA walk.

4 Probability Tools

In this section we provide some probability results that we will use in our analysis in the following sections.

The next lemma is a variant of Wald's equation. The more familiar version of this result is obtained by letting $X_i = 1$ below.

LEMMA 4.1. *Consider a filtration $\mathcal{F} = (\mathcal{F}_i)_{i \in \mathbb{N}}$ and two sequences $(X_i)_{i \in \mathbb{N}}$ and $(Y_i)_{i \in \mathbb{N}}$ of random variables that are adapted to \mathcal{F} . Suppose that X_i and Y_i are bounded for all $i \in \mathbb{N}$, that X_i is \mathcal{F}_{i-1} measurable, and that $\mathbb{E}[Y_i | \mathcal{F}_{i-1}] \geq \alpha \cdot X_i$ for some constant $\alpha > 0$. Let τ be a stopping time with respect to \mathcal{F} and suppose that τ has bounded expectation. Then*

$$\mathbb{E} \left[\sum_{i=1}^{\tau} Y_i \right] \geq \alpha \cdot \mathbb{E} \left[\sum_{i=1}^{\tau} X_i \right].$$

Proof. For $i \in \mathbb{N}$, let $Z_i := Y_i - \alpha X_i$. Observe that $(Z_i)_{i \in \mathbb{N}}$ is a submartingale difference sequence with respect to \mathcal{F} and that Z_i is bounded for all $i \in \mathbb{N}$. Since, additionally, $\mathbb{E}[\tau] < \infty$, we can apply the optional stopping theorem to obtain $\mathbb{E}[\sum_{i=1}^{\tau} Z_i] \geq 0$. Substituting the definition of Z_i and using linearity of expectation yields the desired statement. Note that we can apply the linearity of expectation because $\mathbb{E}[\sum_{i=1}^{\tau} X_i] < \infty$ and $\mathbb{E}[\sum_{i=1}^{\tau} Y_i] < \infty$, since X_i and Y_i are bounded and $\mathbb{E}[\tau] < \infty$. \square

The next Chernoff-like bound gives concentration guarantees for the difference between two sums of (possibly not independent) binary random variables. It follows from a simple generalization of standard Chernoff bounds to the dependent setting. Its proof can be found in [Appendix A](#).

LEMMA 4.2. *Let $n \in \mathbb{N}$. Consider a filtration $\mathcal{F} = (\mathcal{F}_i)_{i=1}^n$ and two sequences $(X_i)_{i=1}^n$ and $(Y_i)_{i=1}^n$ of binary random variables that are adapted to \mathcal{F} . For $i \in \{1, \dots, n\}$, define the random variables*

$$\begin{aligned} X_{[1,i]} &:= \sum_{1 \leq j \leq i} X_j, & P_{[1,i]} &:= \sum_{1 \leq j \leq i} \mathbb{E}[X_i | \mathcal{F}_{i-1}], \\ Y_{[1,i]} &:= \sum_{1 \leq j \leq i} Y_j, & Q_{[1,i]} &:= \sum_{1 \leq j \leq i} \mathbb{E}[Y_i | \mathcal{F}_{i-1}]. \end{aligned}$$

Let $I \subseteq \{1, \dots, n\}$ and suppose there exists a constant $\alpha > 0$ such that $P_{[1,i]} \geq (1 + \alpha) \cdot Q_{[1,i]}$ for all $i \in I$. Then for any $0 < \epsilon < 1$ and $0 < \delta < 1$,

$$\begin{aligned} \Pr [\forall i \in I: X_{[1,i]} - Y_{[1,i]} \\ \geq (1 - \delta) \cdot (P_{[1,i]} - Q_{[1,i]} - \log \epsilon^{-1})] \\ = 1 - O\left(\delta^{-2} \cdot \epsilon^{\frac{\delta^2 \cdot (1 - \delta)}{6 \cdot (1 + 2/\alpha)^2}}\right). \end{aligned}$$

5 Bound with Conductance

Our first main result, [Theorem 1.1](#), follows from the following bound on the infection time of the 2-BIPS process.

THEOREM 5.1. *For any connected regular graph $G = (V, E)$ with n nodes and conductance ϕ , the infection time of the 2-BIPS process on G is $O(\phi^{-1} \cdot \log n)$ w.h.p.*

[Theorem 1.1](#) follows by combining [Theorem 5.1](#) with [Theorem 3.1](#) on the duality between CoBRA walks and the BIPS process. The rest of this section is devoted to the proof of [Theorem 5.1](#).

5.1 Expectation Bounds for BIPS. In this section, we derive lower bounds on the expected change in the number of infected nodes, during a single round of the 2-BIPS process on graph G .

In the following, we assume $G = (V, E)$ is a connected d -regular graph, $I_0 \subseteq V$ is the set of persistently infected nodes, and I_t is the set of infected nodes after the first t rounds of the 2-BIPS process on G .

LEMMA 5.1. *Let $t \in \mathbb{N}$ and $S \subseteq V$. Then*

$$\begin{aligned} \mathbb{E}[|I_{t+1}| - |I_t| \mid I_t = S] \\ = \sum_{u \in V} \gamma_{u,S}(1 - \gamma_{u,S}) + \sum_{u \in I_0} (1 - \gamma_{u,S})^2. \end{aligned}$$

Proof. Suppose $I_t = S$. Recall $\gamma_{u,S} := \deg_S(u)/d$. A node $u \in S \setminus I_0$ becomes uninfected in round $t + 1$ with probability $(1 - \gamma_{u,S})^2$. Similarly, a node $u \in V \setminus S$ becomes infected with probability $1 - (1 - \gamma_{u,S})^2$. Using that, we calculate

$$\begin{aligned} \mathbb{E}[|I_{t+1}| - |I_t| \mid I_t = S] \\ = \sum_{u \in V \setminus S} (1 - (1 - \gamma_{u,S})^2) - \sum_{u \in S \setminus I_0} (1 - \gamma_{u,S})^2 \\ = \sum_{u \in V \setminus S} (2\gamma_{u,S} - \gamma_{u,S}^2) - \sum_{u \in S} (1 - \gamma_{u,S})^2 \\ \quad + \sum_{u \in I_0} (1 - \gamma_{u,S})^2 \\ = \sum_{u \in V \setminus S} \gamma_{u,S} + \sum_{u \in V \setminus S} \gamma_{u,S} \cdot (1 - \gamma_{u,S}) \end{aligned}$$

$$\begin{aligned} - \sum_{u \in S} (1 - \gamma_{u,S})^2 + \sum_{u \in I_0} (1 - \gamma_{u,S})^2 \\ = \frac{1}{d} \cdot \mathbb{E}(V \setminus S, S) + \sum_{u \in V \setminus S} \gamma_{u,S}(1 - \gamma_{u,S}) - \sum_{u \in S} (1 - \gamma_{u,S}) \\ \quad + \sum_{u \in S} \gamma_{u,S} \cdot (1 - \gamma_{u,S}) + \sum_{u \in I_0} (1 - \gamma_{u,S})^2 \\ = \frac{1}{d} \cdot \mathbb{E}(V \setminus S, S) - \frac{1}{d} \cdot \mathbb{E}(S, V \setminus S) \\ \quad + \sum_{u \in V} \gamma_{u,S} \cdot (1 - \gamma_{u,S}) + \sum_{u \in I_0} (1 - \gamma_{u,S})^2 \\ = \sum_{u \in V} \gamma_{u,S} \cdot (1 - \gamma_{u,S}) + \sum_{u \in I_0} (1 - \gamma_{u,S})^2. \end{aligned}$$

This finishes the proof. \square

The next result is an immediate corollary of [Lemma 5.1](#).

LEMMA 5.2. *Let $t \in \mathbb{N}$ and $S \subset V$. Then*

$$\mathbb{E}[|I_{t+1}| - |I_t| \mid I_t = S] \geq 1/(2d).$$

Moreover, if $(S, V \setminus S, E)$ is a bipartite graph, then $\mathbb{E}[|I_{t+1}| - |I_t| \mid I_t = S] \geq 1$.

Proof. Suppose that $I_t = S \subset V$. If $(S, V \setminus S, E)$ is a bipartite graph, then $\gamma_{u,S} = 0$ if $u \in S$, and $\gamma_{u,S} = 1$ if $u \in V \setminus S$. Then [Lemma 5.1](#) implies

$$\mathbb{E}[|I_{t+1}| - |I_t| \mid I_t = S] = 0 + \sum_{u \in I_0} 1 \geq 1.$$

If $(S, V \setminus S, E)$ is not a bipartite graph, then there is a node $u \in V \setminus S$ that has at least one neighbor in each of the sets S and $V \setminus S$. By [Lemma 5.1](#), the contribution of this node to $\mathbb{E}[|I_{t+1}| - |I_t| \mid I_t = S]$ is

$$\gamma_{u,S} \cdot (1 - \gamma_{u,S}) \geq \frac{1}{d} \cdot \left(1 - \frac{1}{d}\right) \geq \frac{1}{2d},$$

finishing the proof. \square

The next key claim provides lower bounds on the first of the two sums in [Lemma 5.1](#), in terms of the conductance ϕ .

LEMMA 5.3. *Let $S \subseteq V$.*

- (a) $\sum_{u \in V} \gamma_{u,S} \cdot (1 - \gamma_{u,S}) \geq \frac{\phi}{4} \cdot \min\{|S|, n - 2|S|\}$.
- (b) *If $|S| = n/2$ and there is a $u \in S$ with $\gamma_{u,S} \geq \frac{1}{2}$, then $\sum_{u \in V} \gamma_{u,S} \cdot (1 - \gamma_{u,S}) \geq \frac{\phi}{2}$.*

Proof. (a) Let $B := \{u \in V \setminus S : \gamma_{u,S} \geq 1/2\}$, and let $b := \min\{|S|, |B|\}$. Define set \hat{S} as the union of S with an arbitrary b -subset of B . We have

$$(5.1) \quad \begin{aligned} \mathbb{E}(\hat{S}, V \setminus \hat{S}) &= \mathbb{E}(S, V \setminus \hat{S}) + \mathbb{E}(\hat{S} \setminus S, V \setminus \hat{S}) \\ &\leq \mathbb{E}(S, V \setminus \hat{S}) + \mathbb{E}(\hat{S} \setminus S, V \setminus S), \end{aligned}$$

where the inequality holds because, by construction of \hat{S} , $V \setminus \hat{S} \subseteq V \setminus S$. On the other hand, by definition of the conductance and since $b \leq |S|$, we have

$$(5.2) \quad \begin{aligned} \mathbb{E}(\hat{S}, V \setminus \hat{S}) &\geq \phi \cdot d \cdot \min\{|\hat{S}|, |V| - |\hat{S}|\} \\ &= \phi \cdot d \cdot \min\{|S| + b, n - |S| - b\} \\ &\geq \phi \cdot d \cdot \min\{|S|, n - 2|S|\}. \end{aligned}$$

From (5.1) and (5.2), $\mathbb{E}(S, V \setminus \hat{S}) + \mathbb{E}(\hat{S} \setminus S, V \setminus S) \geq \phi \cdot d \cdot \min\{|S|, n - 2|S|\}$. Hence, to prove the claim it suffices to show that $\mathbb{E}(S, V \setminus \hat{S}) + \mathbb{E}(\hat{S} \setminus S, V \setminus S) \leq 4d \cdot \sum_{u \in V} \gamma_{u,S}(1 - \gamma_{u,S})$. We do so by distinguishing two cases.

Case 1: $b < s$.

In this case, we have $\gamma_{u,S} \leq 1/2$ for all $u \in V \setminus \hat{S}$. Moreover, by construction we have $\gamma_{u,S} \geq 1/2$ for all $u \in \hat{S} \setminus S$. Using that, we obtain

$$\begin{aligned} &\mathbb{E}(S, V \setminus \hat{S}) + \mathbb{E}(\hat{S} \setminus S, V \setminus S) \\ &= \sum_{u \in V \setminus \hat{S}} \gamma_{u,S} \cdot d + \sum_{u \in \hat{S} \setminus S} (1 - \gamma_{u,S}) \cdot d \\ &\leq 2d \left(\sum_{u \in V \setminus \hat{S}} \gamma_{u,S}(1 - \gamma_{u,S}) + \sum_{u \in \hat{S} \setminus S} \gamma_{u,S}(1 - \gamma_{u,S}) \right) \\ &\leq 2d \cdot \sum_{u \in V} \gamma_{u,S} \cdot (1 - \gamma_{u,S}). \end{aligned}$$

Case 2: $b = s$.

By construction of \hat{S} , we have $\mathbb{E}(\hat{S} \setminus S, V) = b \cdot d$. Together with the case assumption this implies that $\mathbb{E}(S, V \setminus S) \leq \mathbb{E}(S, V) = s \cdot d = b \cdot d = \mathbb{E}(\hat{S} \setminus S, V)$. With this, we calculate

$$\begin{aligned} \mathbb{E}(S, V \setminus \hat{S}) &= \mathbb{E}(S, V \setminus S) - \mathbb{E}(S, \hat{S} \setminus S) \\ &\leq \mathbb{E}(\hat{S} \setminus S, V) - \mathbb{E}(S, \hat{S} \setminus S) \\ &= \mathbb{E}(\hat{S} \setminus S, V) - \mathbb{E}(\hat{S} \setminus S, S) \\ &= \mathbb{E}(\hat{S} \setminus S, V \setminus S). \end{aligned}$$

Then the desired inequality follows:

$$\begin{aligned} \mathbb{E}(S, V \setminus \hat{S}) + \mathbb{E}(\hat{S} \setminus S, V \setminus S) &\leq 2 \cdot \mathbb{E}(\hat{S} \setminus S, V \setminus S) \\ &= 2 \cdot \sum_{u \in \hat{S} \setminus S} d \cdot (1 - \gamma_{u,S}) \\ &\leq 4d \cdot \sum_{u \in \hat{S} \setminus S} \gamma_{u,S} \cdot (1 - \gamma_{u,S}) \\ &\leq 4d \cdot \sum_{u \in V} \gamma_{u,S} \cdot (1 - \gamma_{u,S}). \end{aligned}$$

(b) Let $B := \{u \in S : \gamma_{u,S} \geq 1/2\} \subseteq S$. By the lemma's assumptions, $1 \leq |B| \leq n/2$. Together with $\text{vol}(B) = d \cdot |B|$ and the definition of conductance, this gives

$$\mathbb{E}(B, V \setminus B) \geq \phi \cdot d \cdot |B| \geq d \cdot \phi.$$

We can also bound $\mathbb{E}(B, V \setminus B)$ from above as follows:

$$\begin{aligned} \mathbb{E}(B, V \setminus B) &= \mathbb{E}(B, V \setminus S) + \mathbb{E}(B, S \setminus B) \\ &\leq \mathbb{E}(B, V \setminus S) + \mathbb{E}(S, S \setminus B) \\ &= \sum_{u \in B} d \cdot (1 - \gamma_{u,S}) + \sum_{u \in S \setminus B} d \cdot \gamma_{u,S} \\ &\leq 2 \sum_{u \in B} d \cdot \gamma_{u,S}(1 - \gamma_{u,S}) + 2 \sum_{u \in S \setminus B} d \cdot \gamma_{u,S}(1 - \gamma_{u,S}) \\ &= 2d \sum_{u \in S} \gamma_{u,S} \cdot (1 - \gamma_{u,S}) \\ &\leq 2d \sum_{u \in V} \gamma_{u,S} \cdot (1 - \gamma_{u,S}). \end{aligned}$$

Here, the first inequality uses $B \subseteq S$ and the second inequality uses $\gamma_{u,S} \geq 1/2$ for $u \in B$ and $1 - \gamma_{u,S} \geq 1/2$ for $u \in S \setminus B$. Combining the above lower and upper bounds for $\mathbb{E}(B, V \setminus B)$ proves the claim. \square

Combining Lemmas 5.1 and 5.3, we obtain the following lower bounds on the expected increase in the number of infected nodes in a round (see also Figure 1).

LEMMA 5.4. *Let $S \subseteq V$ and $s := |S|$.*

$$(a) \quad \mathbb{E}[|I_{t+1}| - |I_t| \mid I_t = S] \geq \frac{\phi}{4} \cdot \min\{s, |n - 2s|, n - s\}.$$

$$(b) \quad \text{If } s = n/2, \text{ then } \mathbb{E}[|I_{t+1}| - |I_t| \mid I_t = S] \geq \frac{\phi}{4}.$$

Proof. (a) We have two cases: If $s \leq n/2$, then Lemma 5.3(a) implies that $\sum_{u \in V} \gamma_{u,S} \cdot (1 - \gamma_{u,S}) \geq \frac{\phi}{4} \cdot \min\{s, n - 2s\}$. If $s > n/2$, then Lemma 5.3(a) applied to the set $V \setminus S$ instead of the set S implies

$$\sum_{u \in V} \gamma_{u, V \setminus S} (1 - \gamma_{u, V \setminus S}) \geq \frac{\phi}{4} \cdot \min\{n - s, n - 2(n - s)\},$$

and this is equivalent to

$$\sum_{u \in V} \gamma_{u,S}(1 - \gamma_{u,S}) \geq \frac{\phi}{4} \cdot \min\{n - s, 2s - n\}.$$

Combining both cases with [Lemma 5.1](#) proves (a).

(b) We have two cases: If there is a $u \in S$ with $\gamma_{u,S} \geq 1/2$, [Lemma 5.3\(b\)](#) gives $\sum_{u \in V} \gamma_{u,S}(1 - \gamma_{u,S}) \geq \phi/2$. Otherwise, we have $\gamma_{u,S} \leq 1/2$ for all $u \in I_0 \subseteq S$. This implies $\sum_{u \in I_0} (1 - \gamma_{u,S})^2 \geq 1/4 \geq \phi/4$, where we used $|I_0| \geq 1$ and $\phi \leq 1$. Therefore in both cases,

$$\sum_{u \in V} \gamma_{u,S} \cdot (1 - \gamma_{u,S}) + \sum_{u \in I_0} (1 - \gamma_{u,S})^2 \geq \frac{\phi}{4}.$$

Combining that with [Lemma 5.1](#) proves (b). \square

5.2 Tail Bounds for BIPS. In this section, we derive high probability lower bounds on the increase in the number of infected nodes, over the course of several rounds of the 2-BIPS process on graph G .

The next result is obtained by using [Lemma 4.2](#).

LEMMA 5.5. *For $t_1, t_2 \in \mathbb{N}_0$, let*

$$\Delta_{(t_1, t_2]} := \sum_{t=t_1+1}^{t_2} \mathbb{E}[|I_t| - |I_{t-1}| \mid I_{t-1}].$$

For any $t_0 \in \mathbb{N}_0$, $S \subseteq V$, $0 < \epsilon < 1$, and $0 < \delta < 1$,

$$\begin{aligned} \Pr[\forall t > t_0: |I_t| - |I_{t_0}| \\ &\geq (1 - \delta) (\Delta_{(t_0, t]} - \log \epsilon^{-1}) \mid I_{t_0} = S] \\ &= 1 - O\left(\delta^{-2} \cdot \epsilon^{\delta^2 \cdot (1 - \delta)/150}\right). \end{aligned}$$

Proof. As already mentioned, we will apply [Lemma 4.2](#) to derive the above tail bound. For that, we must decompose the change $|I_t| - |I_{t-1}|$ in the number of infected nodes during a round t , into a difference $\sum_i X_i^{(t)} - \sum_i Y_i^{(t)}$ for suitable binary random variables. The natural choice would be to let $X_u^{(t)} = 1$ if u is uninfected and becomes infected during round t (such that it contributes positively to $|I_t| - |I_{t-1}|$) and, similarly, $Y_u^{(t)} = 1$ if u is infected and becomes uninfected during round t (such that it contributes negatively to $|I_t| - |I_{t-1}|$). Unfortunately, this might yield situations where no suitable α exists, in particular, if both $\sum_u X_u^{(t)}$ and $\sum_u Y_u^{(t)}$ are large but their difference is small. To deal with this problem, we choose more carefully which nodes are seen as contributing positively or negatively to $|I_t| - |I_{t-1}|$. The idea is that a node u with large $\gamma_{u, I_{t-1}}$ is very likely to be infected at the end of round t , so we will consider it as a negative contribution if u

is not infected, while if $\gamma_{u, I_{t-1}}$ is small, we will consider it as a positive contribution if u is infected at the end of the round.

We formalize this idea as follows. For any set $S \subseteq V$, let $\top(S)$ be the set of the $|S|$ (“top”) nodes $u \in V$ with the largest value $\gamma_{u,S}$, breaking ties arbitrarily (but deterministically). For any $t \in \mathbb{N}$ and $u \in V$, define the random variables $X_u^{(t)}$ and $Y_u^{(t)}$ as follows.

- If $u \notin \top(I_{t-1})$: $X_u^{(t)} = 1$ if $u \in I_t$; $X_u^{(t)} = 0$ if $u \notin I_t$; and $Y_u^{(t)} = 0$.
- If $u \in \top(I_{t-1})$: $Y_u^{(t)} = 1$ if $u \notin I_t$; $Y_u^{(t)} = 0$ if $u \in I_t$; and $X_u^{(t)} = 0$.

We can now express $|I_t|$ as

$$\begin{aligned} |I_t| &= \sum_{u \notin \top(I_{t-1})} X_u^{(t)} + \sum_{u \in \top(I_{t-1})} (1 - Y_u^{(t)}) \\ &= \sum_{u \in V} X_u^{(t)} - \sum_{u \in V} Y_u^{(t)} + |\top(I_{t-1})| \\ &= \sum_{u \in V} (X_u^{(t)} - Y_u^{(t)}) + |I_{t-1}|. \end{aligned}$$

Thus

$$(5.3) \quad |I_t| - |I_{t-1}| = \sum_{u \in V} (X_u^{(t)} - Y_u^{(t)}).$$

Next we show

$$(5.4) \quad \sum_{u \in V} \mathbb{E}[X_u^{(t)} \mid I_{t-1}] \geq \frac{3}{2} \cdot \sum_{u \in V} \mathbb{E}[Y_u^{(t)} \mid I_{t-1}].$$

Taking the expectation in (5.3) while conditioning on $I_{t-1} = S \subseteq V$ yields

$$\begin{aligned} (5.5) \quad \sum_{u \in V} \mathbb{E}[X_u^{(t)} - Y_u^{(t)} \mid I_{t-1} = S] \\ &= \mathbb{E}[|I_t| - |I_{t-1}| \mid I_{t-1} = S] \\ &\geq \sum_{u \in V} \gamma_{u,S} \cdot (1 - \gamma_{u,S}), \end{aligned}$$

where the last inequality follows from [Lemma 5.1](#). We also have

$$\sum_{u \in V} \mathbb{E}[Y_u^{(t)} \mid I_{t-1} = S] = \sum_{u \in \top(S)} (1 - \gamma_{u,S})^2.$$

We bound the right side of this identity by distinguishing two cases: If $\gamma_{u,S} \geq 1/2$ for all $u \in \top(S)$, then the right side is at most $\sum_{u \in \top(S)} \gamma_{u,S}(1 - \gamma_{u,S})$. Otherwise, we must have $\gamma_{u,S} < 1/2$ for all $u \in V \setminus \top(S)$, and thus,

$$\begin{aligned} \sum_{u \in \top(S)} (1 - \gamma_{u,S})^2 &\leq \sum_{u \in \top(S)} (1 - \gamma_{u,S}) = |\top(S)| - \sum_{u \in \top(S)} \gamma_{u,S} \\ &= \sum_{u \in V \setminus \top(S)} \gamma_{u,S} \leq 2 \sum_{u \in V \setminus \top(S)} \gamma_{u,S} \cdot (1 - \gamma_{u,S}), \end{aligned}$$

where in the penultimate step we used $\sum_{u \in V} \gamma_{u,S} = |S| = |\mathbb{T}(S)|$, and in the last step we used $1 - \gamma_{u,S} \geq 1/2$. We conclude that in both cases, we have

$$\sum_{u \in V} \mathbb{E}[Y_u^{(t)} \mid I_{t-1} = S] \leq 2 \sum_{u \in V} \gamma_{u,S} \cdot (1 - \gamma_{u,S}).$$

Substituting that to (5.5) and rearranging, gives (5.4).

We now specify the random variable sequences on which we apply Lemma 4.2, and show that they satisfy the necessary conditions. Suppose $I_{t_0} = S$ and assume, without loss of generality, that $V = \{1, \dots, n\}$. For each round $t \geq t_0$ and node $u \in V$ we define the single subscript variables $X_i := X_u^{(t)}$ and $Y_i := Y_u^{(t)}$, where $i = (t - t_0 - 1) \cdot n + u$. Note that the nodes' random choices generate a filtration $\mathcal{F} = (\mathcal{F}_i)_{i \in \mathbb{N}}$ such that $(X_i)_{i \in \mathbb{N}}$ and $(Y_i)_{i \in \mathbb{N}}$ are adapted to \mathcal{F} . Using this notation in (5.3) and summing over all rounds $t_0 < t \leq t_0 + k$ for $k \in \mathbb{N}$, we get

$$(5.6) \quad |I_{t_0+k}| - |I_{t_0}| = \sum_{i=1}^{k \cdot n} (X_i - Y_i).$$

Moreover, since for any $k \in \mathbb{N}_0$ and $k \cdot n < i \leq (k+1) \cdot n$ the random set I_{t_0+k} is fully determined by $\mathcal{F}_{k \cdot n} \subseteq \mathcal{F}_{i-1}$ and, given I_{t_0+k} , the random variables X_i and Y_i are independent of \mathcal{F}_{i-1} , we have $\mathbb{E}[X_i \mid I_{t_0+k}] = \mathbb{E}[X_i \mid \mathcal{F}_{i-1}]$ and $\mathbb{E}[Y_i \mid I_{t_0+k}] = \mathbb{E}[Y_i \mid \mathcal{F}_{i-1}]$. Combining this with (5.3) and (5.4) we get respectively

$$(5.7) \quad \sum_{i=1}^{k \cdot n} \mathbb{E}[X_i - Y_i \mid \mathcal{F}_{i-1}] \\ = \sum_{t=t_0+1}^{t_0+k} \mathbb{E}[|I_t| - |I_{t-1}| \mid I_{t-1}] = \Delta_{(t_0, t_0+k)},$$

and

$$(5.8) \quad \sum_{i=1}^{k \cdot n} \mathbb{E}[X_i \mid \mathcal{F}_{i-1}] \geq \frac{3}{2} \cdot \sum_{i=1}^{k \cdot n} \mathbb{E}[Y_i \mid \mathcal{F}_{i-1}].$$

We now apply Lemma 4.2 to the random variable sequences $(X_i)_{i=1}^{\lambda n}$ and $(Y_i)_{i=1}^{\lambda n}$ for an arbitrary large λ (the results will not depend on λ , so we can let $\lambda \rightarrow \infty$). Using the notation of Lemma 4.2, we can restate (5.6), (5.7), and (5.8), respectively, as

$$\begin{aligned} X_{[1, k \cdot n]} - Y_{[1, k \cdot n]} &= |I_{t_0+k}| - |I_{t_0}|, \\ P_{[1, k \cdot n]} - Q_{[1, k \cdot n]} &= \Delta_{(t_0, t_0+k)}, \\ P_{[1, k \cdot n]} &\geq \left(1 + \frac{1}{2}\right) \cdot Q_{[1, k \cdot n]}, \end{aligned}$$

for $k \in \mathbb{N}$. Thus, for any $0 < \epsilon < 1$ and $0 < \delta < 1$

Lemma 4.2 yields

$$\begin{aligned} \Pr [\forall k: |I_{t_0+k}| - |I_{t_0}| \\ \geq (1 - \delta) \cdot (\Delta_{(t_0, t_0+k)} - \log \epsilon^{-1}) \mid I_t = S] \\ = 1 - O\left(\delta^{-2} \cdot \epsilon^{\frac{\delta^2 \cdot (1-\delta)}{150}}\right), \end{aligned}$$

finishing the proof. \square

The next lemma uses Lemmas 5.4 and 5.5 to show high-probability lower bounds on the spread of infection in various different ‘‘regimes,’’ depending on the number of currently infected nodes.

LEMMA 5.6. *Assume $I_{t_0} = S$, and let $s := |S|$. Let $\sigma := c \cdot \log n$ for a sufficiently large constant c .*

- (a) $\min\{k: |I_{t_0+k}| \geq \min\{s + \sigma, n\}\} = O(\phi^{-1} \log n)$ w.h.p.
- (b) If $s \in [\sigma, \frac{n}{3}]$ then w.h.p. $\min\{k: |I_{t_0+k}| \geq \min\{2s, \frac{n}{3}\}\} = O(\phi^{-1})$.
- (c) If $s \in [\frac{n}{3}, \frac{n}{2} - \sigma]$ then w.h.p. $\min\{k: \frac{n}{2} - |I_{t_0+k}| \leq \frac{n/2-s}{2}\} = O(\phi^{-1})$.
- (d) If $s \in [\frac{n}{2} + \sigma, \frac{2n}{3}]$ then w.h.p. $\min\{k: |I_{t_0+k}| - \frac{n}{2} \geq \min\{2(s - \frac{n}{2}), \frac{n}{6}\}\} = O(\phi^{-1})$.
- (e) If $s \in [\frac{2n}{3}, n - \sigma]$ then w.h.p. $\min\{k: n - |I_{t_0+k}| \leq \frac{n-s}{2}\} = O(\phi^{-1})$.

Proof. Let $\Delta_{(t_0, t]} := \sum_{i=t_0+1}^t \mathbb{E}[|I_i| - |I_{i-1}| \mid I_{i-1}]$, as in Lemma 5.5. By applying Lemma 5.5 with $\epsilon = n^{-c}$ and $\delta = 1/2$, we obtain

$$(5.9) \quad \Pr \left[\forall t \geq t_0: |I_t| - s \geq \frac{1}{2} \cdot (\Delta_{(t_0, t]} - \sigma) \right] \\ = 1 - O\left(n^{-c/1200}\right).$$

We will use this probability bound to prove the lemma's statements.

(a) Let $T := \min\{t \geq 0: |I_t| = n\}$, i.e., T is the total number of rounds before all nodes are infected. If $i \leq T$, Lemma 5.4 implies $\mathbb{E}[|I_i| - |I_{i-1}| \mid I_{i-1}] \geq \phi/4$. Thus, $\Delta_{(t_0, t]} \geq (t - t_0) \cdot \phi/4$, for $t_0 \leq t \leq T$. Combined with (5.9) this yields

$$(5.10) \quad \Pr \left[\forall t \in \{t_0, \dots, T\}: \right. \\ \left. |I_t| - s \geq \frac{1}{2} \left((t - t_0) \frac{\phi}{4} - \sigma \right) \right] \\ = 1 - O\left(n^{-c/1200}\right).$$

For $t = t_0 + 12\sigma \cdot \phi^{-1} = t_0 + O(\phi^{-1} \cdot \log n)$, (5.10) implies that with probability $1 - O(n^{-c/1200})$, we have $|I_t| - s \geq \sigma$ or $t \geq T$. This implies (a).

(b) Let $T_1 := \min\{t \geq t_0: |I_t| \geq \frac{n}{3} \text{ or } |I_t| < s - \frac{\sigma}{2}\}$, i.e., T_1 is the first round after t_0 in which either at least $n/3$ nodes are infected or the number of infected nodes dropped by more than $\sigma/2$ compared to S . If $t_0 < i \leq T_1$, [Lemma 5.4\(a\)](#) implies $\mathbb{E}[|I_i| - |I_{i-1}| \mid I_{i-1}] \geq |I_{i-1}| \cdot \phi/4 \geq (s - \sigma/2) \cdot \phi/4$. Thus, $\Delta_{(t_0, t]} \geq (t - t_0) \cdot (s - \sigma/2) \cdot \phi/4$, for $t_0 \leq t \leq T_1$. Combined with [\(5.9\)](#) this yields

$$(5.11) \quad \Pr \left[\forall t \in \{t_0, \dots, T_1\} : \right. \\ \left. |I_t| - s \geq \frac{1}{2} \left((t - t_0) \left(s - \frac{\sigma}{2} \right) \frac{\phi}{4} - \sigma \right) \right] \\ = 1 - O\left(n^{-c/1200}\right).$$

For $t = t_0 + 24\phi^{-1}$, [\(5.11\)](#) implies that with probability $1 - O(n^{-c/1200})$, we have $|I_t| - s \geq 3s - 2\sigma \geq s$ or $t > T_1$. However, $t > T_1$ implies $|I_t| \geq n/3$ or $|I_t| < s - \sigma/2$. But since $\Delta_{(t_0, t]} \geq 0$ for all $t \geq t_0$ (which follows from [Lemma 5.1](#)), [\(5.9\)](#) implies that with probability $1 - O(n^{-c/1200})$, we have $|I_t| \geq s - \sigma/2$ for all $t \geq t_0$. Together with a union bound, we obtain that with probability $1 - O(n^{-c/1200})$, $|I_t| - s \geq s$ or $|I_t| \geq n/3$. This implies [\(b\)](#).

(c) Let $T_2 := \min\{t \geq t_0: \frac{n}{2} - |I_t| \leq \frac{n/2 - s}{2} \text{ or } |I_t| < s - \frac{\sigma}{2}\}$. If $t_0 < i \leq T_2$, [Lemma 5.4\(a\)](#) implies $\mathbb{E}[|I_i| - |I_{i-1}| \mid I_{i-1}] \geq \min\{|I_{i-1}|, n - 2|I_{i-1}|\} \cdot \phi/4$. By our definition of T_2 , the right side of this inequality is at most $\min\{s - \sigma/2, n/2 - s\} \cdot \phi/4 = (n/2 - s) \cdot \phi/4$, where we used the assumption $n/3 \leq s \leq n/2 - \sigma$. Thus, $\Delta_{(t_0, t]} \geq (t - t_0) \cdot (n/2 - s) \cdot \phi/4$ for $t_0 \leq t \leq T_2$. Combined with [\(5.9\)](#) this yields

$$(5.12) \quad \Pr \left[\forall t \in \{t_0, \dots, T_2\} : \right. \\ \left. |I_t| - s \geq \frac{1}{2} \left((t - t_0) \left(\frac{n}{2} - s \right) \frac{\phi}{4} - \sigma \right) \right] \\ = 1 - O\left(n^{-c/1200}\right).$$

For $t = t_0 + 8\phi^{-1}$, [\(5.12\)](#) implies that with probability $1 - O(n^{-c/1200})$, we have $n/2 - |I_t| \leq \frac{n/2 - s}{2}$ or $|I_t| < s - \sigma/2$. Since with probability $1 - O(n^{-c/1200})$, we have $|I_t| \geq s - \sigma/2$ for all $t \geq t_0$ (as argued earlier in [\(b\)](#)), it follows that with probability $1 - O(n^{-c/1200})$, $n/2 - |I_t| \leq \frac{n/2 - s}{2}$. This implies [\(c\)](#).

(d) Let $T_3 := \min\{t \geq t_0: |I_t| \geq \frac{2n}{3} \text{ or } |I_t| < s - \frac{\sigma}{2}\}$. If $t_0 < i \leq T_3$, [Lemma 5.4\(a\)](#) implies that, since $n/2 < s - \sigma/2 \leq |I_{i-1}| \leq 2n/3$, $\mathbb{E}[|I_i| - |I_{i-1}| \mid I_{i-1}] \geq (2|I_{i-1}| - n) \cdot \phi/4 \geq (2s - \sigma - n) \cdot \phi/4$. Thus, $\Delta_{(t_0, t]} \geq (t - t_0) \cdot (2s - \sigma - n) \cdot \phi/4$ for $t_0 \leq t \leq T_3$. Combined

with [\(5.9\)](#) this yields

$$(5.13) \quad \Pr \left[\forall t \in \{t_0, \dots, T_3\} : \right. \\ \left. |I_t| - s \geq \frac{1}{2} \left((t - t_0)(2s - \sigma - n) \frac{\phi}{4} - \sigma \right) \right] \\ = 1 - O\left(n^{-c/1200}\right).$$

For $t = t_0 + 12\phi^{-1}$, [\(5.13\)](#) implies that with probability $1 - O(n^{-c/1200})$, we have $|I_t| - n/2 \geq 2(s - n/2)$ or $|I_t| \geq 2n/3$ or $|I_t| < s - \sigma/2$. Since with probability $1 - O(n^{-c/1200})$, we have $|I_t| \geq s - \sigma/2$ for all $t \geq t_0$ (as argued in [\(b\)](#)), it follows that with probability $1 - O(n^{-c/1200})$, $|I_t| - n/2 \geq 2(s - n/2)$ or $|I_t| \geq 2n/3$. This implies [\(d\)](#).

(e) Let $T_4 := \min\{t \geq t_0: n - |I_t| \leq \frac{n-s}{2} \text{ or } |I_t| < s - \frac{\sigma}{2}\}$. If $t_0 < i \leq T_4$, [Lemma 5.4\(a\)](#) implies that, since $s - \sigma/2 \leq |I_{i-1}| \leq n - \frac{n-s}{2}$, $\mathbb{E}[|I_i| - |I_{i-1}| \mid I_{i-1}] \geq \min\{2|I_{i-1}| - n, n - |I_{i-1}|\} \cdot \phi/4 \geq \min\{2s - \sigma - n, \frac{n-s}{2}\} \cdot \phi/4 \geq \frac{n-s}{2} \cdot \phi/4$, since we assume $2n/3 \leq s \leq n - \sigma$. Thus, $\Delta_{(t_0, t]} \geq (t - t_0) \cdot \frac{n-s}{2} \cdot \phi/4$ for $t_0 \leq t \leq T_4$. Combined with [\(5.9\)](#) this yields

$$(5.14) \quad \Pr \left[\forall t \in \{t_0, \dots, T_4\} : \right. \\ \left. |I_t| - s \geq \frac{1}{2} \left((t - t_0) \frac{n-s}{2} \cdot \frac{\phi}{4} - \sigma \right) \right] \\ = 1 - O\left(n^{-c/1200}\right).$$

For $r = t_0 + 16\phi^{-1}$, [\(5.14\)](#) implies that with probability $1 - O(n^{-c/1200})$, we have $n - |I_t| \leq \frac{n-s}{2}$ or $|I_t| < s - \sigma/2$. Since with probability $1 - O(n^{-c/1200})$, we have $|I_t| \geq s - \sigma/2$ for all $t \geq t_0$ (as argued in [\(b\)](#)), it follows that with probability $1 - O(n^{-c/1200})$, $n - |I_t| \leq \frac{n-s}{2}$. This implies [\(e\)](#). \square

5.3 Proof of [Theorem 5.1](#). We divide the process into phases, as illustrated in [Figure 1](#). The theorem then follows by combining the bounds we get from [Lemma 5.6](#) for the length of the different phases.

Formally, let $\sigma = \Omega(\log n)$ as defined in [Lemma 5.6](#), and let $\tau(i) := \min\{t: |I_t| \geq i\}$. We consider the following phases.

Phase 1: $0 \rightarrow \tau(\sigma)$

Applying [Lemma 5.6\(a\)](#) with $t_0 = 0$ yields that, w.h.p., $\tau(\sigma) = O(\phi^{-1} \cdot \log n)$. Thus, the total time spend in this phase is $O(\phi^{-1} \cdot \log n)$ w.h.p.

Phase 2: $\tau(\sigma) \rightarrow \tau(n/3)$

Applying [Lemma 5.6\(b\)](#) with $t_0 = \tau(2^j \cdot \sigma)$, for each $j \in \{0, \dots, \lfloor \log \frac{n}{3\sigma} \rfloor\}$ yields that $\tau(\min\{2^{j+1} \cdot \sigma, n/3\}) -$




Figure 1: Illustration of the seven phases used to bound the infection time of 2-BIPS, in the proof of [Theorem 5.1](#). The x -axis shows the number of infected nodes, and the y -axis shows the lower bound on the expected change in the number of infected nodes per round given by [Lemma 5.4](#). We show that each phase lasts for $O(\phi^{-1} \cdot \log n)$ rounds w.h.p.

$\tau(2^j \cdot \sigma) = O(\phi^{-1})$ w.h.p. Thus, the total time spend in this phase is $\lceil \log \frac{n}{3\sigma} \rceil \cdot O(\phi^{-1})$ w.h.p.

Phase 3: $\tau(n/3) \rightarrow \tau(n/2 - \sigma)$

Applying [Lemma 5.6\(c\)](#) with $t_0 = \tau(n/2 - 2^{-j} \cdot n/6)$, for each $j \in \{0, \dots, \lceil \log \frac{n}{6\sigma} \rceil\}$ yields that $\tau(n/2 - 2^{-j} \cdot n/6) - \tau(n/2 - 2^{-j-1} \cdot n/6) = O(\phi^{-1})$ w.h.p. Thus, the total time spend in this phase is $\lceil \log \frac{n}{6\sigma} \rceil \cdot O(\phi^{-1})$ w.h.p.

Phase 4: $\tau(n/2 - \sigma) \rightarrow \tau(n/2 + \sigma)$

Applying [Lemma 5.6\(a\)](#) with $t_0 = n/2 - \sigma$ and $t_0 = n/2$ yields that $\tau(n/2 + \sigma) - \tau(n/2 - \sigma) = O(\phi^{-1} \cdot \log n)$ w.h.p. Thus, the total time spend in this phase is $O(\phi^{-1} \cdot \log n)$ w.h.p.

Phase 5: $\tau(n/2 + \sigma) \rightarrow \tau(2n/3)$

Applying [Lemma 5.6\(d\)](#) with $t_0 = \tau(n/2 + 2^j \cdot \sigma)$, for each $j \in \{0, \dots, \lceil \log \frac{n}{6\sigma} \rceil\}$ yields that $\tau(\min\{n/2 + 2^{j+1} \cdot \sigma, 2n/3\}) - \tau(n/2 + 2^j \cdot \sigma) = O(\phi^{-1})$ w.h.p. Thus, the total time spend in this phase is $\lceil \log \frac{n}{6\sigma} \rceil \cdot O(\phi^{-1})$ w.h.p.

Phase 6: $\tau(2n/3) \rightarrow \tau(n - \sigma)$

Applying [Lemma 5.6\(e\)](#) with $t_0 = \tau(n - 2^{-j} \cdot n/3)$, for each $j \in \{0, \dots, \lceil \log \frac{n}{3\sigma} \rceil\}$ yields that $\tau(n - 2^{-j-1} \cdot n/3) - \tau(n - 2^{-j} \cdot n/3) = O(\phi^{-1})$ w.h.p. Thus, the total time spend in this phase is $\lceil \log \frac{n}{3\sigma} \rceil \cdot O(\phi^{-1})$ w.h.p.

Phase 7: $\tau(n - \sigma) \rightarrow \tau(n)$

Applying [Lemma 5.6\(a\)](#) with $t_0 = n - \sigma$ yields that $\tau(n) - \tau(n - \sigma) = O(\phi^{-1} \cdot \log n)$ w.h.p.. Thus, the total time spend in this phase is $O(\phi^{-1} \cdot \log n)$ w.h.p.

Combining all the above and using a union bound, we obtain that, w.h.p.,

$$\begin{aligned} \tau(n) &= 3 \cdot O(\phi^{-1} \cdot \log n) \\ &\quad + \left(2 \left\lceil \log \frac{n}{3\sigma} \right\rceil + 2 \left\lceil \log \frac{n}{6\sigma} \right\rceil \right) \cdot O(\phi^{-1}) \\ &= O(\phi^{-1} \cdot \log n). \end{aligned}$$

This completes the proof. \square

6 General Bound

Our second main result, [Theorem 1.2](#), follows from the next bound on the infection time of the 2-BIPS process.

THEOREM 6.1. *For any connected d -regular graph $G = (V, E)$ with n nodes, the infection time of the 2-BIPS process is $O(n)$ in expectation and $O(n \log n)$ w.h.p.*

[Theorem 1.2](#) is obtained by combining the above result with [Theorem 3.1](#). The rest of this section is thus devoted to the proof of [Theorem 6.1](#).

We describe a partition of the 2-BIPS process into phases, such that, roughly speaking, the expected increase in the number of infected nodes in a phase is at least proportional to the length of the phase. We then show that an expected number of $O(n)$ rounds suffice to infect all nodes, by using the variant of Wald's theorem we showed in [Lemma 4.1](#). This variant can deal with the complication that the length of a phase is itself a random variable.

We distinguish two types of phases, depending on the expected increase in the number of infected nodes in the first round of the phase. Suppose that a phase starts in round t , and $I_{t-1} = S$. Consider $\mathbb{E}[|I_t| - |I_{t-1}| \mid I_{t-1} = S]$; more precisely, we consider a quantity $\psi(S)$ which is very close to that expectation.

If $\psi(S) = \Omega(1)$ then the phase consists of a single round, round t . This satisfies the requirement that the expected increase in the number of infected nodes in a phase is at least proportional to the phase length.

If $\psi(S) = o(1)$ then [Lemma 5.1](#) implies that one of the following two cases applies: (1) the number of edges in the cut between S and $V \setminus S$ is small, namely $o(d)$; or (2) the number of edges with both endpoints in S or both in $S \setminus V$ is $o(d)$.

If case (1) holds, then the phase lasts until one of the following happens: (i) the number of infected nodes increases by $1/\psi(S)$; or (ii) some node from S becomes uninfected; or (iii) $2/\psi(S)$ rounds have passed since the beginning of the phase. In [Lemma 6.1](#) we show that (i) happens with at least some constant probability, while the probability that (ii) happens and $k \geq 1$ nodes from S are uninfected at the end of the phase is $\psi(S)^{\Omega(k)}$. These two imply that the expected increase in the number of infected nodes in the phase is at least proportional to the length of the phase.

In the above case (1), it is unlikely that some node from S becomes uninfected during the phase, because there are very few edges between S and the rest of the graph. If case (2) holds, it is no longer true that nodes in S are likely to remain infected during the phase. Instead they are likely to switch between the infected and uninfected states in each round. It turns out that

this case can be treated similarly to case (1), if in place of set S we consider the set $\{u: \gamma_{u,S} > 1/2\}$ in odd rounds from the beginning of the phase. The conditions for terminating the phase are the same as the conditions (i)–(iii) from case (1), except that condition (ii) now states that: some node from S is uninfected at the end of an even round from the beginning of the phase, or some node from set $\{u: \gamma_{u,S} > 1/2\}$ is uninfected at the end of an odd round. Using [Lemma 6.1](#) as before, we conclude that the expected increase in the number of infected nodes in the phase is at least proportional to the length of the phase.

We now give the detailed proof. For any $S \subseteq V$, let

$$\psi(S) := \sum_{u \in V} \min\{\gamma_{u,S}, 1 - \gamma_{u,S}\},$$

and

$$\xi(S) := \{u \in V: \gamma_{u,S} > 1/2\}.$$

Note that $\psi(S) = \sum_{u \in \xi(S)} (1 - \gamma_{u,S}) + \sum_{u \in V \setminus \xi(S)} \gamma_{u,S}$. Note also that $\psi(S)$ is within a factor of two of the sum $\sum_{u \in V} \gamma_{u,S} (1 - \gamma_{u,S})$ from [Lemma 5.1](#).

We consider the following partition of the 2-BIPS process into phases. Let ℓ_i , for $i \in \mathbb{N}_0$, denote the total length of the first i phases. So, phase $i \in \mathbb{N}$ consists of the rounds $\ell_{i-1} + 1$ up to ℓ_i . Assuming $I_{\ell_{i-1}} = S$, we define ℓ_i inductively as follows. Let $\epsilon > 0$ be a sufficiently small constant (to be chosen later).

- If $\psi(S) \geq \epsilon$ or $\psi(S) = 0$, then $\ell_i = \ell_{i-1} + 1$, i.e., phase i consists of just a single round.
- If $0 < \psi(S) < \epsilon$, then ℓ_i is the smallest round that is greater than ℓ_{i-1} and satisfies at least one of the conditions below:

- (i) $|I_{\ell_i}| \geq |S| + 1/\psi(S)$;
- (ii) $S' \setminus I_{\ell_i} \neq \emptyset$, where

$$(6.15) \quad S' := \begin{cases} S, & \text{if } \mathbb{E}(S, V \setminus S) < d/2 \\ & \text{or } \ell_i - \ell_{i-1} \text{ is even;} \\ \xi(S), & \text{otherwise.} \end{cases}$$

- (iii) $\ell_i \geq \ell_{i-1} + 2/\psi(S)$.

The next lemma analyzes the case where $0 < \psi(S) < \epsilon$.

LEMMA 6.1. *Consider a phase $i \geq 1$. Let $S \subseteq V$ such that $0 < \psi(S) < \epsilon$, and let S' be defined as in (6.15).*

- (a) $\Pr[|S' \setminus I_{\ell_i}| \geq k \mid I_{\ell_{i-1}} = S] \leq 2\psi(S)^{2k-1}$, for any $k \geq 1$.
- (b) $\Pr[|I_{\ell_i}| \geq |S| + 1/\psi(S) \mid I_{\ell_{i-1}} = S] \geq \epsilon'$, for some constant $\epsilon' > 0$.

6.1 Proof of [Lemma 6.1](#). First we prove the result for the case in which $\mathbb{E}(S, V \setminus S) < d/2$, and then for the case of $\mathbb{E}(S, V \setminus S) \geq d/2$.

The Case of Small Cut: $\mathbb{E}(S, V \setminus S) < d/2$

In this case, we have $S' = S$. Moreover, $\gamma_{u,S} > 1/2$ for all $u \in S$, and $\gamma_{u,S} < 1/2$ for all $u \in V \setminus S$.

We prove (a) first. For a round $t \geq 1$, fix I_{t-1} and assume $S \subseteq I_{t-1}$. We show by induction on $k \geq 0$ that

$$(6.16) \quad \Pr[|S \setminus I_t| \geq k] \leq \psi(S)^{2k}.$$

This holds trivially for $k = 0$, because the right side is 1. Suppose that $\Pr[|S \setminus I_t| \geq k] \leq \psi(S)^{2k}$ for some $k \geq 0$. We show that $\Pr[|S \setminus I_t| \geq k+1] \leq \psi(S)^{2(k+1)}$:

$$\begin{aligned} \Pr[|S \setminus I_t| \geq k+1] &\leq \sum_{u \in S} \Pr[u \notin I_t] \cdot \Pr[|S \setminus \{u\} \setminus I_t| \geq k] \\ &\leq \sum_{u \in S} (1 - \gamma_{u,S})^2 \cdot \Pr[|S \setminus I_t| \geq k] \\ &\leq \sum_{u \in S} (1 - \gamma_{u,S})^2 \cdot \psi(S)^{2k}, \end{aligned}$$

by the induction hypothesis. Also,

$$\sum_{u \in S} (1 - \gamma_{u,S})^2 \leq \left(\sum_{u \in S} (1 - \gamma_{u,S}) \right)^2 \leq \psi(S)^2,$$

where the last relation is obtained by using the case hypothesis, which implies that $1 - \gamma_{u,S} < 1/2$ for all $u \in S$. It follows that $\Pr[|T \setminus I_t| \geq k+1] \leq \psi(S)^{2(k+1)}$, completing the proof of (6.16).

From (6.16) and a union bound, the probability that $|S \setminus I_t| \geq k$ for at least one of the $\ell_i \leq 2/\psi(S)$ rounds of phase i is at most $(2/\psi(S)) \cdot \psi(S)^{2k}$. This implies (a), since $S' = S$.

Next we prove (b). For a round $t \geq 1$, fix I_{t-1} and suppose that $S \subseteq I_{t-1}$ and $|I_{t-1} \setminus S| = k$. Given that, we compute in the conditional distribution of $|I_t \setminus S|$.

First, suppose that $|I_{t-1} \setminus S| = 0$, i.e., $I_{t-1} = S$.

$$(6.17) \quad \begin{aligned} \Pr[|I_t \setminus S| \geq 1] &= 1 - \prod_{u \in V \setminus S} (1 - \gamma_{u,S})^2 \\ &\geq 1 - e^{-\sum_{u \in V \setminus S} 2\gamma_{u,S}} \\ &\geq 1 - e^{-2\psi(S)} \\ &\geq 2\psi(S) - 2\psi(S)^2, \end{aligned}$$

where the penultimate inequality is obtained using the case hypothesis, which implies $\gamma_{u,S} < 1/2$ if $u \in V \setminus S$.

Suppose now that $|I_{t-1} \setminus S| = k \geq 1$. We have

$$\begin{aligned} \mathbb{E}[|I_t \setminus S|] &= \sum_{u \in V \setminus S} \left(1 - (1 - \gamma_{u, I_{t-1}})^2\right) \\ &= \sum_{u \in V \setminus S} \gamma_{u, I_{t-1}} (2 - \gamma_{u, I_{t-1}}). \end{aligned}$$

For any $u \in V \setminus S$,

$$\gamma_{u, I_{t-1}} \leq \psi(S) + k/d,$$

because u has a fraction $\gamma_{u, S} \leq \psi(S)$ of its neighbors in S , and at most k neighbors in $I_{t-1} \setminus S$. Also,

$$\sum_{u \in V \setminus S} \gamma_{u, I_{t-1}} \geq k.$$

Substituting these two bounds above gives

$$\mathbb{E}[|I_t \setminus S|] \geq k \cdot (2 - \psi(S) - k/d) \geq 3k/2,$$

for $\psi(S) < \epsilon \leq 1/4$ and $k \leq d/4$. Then a Chernoff bound gives,

$$(6.18) \quad \begin{aligned} \Pr[|I_t \setminus S| \geq 4k/3] \\ &= \Pr[|I_t \setminus S| \geq (1 - 1/8) \cdot 3k/2] \\ &\geq 1 - e^{-ck}, \end{aligned}$$

for $c := \frac{1}{86} < \frac{(1/8)^2 \cdot 3/2}{2}$.

We now combine the above results to show (b). For $k \geq 0$, let $x_k = \min\{t > \ell_{i-1} : |I_t \setminus S| \geq k\}$, and let $x_{bad} = \min\{t > \ell_{i-1} : S \setminus I_t \neq \emptyset\}$.

Let \mathcal{E}_1 denote the following event: “ $x_1 < x_{bad}$ and $x_1 \leq \frac{1}{\psi(S) - \psi(S)^2}$ ”. The expectation of $\min\{x_1, x_{bad}\}$ is at most $\frac{1}{2\psi(S) - 2\psi(S)^2}$, by (6.17). Thus, by Markov’s inequality, $\min\{x_1, x_{bad}\} \leq \frac{1}{\psi(S) - \psi(S)^2}$, with probability at least $1/2$. Moreover, by (6.16) (applied for $k = 1$) and (6.17), the probability that $x_1 < x_{bad}$ is at least $1 - \frac{\psi(S)^2}{2\psi(S) - 2\psi(S)^2}$. It follows

$$(6.19) \quad \Pr[\mathcal{E}_1] \geq \frac{1}{2} \cdot \left(1 - \frac{\psi(S)}{2 - 2\psi(S)}\right) \geq \frac{1}{3},$$

for $\psi(S) < \epsilon \leq 1/4$.

Suppose now that \mathcal{E}_1 occurs. Let \mathcal{E}_2 be the event: “ $x_{1/c} < x_{bad}$ and in each of the rounds $x_1 + 1$ up to $x_{1/c}$ the number of infected nodes increases by at least one.” Then by (6.18) and (6.16),

$$(6.20) \quad \begin{aligned} \Pr[\mathcal{E}_2] &\geq \prod_{1 \leq k \leq 1/c-1} (1 - e^{-ck} - \psi(S)^2) \\ &\geq \prod_{1 \leq k \leq 1/c-1} ck/2 \\ &\geq (c/2)^{1/c} \cdot (1/c)! \\ &\geq (1/2e)^{1/c}, \end{aligned}$$

where the second inequality above holds for $\psi(S)^2 < \epsilon^2 \leq c/4$ and is obtained using $1 - e^{-ck} \geq ck - (ck)^2/2$;

and the last inequality above is obtained using Stirling’s formula.

Given the event $\mathcal{E}_1 \cap \mathcal{E}_2$, let \mathcal{E}_3 be the event: “ $x_{1/\psi(S)} < x_{bad}$ and in each of the rounds $x_{1/c} + 1$ up to $x_{1/\psi(S)}$ the number of infected nodes $u \in V \setminus S$ increases by at least a factor of $4/3$.” Then by (6.18) and (6.16),

$$(6.21) \quad \begin{aligned} \Pr[\mathcal{E}_3] \\ &\geq \prod_{0 \leq i \leq \log_{4/3}(1/\psi(S)) - 1} (1 - e^{-(4/3)^i} - \psi(S)^2) \\ &\geq 1 - \sum_{i \geq 0} e^{-(4/3)^i} - \psi(S)^2 \cdot \log_{4/3}(1/\psi(S)) \\ &\geq 1/25 - \psi(S)^2 \cdot \log_{4/3}(1/\psi(S)) \\ &\geq 1/30, \end{aligned}$$

for $\psi(S) < \epsilon$, for a small enough constant ϵ .

Finally, we observe that event $\mathcal{E}_1 \cap \mathcal{E}_2 \cap \mathcal{E}_3$ implies $|I_{\ell_i}| \geq |S| + 1/\psi(S)$ and

$$\ell_i - \ell_{i-1} \leq \frac{1}{\psi(S) - \psi(S)^2} + \frac{1}{c} + \log_{4/3}(1/\psi(S)),$$

which is less than $2/\psi(S)$ for $\psi(S) < \epsilon$, for a small enough constant ϵ . Finally, from (6.19), (6.20), (6.21),

$$\Pr[\mathcal{E}_1 \cap \mathcal{E}_2 \cap \mathcal{E}_3] \geq \frac{(1/2e)^{1/c}}{90}.$$

It follows that the same lower bound applies also to $\Pr[|I_{\ell_i}| \geq |S| + 1/\psi(S)]$. This completes the proof of (b).

The Case of Large Cut: $\mathbb{E}(S, V \setminus S) \geq d/2$

The analysis for this case is a straightforward adaptation of the proof for the case of $\mathbb{E}(S, V \setminus S) < d/2$ (presented earlier), given the following claim which relates S and $\xi(S)$. Note that the claim does not assume that $\mathbb{E}(S, V \setminus S) \geq d/2$.

CLAIM 6.1. *For any $S \subseteq V$, if $\psi(S) < 1/2$, then $\xi(\xi(S)) = S$, $\psi(\xi(S)) = \psi(S)$, and $|\xi(S)| = |S|$.*

Proof. Suppose, for contradiction, that there is some $u \in S \setminus \xi(\xi(S))$. Since $u \notin \xi(\xi(S))$, we have that $\gamma_{u, \xi(S)} \leq 1/2$. That is, for at least $d/2$ of the neighbors v of u , we have $v \notin \xi(S)$ and thus $\gamma_{v, S} \leq 1/2$. Moreover, for each of those nodes v , we have $\gamma_{v, S} \geq 1/d$, as they are neighbors of $u \in S$. Counting just the contribution of these nodes v to $\psi(S)$ we obtain

$$\psi(S) \geq (d/2) \cdot (1/d) = 1/2,$$

which contradicts the assumption that $\psi(S) < 1/2$. We conclude that $S \subseteq \xi(\xi(S))$.

Suppose now that there is some $u \in \xi(\xi(S)) \setminus S$; we use a similar argument as above. Since $u \in \xi(\xi(S))$, we have that $\gamma_{u, \xi(S)} > 1/2$. That is, for more than $d/2$ of the neighbors v of u , we have $v \in \xi(S)$ and thus $\gamma_{v, S} > 1/2$. Moreover, for each of those v we have $\gamma_{v, S} \leq 1 - 1/d$, as they are neighbors of $u \notin S$. Counting just the contribution of those nodes v to $\psi(S)$ we obtain

$$\psi(S) > (d/2) \cdot (1/d) = 1/2,$$

which contradicts the assumption that $\psi(S) < 1/2$. Therefore $\xi(\xi(S)) \subseteq S$, and since we have shown earlier that $S \subseteq \xi(\xi(S))$, we conclude that $\xi(\xi(S)) = S$.

Next we prove $\psi(\xi(S)) = \psi(S)$. Define

$$\begin{aligned} \nu(S) := & \{ (u, v) : \{u, v\} \in E, u \in \xi(S), v \in V \setminus S \} \\ & \cup \{ (u, v) : \{u, v\} \in E, u \in V \setminus \xi(S), v \in S \}. \end{aligned}$$

Observe that

$$\psi(S) = |\nu(S)|/d.$$

From the definition of ν and the property $\xi(\xi(S)) = S$ we showed earlier, it follows that $(u, v) \in \nu(\xi(S))$ if and only if $(v, u) \in \nu(S)$, thus

$$|\nu(\xi(S))| = |\nu(S)|.$$

Therefore,

$$\psi(\xi(S)) = |\nu(\xi(S))|/d = |\nu(S)|/d = \psi(S).$$

Last, we prove $|\xi(S)| = |S|$. We have

$$(6.22) \quad \begin{aligned} |\xi(S)| &= \sum_{u \in \xi(S)} \gamma_{u, S} + \sum_{u \in \xi(S)} (1 - \gamma_{u, S}) \\ &\leq |S| + \psi(S). \end{aligned}$$

By substituting S with $\xi(S)$, we get

$$|\xi(\xi(S))| \leq |\xi(S)| + \psi(\xi(S)).$$

Substituting also $\xi(\xi(S)) = S$ and $\psi(\xi(S)) = \psi(S)$ (as shown earlier), yields

$$|S| \leq |\xi(S)| + \psi(S).$$

Combining this with (6.22), and using that $\psi(S) < 1$, we obtain $|\xi(S)| = |S|$. This completes the proof of Claim 6.1. \square

To prove part (a) of Lemma 6.1, we show similar to (6.16), that if $S \subseteq I_{t-1}$ then $\Pr[|\xi(S) \setminus I_t| \geq k \mid I_{t-1}] \leq \psi(S)^{2k}$, and if $\xi(S) \subseteq I_{t-1}$ then $\Pr[|S \setminus I_t| \geq k \mid I_{t-1}] \leq \psi(S)^{2k}$. To show that we also use the properties $\xi(\xi(S)) = S$ and $\psi(\xi(S)) = \psi(S)$, from Claim 6.1.

To prove (b), we show that (6.17) and (6.18) still hold if we replace S by $\xi(S)$, or if we condition on $\xi(S) \subseteq I_{t-1}$ instead of $S \subseteq I_{t-1}$. For that again we use that $\xi(\xi(S)) = S$ and $\psi(\xi(S)) = \psi(S)$, and also that $|\xi(S)| = |S|$, by Claim 6.1. The rest of the proof is essentially the same as for the case of $E(S, V \setminus S) < d/2$. \square

6.2 Proof of Theorem 6.1. We apply Lemma 4.1 for the random variable sequences $(X_i)_{i \in \mathbb{N}}$ and $(Y_i)_{i \in \mathbb{N}}$, where X_i is an upper bound on the length $\ell_i - \ell_{i-1}$ of phase i , and Y_i is the increase in the number of infected nodes during that phase. Formally, for $i \in \mathbb{N}$,

$$\begin{aligned} X_i &:= \begin{cases} 2/\psi(I_{\ell_{i-1}}), & \text{if } 0 < \psi(I_{\ell_{i-1}}) < \epsilon; \\ 1, & \text{otherwise,} \end{cases} \\ Y_i &:= |I_{\ell_i}| - |I_{\ell_{i-1}}|. \end{aligned}$$

Let $T := \min \{ t : I_t = V \}$ be the number of rounds until all nodes are infected (i.e., the infection time), and let $\tau := \min \{ i : I_{\ell_i} = V \}$ be the number of phases until all nodes are infected. Then

$$(6.23) \quad \sum_{i=1}^{\tau} X_i \geq T, \quad \text{and} \quad \sum_{i=1}^{\tau} Y_i = n - 1.$$

Consider the filtration $(\mathcal{F}_i)_{i \in \mathbb{N}}$ that describes the outcome of the first i phases. Let $i \in \mathbb{N}$ and assume $i \leq \tau$. We distinguish three cases: If $\psi(I_{\ell_{i-1}}) \geq \epsilon$, then $\ell_i = \ell_{i-1} + 1$, and

$$\begin{aligned} \mathbb{E}[Y_i \mid \mathcal{F}_{i-1}] &\geq \sum_{u \in V} \gamma_{u, I_{\ell_{i-1}}} (1 - \gamma_{u, I_{\ell_{i-1}}}) \\ &\geq \psi(I_{\ell_{i-1}})/2 \\ &\geq \epsilon/2, \end{aligned}$$

where the first inequality follows from Lemma 5.1, and the second inequality follows from the definition of ψ . If $0 < \psi(I_{\ell_{i-1}}) < \epsilon$, then Lemma 6.1 implies

$$\begin{aligned} \mathbb{E}[Y_i \mid \mathcal{F}_{i-1}] &\geq \sum_{k \in \mathbb{N}} (-k) \cdot 2\psi(I_{\ell_{i-1}})^{2k-1} + \frac{\epsilon'}{\psi(I_{\ell_{i-1}})} \\ &\geq - \sum_{k \in \mathbb{N}} k 2\epsilon^{2(k-1)} \psi(I_{\ell_{i-1}}) + \frac{\epsilon'}{\psi(I_{\ell_{i-1}})} \\ &\geq -3\psi(I_{\ell_{i-1}}) + \frac{\epsilon'}{\psi(I_{\ell_{i-1}})} \\ &\geq \frac{\epsilon'}{2\psi(I_{\ell_{i-1}})}, \end{aligned}$$

where the last two inequalities hold if constant ϵ is sufficiently small. The last case we need to consider is the case in which $\psi(I_{\ell_{i-1}}) = 0$. In this case, $\ell_i = \ell_{i-1} + 1$, and $(I_{\ell_{i-1}}, V \setminus I_{\ell_{i-1}}, E)$ is a bipartite graph. Lemma 5.2 then gives $\mathbb{E}[Y_i \mid \mathcal{F}_{i-1}] \geq 1$. Therefore, in all three cases above, $\mathbb{E}[Y_i \mid \mathcal{F}_{i-1}] \geq \alpha \cdot X_i$, for a suitable constant $\alpha > 0$. From Lemma 4.1 then we obtain that $\mathbb{E}[\sum_{i=1}^{\tau} Y_i] \geq \alpha \cdot \mathbb{E}[\sum_{i=1}^{\tau} X_i]$. Substituting to that the equations from (6.23), gives $n - 1 \geq \alpha \cdot \mathbb{E}[T]$, completing the proof of Theorem 6.1. \square

7 Comparison with Rumor Spreading

Our bounds on the cover time of a 2-COBRA walk from [Theorems 1.1](#) and [1.2](#) coincide with the worst-case time that PUSH rumor spreading needs to inform all nodes on a regular graph [[19, 20](#)]. This raises the question of whether 2-COBRA walks and PUSH rumor spreading perform similar in general. Our next theorem, a slightly reformulated version of [Theorem 1.3](#), implies that this is not the case. To show this, we construct a family of tree-like regular graphs for which the cover time of the 2-COBRA walk exceeds the PUSH rumor spreading time by a super-polylogarithmic factor.

THEOREM 7.1. *Let $n \in \mathbb{N}$ be sufficiently large and let $d := 2^{\sqrt{\log n \cdot \log \log n}}$. There is a $(d+1)$ -regular graph with $2n$ nodes such that:*

- (a) *The 2-COBRA walk has an expected cover time of $d \cdot 2^{\Omega(\sqrt{\log n / \log \log n})}$, and*
- (b) *PUSH rumor spreading informs all nodes in time $O(d \cdot \log n)$ w.h.p.*

We prove [Theorem 7.1](#) for a family of *almost* regular graphs $S_{d,h}$, where $d, h \in \mathbb{N}$ and $d \geq 2$. Here, $S_{d,h}$ is obtained from a full d -ary tree of height h by replacing each leaf u by a clique of $d+1$ nodes and adding an edge between u 's parent and one (arbitrary) node of the clique. This graph has $n = (d^{h+1} - 1)/(d - 1)$ nodes and is almost $(d+1)$ -regular (only the root and d nodes per leaf-clique have degree d). To obtain a regular graph from $S_{d,h}$, one can, for example, consider two copies of $S_{d,h}$ and match the two roots as well as the clique nodes of degree d from the two copies. Our analysis immediately generalizes to this $(d+1)$ -regular graph with $2n$ nodes.

We prove that the lower bound from [Theorem 7.1\(a\)](#) holds for the 2-COBRA walk on $S_{d,h}$ ([Lemma 7.2](#)) and that the upper bound from [Theorem 7.1\(b\)](#) holds for PUSH rumor spreading on $S_{d,h}$ ([Lemma 7.1](#)). The latter follows easily from a result by Feige et al. [[19](#)] and is given in the next lemma.

LEMMA 7.1. *Starting from an arbitrary node, the PUSH protocol informs all nodes in $S_{d,h}$ in time $O(d \cdot \log n)$ w.h.p.*

Proof. From [[19, Theorem 2.2](#)] we get a bound of $O(\Delta \cdot (\text{diam}(G) + \log n))$ for PUSH rumor spreading to inform all nodes on any graph with n nodes, maximum degree Δ , and diameter $\text{diam}(G)$. The lemma's statement follows since $\deg(S_{d,h}) = d+1$ and $\text{diam}(S_{d,h}) = 2 \cdot (h+1) = O(\log n)$. \square

The lower bound for 2-COBRA on $S_{d,h}$ is stated in the following lemma.

LEMMA 7.2. *Let $n \in \mathbb{N}$ be sufficiently large, let $d := 2^{\sqrt{\log n \cdot \log \log n}}$, and let $h := \sqrt{\log n / \log \log n} - o(1)$.² The expected cover time of the 2-COBRA walk on $S_{d,h}$ is $d \cdot 2^{\Omega(\sqrt{\log n / \log \log n})}$.*

Our main tool to prove [Lemma 7.2](#) is [Lemma 7.3](#), which provides a recursive lower bound on the time T_ℓ it takes a node u at depth ℓ to activate its parent. Slightly simplified, we show that $T_\ell \approx \Omega(\log d) \cdot T_{\ell+1}$. To build up the intuition behind this lower bound, consider a 2-COBRA walk starting from node u . In any round in which u is active, the probability that it activates its parent is roughly $2/d$. Thus, if u is active for $O(\sqrt{d})$ rounds, the probability that its parent got activated at least once is relatively low (of order $O(1/\sqrt{d})$). Thus, to bound the time until u activates its parent with a good probability, it is sufficient to bound the time until u was active for $\Omega(\sqrt{d})$ rounds. In the first round, the 2-COBRA walk most likely activates two of u 's subtrees. Node u becomes inactive and it takes roughly $T_{\ell+1}/2$ rounds before one of the two active subtrees reactivates u . Each time u is reactivated, the number of active subtrees increases by at most two. With k active subtrees, the time until u is reactivated is roughly $T_{\ell+1}/k$. Thus, it takes about $\sum_{k=1}^{\Theta(\sqrt{d})} T_{\ell+1}/k = \Theta(\log d) \cdot T_{\ell+1}$ rounds for u to be activated for at least $\Theta(\sqrt{d})$ rounds.

The formalization of this proof idea needs some care, since we have to deal, for example, with occasional subtrees that are much faster in reactivating u . Taking such technicalities into account, we can prove the actual, slightly more involved recursive lower bound stated in the following lemma.

LEMMA 7.3. *Consider the values $n, d, h \in \mathbb{N}$ from [Lemma 7.2](#). Let the random variable T_ℓ denote the time until a 2-COBRA walk started from a node u of $S_{d,h}$ at depth ℓ activates its parent for the first time. For any $\ell \in \{1, \dots, h-1\}$, $\Pr[T_\ell \geq \tau_\ell] \geq 1 - d^{-1/2}$, where*

$$\tau_\ell = \begin{cases} \Omega(d \cdot \log d), & \text{if } \ell = h-1; \\ \frac{(1 - o(1)) \log d}{2 \log(\tau_{\ell+1}/d)} \cdot \tau_{\ell+1}, & \text{if } 1 \leq \ell < h-1. \end{cases}$$

It remains to prove [Lemma 7.3](#) and, with its help, [Lemma 7.2](#). We start with the comparatively easier proof of [Lemma 7.2](#), and provide the proof of [Lemma 7.3](#) afterward.

²The $o(1)$ term in the expression for h ensures that the number of nodes equals $n = (d^{h+1} - 1)/(d - 1)$, as required by our construction of $S_{d,h}$.

Proof of Lemma 7.2. By unfolding the recurrence for τ_ℓ in Lemma 7.3, we obtain

$$\tau_\ell = \Omega \left(\frac{d \cdot \log d}{(h-1-\ell)!} \cdot \left(\frac{(1-o(1)) \cdot \log d}{2 \log \log d} \right)^{h-1-\ell} \right).$$

Substituting $\ell = 1$, and using Stirling's formula $n! \leq e \cdot \sqrt{n} \cdot (n/e)^n$, gives

$$\begin{aligned} \tau_1 &= \Omega \left(\frac{d \cdot \log d}{(h-2)!} \cdot \left(\frac{(1-o(1)) \cdot \log d}{2 \log \log d} \right)^{h-2} \right) \\ &= \Omega \left(\frac{d \cdot \log d}{\sqrt{h-2}} \cdot \left(\frac{e \cdot (1-o(1)) \cdot \log d}{2 \log \log d \cdot (h-2)} \right)^{h-2} \right). \end{aligned}$$

Finally, by substituting $d = 2^{\sqrt{\log n \cdot \log \log n}}$ and $h = \sqrt{\log n / \log \log n} - o(1)$, we obtain $\tau_1 \geq d \cdot 2^{\Omega(h)}$. With this, Lemma 7.3 implies $T_1 \geq d \cdot 2^{\Omega(h)}$ with probability $1 - d^{-1/2}$ and, thus, $\mathbb{E}[T_1] \geq d \cdot 2^{\Omega(h)}$. Therefore the same $d \cdot 2^{\Omega(h)}$ lower bound holds for the expected cover time of the 2-CoBRA walk, when the start node is a child of the root in $S_{d,h}$, completing the proof of Lemma 7.2. \square

Proof of Lemma 7.3. In the rest of the analysis, we assume that n (and, thus, also h and d) is at least some suitably large constant. The proof is by induction over ℓ . We start with the inductive step; the base case $\ell = h-1$ is considered at the end of the proof.

For the inductive step, assume there is an $\ell \in \{1, \dots, h-2\}$ for which $\Pr[T_{\ell+1} \geq \tau_{\ell+1}] \geq 1 - d^{-1/2}$. We have to show that $\Pr[T_\ell \geq \tau_\ell] \geq 1 - d^{-1/2}$. To this end, consider a 2-CoBRA walk that starts from a node u at depth ℓ . Let C denote the set of u 's children and let p be u 's parent. By definition of the 2-CoBRA walk, T_ℓ is the first round in which u is active and p is among the 2 samples of u . Without loss of generality, we assume each $v \in C$ becomes *permanently active* as soon as

- v samples its parent u for the first time or
- v is sampled by u for the second time.

Note that this assumption can only speed up the process and, thus, strengthens the lower bound.

For $i \in \mathbb{N}$, let t_i be the i th round in which node u is sampled by at least one of its children. We also set $t_0 := 0$. Thus, for each $i \in \mathbb{N}_0$, node u is active at the beginning of round $t_i + 1$. Define \mathcal{P} as the event that u does not sample its parent p in any of the rounds $t_i + 1$ with $0 \leq i \leq \sqrt{d}/4$. This event implies $T_\ell > t_{\sqrt{d}/4} + 1$. If u is active, the probability that p is not sampled is

$(1 - 1/(d+1))^2$. This yields

$$\begin{aligned} (7.24) \quad \Pr[\mathcal{P}] &= \left(1 - \frac{1}{d+1}\right)^{2 \cdot (\sqrt{d}/4 + 1)} \\ &\geq \left(1 - \frac{1}{d}\right)^{2 \cdot (\sqrt{d}/4 + 1)} \\ &\geq 1 - \frac{1}{2 \cdot \sqrt{d}} - \frac{2}{d}, \end{aligned}$$

where the last step used Bernoulli's inequality.

Define \mathcal{D} as the event that the total number of distinct nodes sampled by u during rounds $t_i + 1$, $0 \leq i \leq \sqrt{d}/4$, is at least $2 \cdot (\sqrt{d}/4 + 1) - \log d$. Observe that u performs in total $N := 2 \cdot (\sqrt{d}/4 + 1)$ independent samples in those rounds. The probability that the i th of these samples equals one of the at most $i-1$ previous distinct samples is at most $(i-1)/(d+1)$. Thus the expected number of samples that equal some of the previous samples is at most $\sum_{i=1}^N \frac{i-1}{d+1} < 1$. A standard Chernoff bound gives

$$(7.25) \quad \Pr[\mathcal{D}] \geq 1 - 2^{-\log d} = 1 - \frac{1}{d}.$$

Assume a child $v \in C$ of u is sampled by u for the first time in round x and for the second time in round y . Moreover, let z be the first round in which v samples u (note that $z > x$). We say v is *fast* if $z < y$ and $z - x < \tau_{\ell+1}$. Consider the event \mathcal{H} that the number of fast (distinct) nodes $v \in C$ sampled by u during rounds $t_i + 1$, $0 \leq i \leq \sqrt{d}/4$, is at most $\log d$. The induction hypothesis states $T_{\ell+1} \geq \tau_{\ell+1}$ with probability at least $1 - d^{-1/2}$. Using a standard Chernoff bound, this implies

$$(7.26) \quad \Pr[\mathcal{H}] \geq 1 - 2^{-\log d} = 1 - 1/d.$$

Let π_v denote the round in which a child $v \in C$ becomes permanently active. That is, v sampled u for the first time or u sampled v for the second time. Consider the list $(v_1, r_1), (v_2, r_2), \dots$ of all tuples $v_i \in C$ and $r_i \geq \pi_{v_i} + 1$ order non-decreasingly by r_i (breaking ties by node IDs). Let

$$a_k := |\{(v_i, r_i) : 1 \leq i \leq k, v_i \text{ samples } u \text{ in round } r_i\}|,$$

for $k \in \mathbb{N}_0$. The probability for v_i to sample u in a round r_i is $(1 - (1 - \frac{1}{d+1})^2)$. Thus, $\mathbb{E}[a_k] = k \cdot (1 - (1 - \frac{1}{d+1})^2) < 2k/d$. Define the event \mathcal{C} that $a_k < 8k/d$ for all $k \geq d \cdot \ln d$. Using the union bound over all k and

applying a standard Chernoff bound for each k gives

$$\begin{aligned}
(7.27) \quad \Pr[\mathcal{C}] &\geq 1 - \sum_{k \geq d \cdot \ln d} \Pr \left[a_k \geq \frac{8k}{d} \right] \geq 1 - \sum_{k \geq d \cdot \ln d} e^{-\frac{2k}{d}} \\
&= 1 - e^{-\frac{2d \cdot \ln d}{d}} \cdot \sum_{i \geq 0} e^{-\frac{2i}{d}} \\
&= 1 - d^{-2} \cdot \frac{1}{1 - e^{-\frac{2}{d}}} \geq 1 - \frac{1}{d}.
\end{aligned}$$

Next we show that the event $\mathcal{P} \cap \mathcal{D} \cap \mathcal{H} \cap \mathcal{C}$ implies $T_\ell \geq \tau_\ell$. Once this is proven, the lemma's statement follows by observing that $\Pr[\mathcal{P} \cap \mathcal{D} \cap \mathcal{H} \cap \mathcal{C}] \geq 1 - d^{-1/2}$, which follows from (7.24) to (7.27) by a union bound. In order to prove that $\mathcal{P} \cap \mathcal{D} \cap \mathcal{H} \cap \mathcal{C}$ implies $T_\ell \geq \tau_\ell$, define s_j for $j \in \mathbb{N}_0$ as the number of the first $j \cdot \tau_{\ell+1}$ rounds in which u is sampled by at least one of its children. By definition (of s_j and t_i) we have $t_{s_j} \leq j \cdot \tau_{\ell+1} < t_{s_j+1}$. We use the following auxiliary result, whose proof can be found at the of this section.

CLAIM 7.1. *Assume the event $\mathcal{P} \cap \mathcal{D} \cap \mathcal{H} \cap \mathcal{C}$ occurs. For any $j \in \mathbb{N}$, either $s_j \leq 33j \cdot (s_{j-1} + \log d) \cdot \tau_{\ell+1}/d$ or the quantity on the right side is at least $\sqrt{d}/4$.*

Now suppose that event $\mathcal{P} \cap \mathcal{D} \cap \mathcal{H} \cap \mathcal{C}$ occurs. By recursively applying **Claim 7.1**, we obtain that the smallest j for which $s_j \geq \sqrt{d}/4$, satisfies $(34j \cdot \tau_{\ell+1}/d)^j \cdot \log d \geq \sqrt{d}/4$, where we used $s_0 = 0$ and that the $\log d$ factor in **Claim 7.1** is negligible except for $j = 1$. Solving the above inequality gives

$$j \geq \frac{\log d/2 - 2 - \log \log d}{\log(\tau_{\ell+1}/d) + \log(34j)} = \frac{(1 - o(1)) \cdot \log d}{2 \log(\tau_{\ell+1}/d)}.$$

It follows

$$t_{\sqrt{d}/4} \geq \frac{(1 - o(1)) \cdot \log d}{2 \log(\tau_{\ell+1}/d)} \cdot \tau_{\ell+1},$$

and, since $T_\ell > t_{\sqrt{d}/4}$ (by event \mathcal{P}), we have

$$T_\ell > \frac{(1 - o(1)) \log d}{2 \log(\tau_{\ell+1}/d)} \cdot \tau_{\ell+1},$$

finishing the induction step.

It remains to consider the base case of the induction, $\ell = h-1$. That is, the start node u is at depth $h-1$. We assume, without loss of generality, that each node $v \in C$ becomes permanently active as soon as it is sampled by u for the *first* time. We will use event \mathcal{P} defined earlier, stating that u does not sample its parent in rounds t_i+1 , $0 \leq i \leq \sqrt{d}/4$.

Given event \mathcal{P} , we lower bound the number of rounds until u is sampled by its children in $\sqrt{d}/4$ different rounds. Each time u gets sampled, it samples

(and, thus, permanently activates) at most two children in the next round. If $2k$ nodes $v \in C$ are active at the beginning of a round, the probability u is activated is

$$1 - \left(1 - \frac{1}{d+1}\right)^{2 \cdot 2k} \leq \frac{4k}{d+1}.$$

It follows that the number of rounds until u is sampled by its children in $\sqrt{d}/4$ rounds stochastically dominates the sum of $\sqrt{d}/4$ geometric random variables, where the k th of these variables has success probability $4k/(d+1)$. This sum has the same distribution as the number of remaining rounds until all coupons are collected, in a coupon collection problem in which all but $\sqrt{d}/4$ out of $d+1$ coupons have been collected and four random coupons are sampled per round. Then standard analyses for coupon collection give that the number of rounds needed is $\Omega(d \cdot \log d)$ with probability $1 - 1/d$. Combined with (7.24), we get $T_{h-1} = \Omega(d \cdot \log d)$ with probability at least $1 - d^{-1/2}$. This completes the proof of **Lemma 7.3**. \square

Proof of Claim 7.1 We prove the statement via induction over $j \in \mathbb{N}$. Define $t := \min\{j \cdot \tau_{\ell+1}, t_{\sqrt{d}/4}\}$ and $s'_j := \min\{s_j, \sqrt{d}/4\}$. Note that, by this definition, we have $t_{s'_j} \leq t < t_{s'_j+1}$.

If $s_{j-1} > \sqrt{d}/4$, then the statement follows trivially from $33j \cdot (s_{j-1} + \log d) \cdot \tau_{\ell+1}/d > s_{j-1} > \sqrt{d}/4$, where the first inequality uses $\tau_{\ell+1} = \Omega(d \cdot \log d)$. So assume $s_{j-1} \leq \sqrt{d}/4$. During the first $(j-1) \cdot \tau_{\ell+1}$ rounds, node u is activated s_{j-1} times (by its children). Being the start node, u is also active in the first round. Moreover, since the event \mathcal{P} holds, u 's parent was not active during these $(j-1) \cdot \tau_{\ell+1}$ rounds and could not sample u . Thus, in total u was active during at most $s_{j-1} + 1$ rounds of the first $(j-1) \cdot \tau_{\ell+1}$ rounds and could activate at most $2(s_{j-1} + 1)$ of its children.

Consider the nodes $v \in C$ that u samples for the first time during the rounds $(j-1) \cdot \tau_{\ell+1} + 1$ up to t . Since event \mathcal{D} holds, at most $\log d$ of them are sampled twice by u . Moreover, by event \mathcal{H} , at most $\log d$ of them are fast. In particular, at most these $2 \log d$ become permanently active before round $j \cdot \tau_{\ell+1} + 1$.

Summing up, at most $2(s_{j-1} + 1) + 2 \log d$ nodes $v \in C$ become permanently active during the first t rounds. This implies that the total number of pairs (v_i, r_i) with $r_i \leq t$ is at most $(2(s_{j-1} + 1) + 2 \log d) \cdot 2t \leq (2(s_{j-1} + 1) + 2 \log d) \cdot 2j \cdot \tau_{\ell+1}$. Then event \mathcal{C} implies that at most $(2(s_{j-1} + 1) + 2 \log d) \cdot 2j \cdot \tau_{\ell+1} \cdot 8/d$ times is node u sampled in the first t rounds by nodes $v \in C$ that are permanently active. Also, at most $2(s_{j-1} + 1) + 2 \log d$ times is u sampled by nodes $v \in C$ before they become permanently active, as each $v \in C$ samples u at most

once before becoming permanently active (precisely, in round π_v). Therefore,

$$\begin{aligned} s'_j &\leq (2(s_{j-1} + 1) + 2 \log d) \cdot 2j \cdot \tau_{\ell+1} \cdot 8/d \\ &\quad + 2(s_{j-1} + 1) + 2 \log d \\ &\leq 33j \cdot (s_{j-1} + \log d) \cdot \tau_{\ell+1}/d, \end{aligned}$$

where the last inequality holds for all sufficiently large d . Since $s'_j = \min\{s_j, \sqrt{d}/4\}$, if the quantity in the last line above is less than $\sqrt{d}/4$, then $s_j = s'_j$, and thus $s_j \leq 33j \cdot (s_{j-1} + \log d) \cdot \tau_{\ell+1}/d$; otherwise, $33j \cdot (s_{j-1} + \log d) \cdot \tau_{\ell+1}/d \geq \sqrt{d}/4$, and the claim holds trivially. This completes the proof of [Claim 7.1](#). \square

APPENDIX

A Proof of [Lemma 4.2](#)

First we show a Chernoff-like bound for a single sum of (possibly not independent) binary random variables.

LEMMA A.1. *Let $n \in \mathbb{N}$. Consider a filtration $\mathcal{F} = (\mathcal{F}_i)_{i=1}^n$ and a sequence $(X_i)_{i=1}^n$ of binary random variables that is adapted to \mathcal{F} . Define the random variables $X_{[1,n]} := \sum_{i=1}^n X_i$ and $M := \sum_{i=1}^n \mathbb{E}[X_i | \mathcal{F}_{i-1}]$. The following statements hold for any $\delta \in (0, 1]$ and $b > 0$:*

- (a) *If M is bounded from above by b then*
 $\Pr[X_{[1,n]} \geq (1 + \delta) \cdot b] \leq e^{-\delta^2 \cdot b/3}$.
- (b) *If M is bounded from below by b then*
 $\Pr[X_{[1,n]} \leq (1 - \delta) \cdot b] \leq e^{-\delta^2 \cdot b/2}$.

Proof. We follow the standard proof technique for Chernoff bounds (see, e.g., [28, Chapter 4.1] or [26, Chapter 4.2]). In the following we only prove (a). The proof of (b) is along the same lines.

For any $t > 0$ we have $\Pr[X_{[1,n]} \geq (1 + \delta) \cdot b] = \Pr[e^{t \cdot X_{[1,n]}} \geq e^{t \cdot (1 + \delta) \cdot b}]$. Markov's inequality yields

$$(A.1) \quad \Pr[X_{[1,n]} \geq (1 + \delta) \cdot b] \leq \frac{\mathbb{E}[e^{t \cdot X_{[1,n]}}]}{e^{t \cdot (1 + \delta) \cdot b}}.$$

Next we show that

$$(A.2) \quad \mathbb{E}[e^{t \cdot X_{[1,n]}}] \leq e^{b \cdot (e^t - 1)}.$$

Once this is proven, combining (A.1) and (A.2) yields $\Pr[X_{[1,n]} \geq (1 + \delta) \cdot b] \leq \frac{e^{b \cdot (e^t - 1)}}{e^{t \cdot (1 + \delta) \cdot b}}$. Choosing $t = \ln(1 + \delta) > 0$ and performing standard calculations (see [26, Theorem 4.4]) gives (a).

It remains to prove (A.2). To this end, we define the sequence $(x_i)_{i=1}^n$ iteratively as follows: Given x_1, \dots, x_{i-1} , let $x_i \in \{0, 1\}$ be such that $\Pr[X_i = x_i | x_1, \dots, x_{i-1}] > 0$ and $\mathbb{E}[e^{t \cdot \sum_{i < j \leq n} X_j} | x_1, \dots, x_i]$ is maximized; here,

conditioning on x_j is a shorthand for conditioning on $X_j = x_j$. For $i \in \{1, \dots, n\}$ let $p_i := \mathbb{E}[X_i | x_1, \dots, x_{i-1}]$. We now argue that for $i \in \{1, \dots, n\}$ we have

$$(A.3) \quad \mathbb{E}\left[e^{t \cdot \sum_{i \leq j \leq n} X_j} \mid x_1, \dots, x_{i-1}\right] \leq e^{p_i \cdot (e^t - 1)} \cdot \mathbb{E}\left[e^{t \cdot \sum_{i < j \leq n} X_j} \mid x_1, \dots, x_i\right].$$

To see this, compute

$$\begin{aligned} &\mathbb{E}\left[e^{t \cdot \sum_{i \leq j \leq n} X_j} \mid x_1, \dots, x_{i-1}\right] \\ &= \Pr[X_i = 1 \mid x_1, \dots, x_{i-1}] \\ &\quad \cdot \mathbb{E}\left[e^{t + t \cdot \sum_{i < j \leq n} X_j} \mid x_1, \dots, x_{i-1}, X_i = 1\right] \\ &\quad + \Pr[X_i = 0 \mid x_1, \dots, x_{i-1}] \\ &\quad \cdot \mathbb{E}\left[e^{t \cdot \sum_{i < j \leq n} X_j} \mid x_1, \dots, x_{i-1}, X_i = 0\right] \\ &= p_i \cdot e^t \cdot \mathbb{E}\left[e^{t \cdot \sum_{i < j \leq n} X_j} \mid x_1, \dots, x_{i-1}, X_i = 1\right] \\ &\quad + (1 - p_i) \cdot \mathbb{E}\left[e^{t \cdot \sum_{i < j \leq n} X_j} \mid x_1, \dots, x_{i-1}, X_i = 0\right] \\ &\leq (1 + p_i \cdot (e^t - 1)) \cdot \mathbb{E}\left[e^{t \cdot \sum_{i < j \leq n} X_j} \mid x_1, \dots, x_{i-1}, x_i\right] \\ &\leq e^{p_i \cdot (e^t - 1)} \cdot \mathbb{E}\left[e^{t \cdot \sum_{i < j \leq n} X_j} \mid x_1, \dots, x_i\right], \end{aligned}$$

where the penultimate inequality holds by our choice of x_i . This proves (A.3).

Applying (A.3) iteratively for $i = 1, \dots, n$ yields

$$\mathbb{E}\left[e^{t \cdot \sum_{1 \leq j \leq n} X_j}\right] \leq \prod_{1 \leq j \leq n} e^{p_j \cdot (e^t - 1)} = e^{\sum_{1 \leq j \leq n} p_j \cdot (e^t - 1)} \leq e^{b \cdot (e^t - 1)},$$

where the last inequality uses $\sum_{1 \leq i \leq n} p_i \leq b$, which follows from the lemma's assumption that $M \leq b$. This yields (A.2) and, as argued above, concludes the proof of (a). \square

We will use the following corollary of [Lemma A.1](#).

LEMMA A.2. *Let $n \in \mathbb{N}$. Consider a filtration $\mathcal{F} = (\mathcal{F}_i)_{i=1}^n$ and a sequence $(X_i)_{i=1}^n$ of binary random variables that is adapted to \mathcal{F} . For $i \in \{1, \dots, n\}$, define the random variables $X_{[1,i]} := \sum_{1 \leq j \leq i} X_j$ and $P_{[1,i]} := \sum_{1 \leq j \leq i} \mathbb{E}[X_j | \mathcal{F}_{j-1}]$. For any $0 < \epsilon < 1$ and $0 < \delta \leq 1$, the following equations hold:*

- (a) $\Pr[\forall i: X_{[1,i]} \leq (1 + \delta) \cdot (P_{[1,i]} + \log \epsilon^{-1})] = 1 - O(\delta^{-2} \cdot \epsilon^{\delta^2/3})$.
- (b) $\Pr[\forall i: X_{[1,i]} \geq (1 - \delta) \cdot (P_{[1,i]} - \log \epsilon^{-1})] = 1 - O(\delta^{-2} \cdot \epsilon^{\delta^2/2})$.

Proof. We only prove (a). Part (b) follows similarly (by using Lemma A.1(b) instead of Lemma A.1(a)). Fix $k \in \mathbb{N}_0$ and $j \in \{1, \dots, n\}$. Define

$$X_j^{(k)} := \begin{cases} X_j, & \text{if } P_{[1,j]} \leq k; \\ 0, & \text{otherwise.} \end{cases}$$

Note that sequence $(X_j^{(k)})_{j=1}^n$ is adapted to \mathcal{F} . Let $\tilde{X}^{(k)} := \sum_{1 \leq j \leq n} X_j^{(k)}$, $\tilde{P}^{(k)} := \sum_{1 \leq j \leq n} \mathbb{E}[X_j^{(k)} \mid \mathcal{F}_{j-1}]$. By definition $\tilde{P}^{(k)} \leq k$, so we can apply Lemma A.1(a) for $b = k + \beta \geq \tilde{P}^{(k)}$, where $\beta \geq 0$ will be specified later. This yields $\Pr[\tilde{X}^{(k)} \geq (1 + \delta) \cdot (k + \beta)] \leq e^{-\delta^2 \cdot (k + \beta)/3}$. Applying a union bound over all $k \in \mathbb{N}_0$ yields

$$\begin{aligned} \text{(A.4)} \quad & \Pr \left[\forall k \in \mathbb{N}_0: \tilde{X}^{(k)} \leq (1 + \delta) \cdot (k + \beta) \right] \\ & \geq 1 - \sum_{k \in \mathbb{N}_0} e^{-\delta^2 \cdot (k + \beta)/3} = 1 - \frac{e^{-\delta^2 \cdot \beta/3}}{1 - e^{-\delta^2/3}} \\ & = 1 - \mathcal{O} \left(\delta^{-2} \cdot e^{-\delta^2 \cdot \beta/3} \right). \end{aligned}$$

Observe that if $\tilde{X}^{(k)} \leq (1 + \delta) \cdot (k + \beta)$ for all $k \in \mathbb{N}_0$, then for any $i \in \{1, \dots, n\}$, we have $X_{[1,i]} \leq (1 + \delta) \cdot (P_{[1,i]} + \beta + 1)$. To see this, fix $P_{[1,i]}$ and choose $k \in \mathbb{N}_0$ such that $k - 1 < P_{[1,i]} \leq k$. Then $X_{[1,i]} \leq \tilde{X}^{(k)}$ by definition of $\tilde{X}^{(k)}$, and thus,

$$X_{[1,i]} \leq \tilde{X}^{(k)} \leq (1 + \delta) \cdot (k + \beta) < (1 + \delta) \cdot (P_{[1,i]} + 1 + \beta).$$

From this and (A.4) we get

$$\begin{aligned} & \Pr \left[\forall i: X_{[1,i]} \leq (1 + \delta) \cdot (P_{[1,i]} + \beta + 1) \right] \\ & \geq \Pr \left[\forall k \in \mathbb{N}_0: \tilde{X}^{(k)} \leq (1 + \delta) \cdot (k + \beta) \right] \\ & = 1 - \mathcal{O} \left(\delta^{-2} \cdot e^{-\delta^2 \cdot \beta/3} \right). \end{aligned}$$

The claim follows by letting $\beta = \log \epsilon^{-1} - 1$. \square

Proof of Lemma 4.2. We apply Lemma A.2(a) and Lemma A.2(b) to $(X_i)_{i=1}^n$ and $(Y_i)_{i=1}^n$, respectively. Thus, for any $0 < \epsilon < 1$ and $0 < \zeta < 1$,

$$\begin{aligned} & \Pr \left[\forall i \in I: X_{[1,i]} \geq (1 - \zeta) \cdot (P_{[1,i]} - \log \epsilon^{-1}) \right] \\ & = 1 - \mathcal{O} \left(\zeta^{-2} \cdot \epsilon^{\zeta^2/2} \right); \\ & \Pr \left[\forall i \in I: Y_{[1,i]} \leq (1 + \zeta) \cdot (Q_{[1,i]} + \log \epsilon^{-1}) \right] \\ & = 1 - \mathcal{O} \left(\zeta^{-2} \cdot \epsilon^{\zeta^2/3} \right). \end{aligned}$$

By the union bound, this implies

$$\begin{aligned} \text{(A.5)} \quad & \Pr \left[\forall i \in I: X_{[1,i]} - Y_{[1,i]} \right. \\ & \geq (1 - \zeta) \cdot (P_{[1,i]} - \log \epsilon^{-1}) \\ & \quad \left. - (1 + \zeta) \cdot (Q_{[1,i]} + \log \epsilon^{-1}) \right] \\ & = 1 - \mathcal{O} \left(\zeta^{-2} \cdot \epsilon^{\zeta^2/2} + \zeta^{-2} \cdot \epsilon^{\zeta^2/3} \right). \end{aligned}$$

By setting $\zeta = \frac{\delta}{1+2/\alpha}$ and $\epsilon = \epsilon^{(1-\delta)/2}$, and using $P_{[1,i]} \geq (1 + \alpha)Q_{[1,i]}$, we calculate

$$\begin{aligned} & (1 - \zeta) \cdot (P_{[1,i]} - \log \epsilon^{-1}) - (1 + \zeta) \cdot (Q_{[1,i]} + \log \epsilon^{-1}) \\ & \geq (1 - \delta) \cdot (P_{[1,i]} - Q_{[1,i]} - \log \epsilon^{-1}), \end{aligned}$$

and

$$\zeta^{-2} \epsilon^{\zeta^2/2} + \zeta^{-2} \epsilon^{\zeta^2/3} \leq 2\zeta^{-2} \epsilon^{\zeta^2/3} = \mathcal{O} \left(\delta^{-2} \epsilon^{\frac{\delta^2(1-\delta)}{6(1+2/\alpha)^2}} \right).$$

Substituting these to (A.5) completes the proof of Lemma 4.2. \square

References

- [1] N. Alon, C. Avin, M. Koucký, G. Kozma, Z. Lotker, and M. R. Tuttle. Many random walks are faster than one. In *Proc. 20th ACM Symposium on Parallelism in Algorithms and Architectures (SPAA)*, pages 119–128, 2008.
- [2] Y. Azar, A. Z. Broder, A. R. Karlin, N. Linial, and S. Phillips. Biased random walks. In *Proc. 24th ACM Symposium on Theory of Computing (STOC)*, pages 1–9, 1992.
- [3] P. Berenbrink, C. Cooper, R. Elsässer, T. Radzik, and T. Sauerwald. Speeding up random walks with neighborhood exploration. In *Proc. 21st ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1422–1435, 2010.
- [4] P. Berenbrink, C. Cooper, and T. Friedetzky. Random walks which prefer unvisited edges: Exploring high girth even degree expanders in linear time. *Random Struct. Algorithms*, 46(1):36–54, 2015.
- [5] P. Berenbrink, R. Elsässer, and T. Friedetzky. Efficient randomised broadcasting in random regular networks with applications in peer-to-peer systems. *Distributed Computing*, 29(5):317–339, 2016.
- [6] F. Chierichetti, S. Lattanzi, and A. Panconesi. Almost tight bounds for rumour spreading with conductance. In *Proc. 42nd ACM Symposium on Theory of Computing (STOC)*, pages 399–408, 2010.
- [7] C. Cooper, A. M. Frieze, and T. Radzik. Multiple random walks in random regular graphs. *SIAM J. Discrete Math.*, 23(4):1738–1761, 2009.
- [8] C. Cooper, T. Radzik, and N. Rivera. The coalescing-branching random walk on expanders and the dual epidemic process. In *Proc. 2016 ACM Symposium on Principles of Distributed Computing (PODC)*, pages 461–467, 2016.

- [9] C. Cooper, T. Radzik, and N. Rivera. Improved cover time bounds for the coalescing-branching random walk on graphs. In *Proc. 29th ACM Symposium on Parallelism in Algorithms and Architectures (SPAA)*, pages 305–312, 2017.
- [10] R. David and U. Feige. Random walks with the minimum degree local rule have $O(n^2)$ cover time. In *Proc. 28th ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1839–1848, 2017.
- [11] A. Demers, D. Greene, C. Hauser, W. Irish, J. Larson, S. Shenker, H. Sturgis, D. Swinehart, and D. Terry. Epidemic algorithms for replicated database maintenance. In *Proc. 6th ACM Symposium on Principles of Distributed Computing (PODC)*, pages 1–12, 1987.
- [12] C. Dutta, G. Pandurangan, R. Rajaraman, and S. Roche. Coalescing-branching random walks on graphs. In *Proc. 25th ACM Symposium on Parallelism in Algorithms and Architectures (SPAA)*, pages 176–185, 2013.
- [13] C. Dutta, G. Pandurangan, R. Rajaraman, and S. Roche. Coalescing-branching random walks on graphs. *ACM Transactions on Parallel Computing (TOPC)*, 2(3):20:1–20:29, 2015.
- [14] R. Elsässer. On the communication complexity of randomized broadcasting in random-like graphs. In *Proc. 18th ACM Symposium on Parallelism in Algorithms and Architectures (SPAA)*, pages 148–157, 2006.
- [15] R. Elsässer and T. Sauerwald. Tight bounds for the cover time of multiple random walks. In *36th International Colloquium on Automata, Languages and Programming (ICALP)*, pages 415–426, 2009.
- [16] R. Elsässer and T. Sauerwald. Cover time and broadcast time. In *Proc. 26th International Symposium on Theoretical Aspects of Computer Science (STACS)*, pages 373–384, 2009.
- [17] U. Feige. A tight upper bound on the cover time for random walks on graphs. *Random Struct. Algorithms*, 6(1):51–54, 1995.
- [18] U. Feige. A tight lower bound on the cover time for random walks on graphs. *Random Struct. Algorithms*, 6(4):433–438, 1995.
- [19] U. Feige, D. Peleg, P. Raghavan, and E. Upfal. Randomized broadcast in networks. *Random Struct. Algorithms*, 1(4):447–460, 1990.
- [20] G. Giakkoupis. Tight bounds for rumor spreading in graphs of a given conductance. In *Proc. 28th International Symposium on Theoretical Aspects of Computer Science (STACS)*, pages 57–68, 2011.
- [21] G. Giakkoupis. Tight bounds for rumor spreading with vertex expansion. In *Proc. 25th ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 801–815, 2014.
- [22] S. Ikeda, I. Kubo, N. Okumoto, and M. Yamashita. Impact of local topological information on random walks on finite graphs. In *Proc. 30th International Colloquium on Automata, Languages and Programming (ICALP)*, pages 1054–1067, 2003.
- [23] G. Innocent, I. Morrison, J. Brownlie, and G. Gettinby. A computer simulation of the transmission dynamics and the effects of duration of immunity and survival of persistently infected animals on the spread of bovine viral diarrhoea virus in dairy cattle. *Epidemiology & Infection*, 119(1):91–100, 1997.
- [24] J. D. Kahn, N. Linial, N. Nisan, and M. E. Saks. On the cover time of random walks on graphs. *Journal of Theoretical Probability*, 2(1):121–128, 1989.
- [25] R. Karp, C. Schindelhauer, S. Shenker, and B. Vocking. Randomized rumor spreading. In *Proc. 41st IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 565–574, 2000.
- [26] M. Mitzenmacher and E. Upfal. *Probability and Computing: Randomized Algorithms and Probabilistic Analysis*. Cambridge University Press, 2005.
- [27] M. Mitzenmacher, R. Rajaraman, and S. Roche. Better bounds for coalescing-branching random walks. In *Proc. 28th ACM Symposium on Parallelism in Algorithms and Architectures (SPAA)*, 2016.
- [28] R. Motwani and P. Raghavan. *Randomized Algorithms*. Cambridge University Press, 1995.
- [29] C. Pu, S. Li, and J. Yang. Epidemic spreading driven by biased random walks. *Physica A: Statistical Mechanics and its Applications*, 432:230–239, 2015.
- [30] T. Sauerwald and H. Sun. Tight bounds for randomized load balancing on arbitrary network topologies. In *Proc. 53rd IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 341–350, 2012.