Spread of Information and Diseases via Random Walks in Sparse Graphs

George Giakkoupis

Inria, Univ Rennes, CNRS, IRISA, Rennes, France
george.giakkoupis@inria.fr

Hayk Saribekyan

University of Cambridge, UK
hs586@cam.ac.uk

Thomas Sauerwald

University of Cambridge, UK
thomas.sauerwald@cl.cam.ac.uk

Abstract

We consider a natural network diffusion process, modeling the spread of information or infectious diseases. Multiple mobile agents perform independent simple random walks on an n-vertex connected graph G. The number of agents is linear in n and the walks start from the stationary distribution. Initially, a single vertex has a piece of information (or a virus). An agent becomes informed (or infected) the first time it visits some vertex with the information (or virus); thereafter, the agent informs (infects) all vertices it visits. Giakkoupis et al. [19] have shown that the spreading time, i.e., the time before all vertices are informed, is asymptotically and w.h.p. the same as in the well-studied randomized rumor spreading process, on any d-regular graph with $d = \Omega(\log n)$. The case of sub-logarithmic degree was left open, and is the main focus of this paper.

First, we observe that the equivalence shown in [19] does not hold for small d: We give an example of a 3-regular graph with logarithmic diameter for which the expected spreading time is $\Omega(\log^2 n/\log\log n)$, whereas randomized rumor spreading is completed in time $\Theta(\log n)$, w.h.p. Next, we show a general upper bound of $\tilde{O}(d \cdot \mathrm{diam}(G) + \log^3 n/d)$,[1] w.h.p., for the spreading time on any d-regular graph. We also provide a version of the bound based on the average degree, for non-regular graphs. Next, we give tight analyses for specific graph families. We show that the spreading time is $O(\log n)$, w.h.p., for constant-degree regular expanders. For the binary tree, we show an upper bound of $O(\log n \cdot \log\log n)$, w.h.p., and prove that this is tight, by giving a matching lower bound for the cover time of the tree by n random walks. Finally, we show a bound of $O(\mathrm{diam}(G))$, w.h.p., for k-dimensional grids ($k \geq 1$ is constant), by adapting a technique by Kesten and Sidoravicius [24, 25].

2012 ACM Subject Classification Mathematics of computing → Stochastic processes; Theory of computation → Random walks and Markov chains; Mathematics of computing → Graph algorithms

Keywords and phrases parallel random walks, information dissemination, infectious diseases

Funding *George Giakkoupis*: Supported in part by ANR Project PAMELA (ANR16-CE23-0016-01).
Hayk Saribekyan: Gates Cambridge Scholarship programme.
Thomas Sauerwald: Supported by the ERC Grant 'Dynamic March'.

1 Introduction

We consider the following natural diffusion process on a connected n-vertex graph G. A collection of mobile agents perform independent parallel (discrete-time) random walks on G, starting from the stationary distribution. Initially, there is a piece of information at some arbitrary source vertex. An agent learns the information the first time it visits some

[1] The tilde notation hides factors of order at most $(\log\log n)^2$.

informed vertex (the vertex may have received the information in the same or a previous round). From that point on, the agent spreads the information to all vertices it visits. We study the time it takes before all vertices have been informed. We will refer to this process as VISIT-EXCHANGE, following the terminology of [19].

The above process suggests a simple message broadcasting algorithm for networks: Vertices correspond to processes, and agents are tokens circulated in the network. In each round, every process sends each of the tokens it received in the previous round to a random neighbor, and if the process knows the message, it transmits the message along with each token. As observed in [19], when the number of agents/tokens is linear in n, this algorithm has similar per round message complexity as standard randomized rumor spreading [17, 23], but in several graphs it outperforms the latter, due to a more fair bandwidth utilization: each edge is equally likely to be used in each round.

A second potential application of VISIT-EXCHANGE is as a basic model for the spread of diseases in populations. One can think of agents as the members of the population, where an infected member can transmit the infection to another either by direct contact, or indirectly. In the latter case a healthy individual contracts the virus by being in a place previously visited by an infected individual [27]. Alternatively, one can think of a larger population, residing on the vertices of the graphs (e.g., vertices are cities), and a few mobile individuals are responsible for transmitting the infection between different cities. Our basic model assumes perfect contagion and no recovery. It is an interesting future direction to analyze a refined model that allows probabilistic transmission and recovery.

Several works have studied the spread of information (or viruses) via mobile agents, performing random walks or more general jump processes, in discrete or continuous time, on various families of graphs [4, 10, 12, 19, 21, 22, 25, 26, 28, 32] (see Section 2 for an overview of this literature). In almost all of these works, the information is transmitted only between agents when they meet at a vertex, and vertices do not store information.

The work closest to ours is [19] (see also [20]). The authors consider VISIT-EXCHANGE with $\Theta(n)$ agents, starting from stationarity, and compare the spreading time to that of randomized rumor spreading [17, 23]. In the latter protocol, information is transmitted between adjacent vertices, without the use of agents, by having each vertex communicate with a random neighbor in each round. It was observed in [19] that there are graphs in which VISIT-EXCHANGE is significantly faster than randomized rumor spreading (logarithmic versus linear spreading time), and examples where the converse is true.

A main result of [19] is that on any d-regular graph with sufficiently large degree $d = \Omega(\log n)$, VISIT-EXCHANGE and randomized rumor spreading have the same asymptotic spreading time. The intuition for this result is the following. We have that: (i) the average number of agents per vertex is constant, since there are $\Theta(n)$ agents in total, (ii) all agents start from stationarity, and (iii) the graph is regular. It follows that, in every round, a constant expected number of agents depart from each vertex, to random neighbors. This should have a similar effect in the spread of information as randomized rumor spreading, where each vertex communicates with a random neighbor in each round.

The intuition above is not hard to formalize, and prove that VISIT-EXCHANGE is at least as fast as rumor spreading asymptotically:[2] If $d \geq c \log n$, for a large enough constant c, a Chernoff bound together with a union bound show that, w.h.p., for every vertex u and round $t \leq \text{poly}(n)$, at least $\Omega(d)$ agents visit the neighbors of u in round $t - 1$. Thus, at least

[2] The proof of the other direction, that rumor spreading is at least as fast as VISIT-EXCHANGE, is significantly more involved.

one agent visits u in round t, with constant probability. This argument, however, does not extend to the case of $d = O(\log n)$. It was thus left as an open problem in [19], whether the same result holds for graphs of degree $d = O(\log n)$.

Our Contribution. First, we answer the above open question from [19] in the negative.

▶ **Observation 1.** *There is a* 3*-regular graph* G *with* n *vertices and diameter* $\Theta(\log n)$*, such that the expected spreading time of* VISIT-EXCHANGE *on* G *is* $\Omega(\log^2 n/\log\log n)$*.*

The spreading time of randomized rumor spreading is $\Theta(\mathrm{diam}(G))$, w.h.p., on any constant degree graph G [17], thus it is logarithmic for the graph above. To simplify the exposition, here we only give an example of a constant-degree, *non-regular* graph G with diameter and spreading time as described in Observation 1. Consider a 3-regular graph R with n vertices and diameter $\Theta(\log n)$ (e.g., a 3-regular expander), and $\sqrt{n}$ path graphs, each of length $\log n/2$. We obtain G by connecting one of the two endpoints of each path graph, to a distinct vertex of R. The diameter of G is clearly logarithmic. The expected spreading time is $\Omega(\log^2 n/\log\log n)$, because with constant probability, at least one path graph P contains no agents initially, and then it takes $\Omega(\log^2 n/\log\log n)$ rounds before the endpoint of P not connected to R gets informed. In Section 3, we replace the paths of length $\log n/2$ with "ladder" graphs to construct a *regular* graph satisfying Observation 1.

A consequence of Observation 1 is that known bounds for rumor spreading do not readily apply to VISIT-EXCHANGE for low-degree regular graphs, thus new bounds are needed. In view of that, we first provide a general upper bound for VISIT-EXCHANGE for regular graphs of degree $d = O(\log n)$, in terms of the graph diameter. Then we provide tight bounds for several interesting graph families. All our results assume that the number of agents is αn, for some arbitrary constant $\alpha > 0$, and the walks start from the stationary distribution. We denote by $T(G)$ the spreading time on graph G. Since all our bounds hold for any source vertex, we do not explicitly specify a source in the notation. Moreover, we omit G and write just T, when the graph is clear from the context. We write *w.h.p. (with high probability)* to denote a probability that is at least $1 - n^{-c}$ for some constant $c > 0$.

▶ **Theorem 2.** *For any* d*-regular graph* G *with* $d = O(\log n)$*,* $T = \tilde{O}(d \cdot \mathrm{diam}(G) + \log^3 n/d)$*, w.h.p., where the tilde notation hides factors of order at most* $(\log\log n)^2$*.*

In the above bound, the dependence on the diameter is best possible (e.g., the spreading time along a cycle of d-cliques is proportional to the path length multiplied by d). An additive term is also needed when the diameter is sub-logarithmic, but it is not clear whether the term $\log^3 n/d$ is tight. Recall that the corresponding upper bound for randomized rumor spreading shown in [17] is $O(d \cdot (\mathrm{diam}(G) + \log n))$. Thus, it would be reasonable to guess that the right additive term is $d \cdot \log n$. However, the example in Observation 1 shows that the term must be at least $\tilde{\Omega}(\log^2 n)$. We conjecture that the tight bound is $\tilde{O}(d \cdot \mathrm{diam}(G) + \log^2 n)$.

The proof of Theorem 2 bounds the time that the information takes to spread along a given (shortest) path in the graph. We divide time into phases of length $\log^2 n$ rounds, and in each phase, we lower-bound the probability that the information spreads along a sub-path of length $\tilde{\Omega}(\log^2 n/d)$. For $d = \omega(\log\log n)$, we show this probability to be $1 - e^{-\Omega(d)}$. Moreover, we ensure that this probability bound holds, essentially, *independently* of previous phases, by considering every other phase. We prove the bound by showing a concentration result on the number of agents at the neighborhood of each individual vertex in the sub-path, at each round of the phase, and then applying a union bound. To boost the above probability to $1 - e^{-\Omega(\log n)}$, we need $\log n/d$ phases, which yields the $\log^3 n/d$ term of the bound. For the case of $d = O(\log \log n)$, we use a similar approach, but argue instead about the number of agents that visit each vertex in the sub-path over an interval of multiple rounds (instead of looking at its neighborhood at each round). The main technical tool we use is an upper bound on the return probability from [30].

For non-regular graphs, a similar analysis as for Theorem 2 yields the following result.

▶ **Theorem 3.** *For any graph G with average degree d_{avg} and minimum degree $d_{\min} = \Omega(d_{\text{avg}})$, $T = O(d_{\text{avg}} \cdot \log^2 n \cdot (\text{diam}(G) + \log n))$, w.h.p.*

Even though this bound is likely not tight, it is interesting because there is no analogue of it for randomized rumor spreading. For example, in the graph consisting of two stars with their centers connected by an edge [19], for which $d_{\text{avg}} = O(1)$, randomized rumor spreading takes linear expected time, whereas VISIT-EXCHANGE takes logarithmic time w.h.p. (and Theorem 3 gives a $\text{poly}(\log n)$ bound).

Next we show that the spreading time on expanders is optimal, i.e., logarithmic.

▶ **Theorem 4.** *For any d-regular expander G with $d \geq 3$ constant, $T = O(\log n)$, w.h.p.*

Unlike the proof of Theorem 2, where we argue about individual vertices, to prove Theorem 4 we argue about the set of all informed vertices at time t, precisely, the subset S_t of informed vertices with at least one uninformed neighbor. By the expansion property, S_t contains at least a constant fraction of all informed vertices. We claim that a constant fraction of vertices in S_t are visited by some agent between rounds t and $t + r$, w.h.p., for any t and a large enough constant r. Since d is constant, this implies that the number of informed vertices increases by a constant factor every r rounds. To prove the above claim, we argue that the probability a given agent visits S_t between t and $t + r$ is proportional to $k = |S_t|$ and r. Thus, S_t is not visited by sufficiently many agents in these r rounds with probability decreasing exponentially in $r \cdot k$. Next, we consider all possible instantiations of S_t, and apply a union bound. Since the set of informed vertices at any time is connected, the number of different instantiations of S_t can be bounded by $d^{\Theta(k)}$. Since d is constant, the claim follows by choosing constant r large enough.

We currently do not know how to extend Theorem 4 to regular expanders of degree $\omega(1) \leq d \leq O(\log n)$ (for $d = \Omega(\log n)$, the result follows from [19]).

Next we study trees. Let $R_{b,h}$ denote a rooted b-ary tree where each vertex at distance less than h from the root has b children and all leaves are at distance h from the root. The total number of vertices is $n = (b^{h+1} - 1)/(b - 1)$.

▶ **Theorem 5.** *For any b-ary tree $R_{b,h}$ with $b \geq 2$, $T = O(h \log h + \log n)$, w.h.p. Furthermore, for the binary tree $R_{2,h}$, $T = \Omega(h \log h) = \Omega(\log n \cdot \log \log n)$, w.h.p.*

Note that the spreading time on $R_{b,h}$ of the push-only version of randomized rumor spreading is $\Theta(b \log n)$, w.h.p. Thus, VISIT-EXCHANGE is slower than push for small b, and faster for larger b. Another interesting implication of Theorem 5 is that the cover time of the tree by n random walks starting from stationarity has a super-linear speedup, compared to the cover time for a single random walk, which is $\Omega(n \log^2 n)$. Our analysis suggests a deeper connection between the cover time (or other quantities) of multiple random walks and the spreading time of VISIT-EXCHANGE, which might deserve further study.

We give now an overview of the proof of Theorem 5, for the binary tree case; the case of $b > 2$ is similar. To prove the upper bound, we fix a path between the root r and a vertex u at distance at most $h - \log h$. We show that information spreads between r and u in at most $O(\log n)$ rounds w.h.p., by showing that agents arrive at each vertex v of the path at roughly constant rate, independently of the other vertices in the path. To achieve this independence, for each v we identify a subset S_v of the descendants of v at distance $\log h$, and count only visits to v by agents that are in S_v a number of $\Theta(\log h)$ rounds ago, and in the meantime have not walked past v. To show a constant rate, instead of $1/\Theta(\log h)$, a careful pipeline argument is used. To bound the time to spread the information in the last $\log h$ levels of the tree, we bound the cover time of a tree of height $\log h$ by h walks starting from the root, which takes $O(h \log h)$ steps w.h.p. (in n). Finally, to show the lower bound of Theorem 5, we bound from below the cover time of the tree by n random walks starting from stationarity.

Last we show that the spreading time on grids is optimal, i.e., asymptotically equal to the diameter. Let $G_{k,n}$ denote the k-dimensional grid with side length $n^{1/k}$ and n vertices in total (for simplicity, we assume $n^{1/k}$ is an integer).

▶ **Theorem 6.** *For any grid graph $G_{k,n}$ that has a constant number of dimensions $k \geq 1$, $T = \Theta(\mathrm{diam}(G))$, w.h.p.*

A weaker version of this result, with additional $\log \log n$ factors, follows from Theorem 2. To get rid of these extra factors, we employ a much more fine-grained analysis.

Our proof of Theorem 6 uses a technique developed by Kesten and Sidoravicius [24, 25], who proved a similar bound for a continuous-time diffusion process, in which information spreads between agents when they meet (it is not stored on vertices). For our discussion here, we assume the 1-dimensional case, i.e., the n-path. We consider a sequence of $\Theta(\log \log n)$ tessellations of space-time (up to time linear in n), where each tessellation consists of square blocks; the length of the block side is constant in the first tessellation, and increases exponentially in each subsequent tessellation. Let Δ be the side length of a block, and let (v, t) be its bottom left corner; i.e., the block contains all points $(v + j, t + j')$, $0 \leq j, j' < \Delta$. Roughly speaking, the block is "good" if a sufficiently large neighborhood of the vertices in the block (namely, vertices $v - 3\Delta$ up to $v + 4\Delta$) is sufficiently densely populated by agents at time $t - \Delta$. This implies that any space-time point in the block has a good probability of containing some agent. Starting from the last tessellation, for which $\Delta = \Theta(\log n)$ and all blocks are good w.h.p., we recursively bound the number of bad blocks in each tessellation, concluding that at most a constant fraction of all blocks in the first tessellation are bad. Moreover, blocks that are far from each other by at least some constant distance (in space-time), satisfy the property of being good independently of one another. We can then use this result to show that any possible information path contains sufficiently many good blocks, which guarantees that information reaches from one end of the n-path to the other. We note that various aspects of our proof are simpler that in the original proof of Kesten and Sidoravicius, mainly because our process stores the information at vertices, resulting in information paths that are easier to analyse.

Road-map. In Section 2 we give an overview of the related work. In Section 3 we provide a graph example that satisfies Observation 1. In Section 4 we prove Theorem 2, in Section 5 we show Theorem 3, and in Section 6 we prove Theorem 4. In Sections 7 and 8 we show respectively the upper and lower bounds of Theorem 5. Finally, in Section 9 we prove Theorem 6.

2 Related Work

Independent parallel random walks have been studied since the late 70s [2], mainly as a way to speed-up cover and hitting times and related graph problems [3, 6, 9, 10, 15, 16]. Similarly, randomized rumor spreading, where information exchange occurs between adjacent vertices (e.g., via push, pull, or push-pull), has been studied for the past 35 years [11, 17, 23], with the more recent results studying the spreading time in social networks [13], and bounds with graph expansion [7].

A closely related diffusion process to ours is the one where information is not stored on vertices, but is transmitted directly between agents when they meet, and initially a single agent is informed. Naturally, the spreading time in this setting is the time until all *agents* are informed. Several works have studied this process [10, 12, 19, 28, 32]. Dimitriou et al. [12], observed that on any graph the expected spreading time is $O(t^* \log m)$, where m is the number of agents (placed at arbitrary vertices, initially), and t^* is the maximum expected meeting time of two walks; this bound is tight for some graphs. Better bounds were also provided for the complete graph and expanders. Cooper et al. [10] showed (among other results) that the expected spreading time on a random d-regular graph converges to $\frac{2n \ln m}{m} \cdot \frac{d-1}{d-2}$, for most starting positions of the m agents. Pettarin et al. [32, 33] considered the k-dimensional grid, $G_{k,n}$, for $k \in \{1, 2\}$, and showed that the spreading time is $\tilde{\Theta}(n/\sqrt{m})$,[3] w.h.p., for m agents starting from stationarity. Lam et al. [28] studied the same problem for $k \geq 3$ dimensions, and showed a phase transition depending on m: for large m the spreading time is $\tilde{\Theta}(n^{1+1/k}/\sqrt{m})$, while for small m it is $\tilde{\Theta}(n/m)$. Giakkoupis et al. [19, 20] considered the process on d-regular graphs, with $m = \Theta(n)$ agents starting from stationarity, and showed that, on any d-regular graph with $d = \Omega(\log n)$, the spreading time is asymptotically at least as large as for VISIT-EXCHANGE, and in some cases strictly larger.

Kesten and Sidoravicius [25, 26] studied a continuous-time variant of the above process on the infinite grid, where the initial number of agents on each vertex is a poisson random variable with constant mean, and the information starts from the origin. They proved a theorem for the shape formed by the contour of the informed agents in the limiting case. In their analysis it is implicit that if the grid is finite, the spreading time is linear in the diameter (see also [21]). Our proof of Theorem 6 uses techniques from their analysis. A very similar process is the frog model, where only informed agents move, while uninformed ones stay at their initial position, until they are hit by an informed agent. At that point they get informed, and start their own walk. This process has been studied on infinite grids [4, 34] and trees [22].

3 Lower Bound for Regular Graphs

We give an example of a graph that satisfies Observation 1. Let $G_{n,\ell,m,k}$ be a 3-regular graph on n vertices constructed as follows. Let H be a ladder graph defined as the Cartesian product of a path graph of k vertices, and a path graph of two vertices (i.e., a single edge). Every vertex of H has degree 3, except for the 4 endpoints that have degree 2. Take a 3-regular graph R of $n - 2mk$ vertices and diameter at most ℓ, and remove an arbitrary set of $2m$ edges (u_i, v_i), $1 \leq i \leq 2m$. Create m copies of H, and denote the four endpoint of the ith copy by $x_{2i-1}, y_{2i-1}, x_{2i}, y_{2i}$, where x_j and y_j are connected by an edge. Then, join every copy of H with R, by adding edges (x_i, u_i) and (y_i, v_i), for all $1 \leq i \leq 2m$. The resulting graph is $G_{n,\ell,m,k}$. By construction, the graph is 3-regular with n vertices. Also,

$$\operatorname{diam}(G_{n,\ell,m,k}) \leq 3\ell + 2k + 2, \qquad (1)$$

[3] The tilde asymptotic notation hides polylogarithmic factors.

because for every edge (u_i, v_i) removed from R, a path $u_i x_i y_i v_i$ of length 3 is created, and the diameter of each copy of H is k.

Recall we assume that the total number of agents in VISIT-EXCHANGE is αn.

▶ **Lemma 7.** *For $m = \lceil\sqrt{n}\rceil$, and $k = \lfloor \log n/(4\alpha)\rfloor$, $\mathbb{E}[T(G_{n,\ell,m,k})] = \Omega(\log^2 n/\log\log n)$.*

Proof. For $i \in \{1, \ldots, m\}$, let U_i be the vertex set of the ith copy of H, and let S_i be the set of its endpoints. The expected number of unique agents that visit S_i in the first $r = \lceil \log^2 n/\log\log n\rceil$ rounds is $4\alpha r$, and since the agents move independently, by a Chernoff bound, with probability at least $1 - 1/n^2$ no more than $8\alpha r$ unique agents visit S_i during the first r rounds. We create a modified process TWEAKED, in which if for some $i \in \{1, \ldots, m\}$ more than $8\alpha r$ agents visit S_i in the first r rounds, we remove the extra agents. By a union bound, VISIT-EXCHANGE and TWEAKED are identical with probability at least $1 - 1/n$, and for the rest of the proof we will consider TWEAKED.

A single agent starts its walk in some vertex of U_i with probability $|U_i|/n = 2k/n$, thus, U_i does not contain any agents at the start of VISIT-EXCHANGE with probability $(1 - 2k/n)^{\alpha n} \geq e^{-2\alpha k}/2 \geq 1/(2m)$. This implies that with a probability at least $1 - (1 - 1/(2m))^m \geq 1 - e^{-2}$ there is a set U_i that does not contain any agent at time 0. In the rest of the proof we condition on this event and assume that for some fixed i, set U_i does not contain any agents at time 0.

Consider a path graph P_l with vertices $0, \ldots, l = \lceil k/2\rceil - 1$. For every agent g that visits P_i, we couple its movement to a new lazy random walk W_g on P_l, with holding probability $1/3$, that starts from vertex 0. While g is in U_i, the position of W_g in P_l is equal to the distance of g from S_i. When g leaves the set U_i, we freeze W_g at vertex 0 and activate it again when g returns.

By the construction of TWEAKED at most $8\alpha r$ unique agents ever visit S_i, and therefore, at most that many walks exist in P_l. By [3], the expected number of steps for the walks W_g to cover P_l is $\Omega(l^2/\log(8\alpha r)) = \Omega(\log^2 n/\log\log n)$. On the other hand, by the time all vertices of U_i are visited by agents, the coupled walks must cover P_l. Combining that with the fact that VISIT-EXCHANGE and TWEAKED are identical w.h.p., and that a set U_i without agents at the start exists with constant probability, we get that $\mathbb{E}[T] = \Omega(\log^2 n/\log\log n)$. ◀

Choosing m and k as in Lemma 7, and also choosing the graph R in the construction of $G_{n,\ell,m,k}$ to have logarithmic diameter, i.e., $\ell = O(\log n)$, we obtain from (1), that $\mathrm{diam}(G_{n,\ell,m,k}) = O(\log n)$, and from Lemma 7, that $\mathbb{E}[T(G_{n,\ell,m,k})] = \Omega(\log^2 n/\log\log n)$. Thus, graph $G = G_{n,\ell,m,k}$ satisfies Observation 1.

4 Upper Bound for Regular Graphs

In this section we prove Theorem 2.

4.1 Preliminaries

Let $G = (V, E)$ be any graph (not necessarily a regular one), and let A be the set of agents in VISIT-EXCHANGE, where $|A| = \alpha \cdot n$ for a constant $\alpha > 0$. The agents in A start their walks from the stationary distribution π. For a vertex u, let $N_u(t)$ be the number of agents that are at vertex u at round t. For an integer $r > 0$ and round t, let

$$\hat{N}_u(t, r) = \mathbb{E}[N_u(t + r) \mid N_v(t),\ \text{for all } v \in V] = \sum_{v\in V} p^r_{v,u} \cdot N_v(t),$$

where $p^r_{v,u}$ is the probability that a random walk starting from v is at u after exactly r rounds.

▶ **Lemma 8.** *For any vertex u, round t, and integer r,*

$$\mathbb{P}\left[\hat{N}_u(t,r) \le |A| \cdot \pi(u)/2\right] \le \exp\left(-\frac{|A| \cdot \pi(u)}{8 \cdot p^{2r}_{u,u}}\right).$$

Proof. Let $X^t_{v,g}$ be an indicator random variable, which is 1 when agent g is at vertex v at round t. Then, $N_v(t) = \sum_{g \in A} X^t_{v,g}$, which implies

$$\hat{N}_u(t,r) = \sum_{v \in V} p^r_{v,u} \sum_{g \in A} X^t_{v,g} = \sum_{g \in A} \sum_{v \in V} p^r_{v,u} \cdot X^t_{v,g} = \sum_{g \in A} Y_g,$$

where Y_g is the internal sum above for agent g. The random variables Y_g, $g \in A$, are independent, since the agents perform independent random walks. We compute the expectation and the second moment of Y_g to argue about the concentration of $\hat{N}_u(t,r)$.

$$\mathbb{E}\left[\hat{N}_u(t,r)\right] = \mathbb{E}\left[N_u(t+r)\right] = |A| \cdot \pi(u),$$

as the agents are initially distributed according to the stationary distribution π.

$$\begin{aligned}
\mathbb{E}\left[Y_g^2\right] &= \mathbb{E}\left[\sum_{v_1,v_2 \in V} p^r_{v_1,u} p^r_{v_2,u} \cdot X^t_{v_1,g} \cdot X^t_{v_2,g}\right] \\
&= \sum_{v \in V} \left(p^r_{v,u}\right)^2 \cdot \mathbb{E}\left[X^t_{v,g}\right], \quad \text{as } g \text{ cannot be in two vertices simultaneously,} \\
&= \sum_{v \in V} p^r_{v,u} \cdot \left(p^r_{v,u} \cdot \pi(v)\right), \quad \text{since } g \text{ is placed according to } \pi, \\
&= \sum_{v \in V} p^r_{u,v} \cdot \left(\pi(u) \cdot p^r_{u,v}\right), \quad \text{by reversibility,} \\
&= \pi(u) \cdot \sum_{v \in V} p^r_{u,v} \cdot p^r_{v,u} \\
&= \pi(u) \cdot p^{2r}_{u,u}.
\end{aligned}$$

We apply [8, Theorem 3.7], setting $\lambda = \mathbb{E}\left[\hat{N}_u(t,r)\right]/2$ and $M = 0$, to obtain

$$\begin{aligned}
\mathbb{P}\left[\hat{N}_u(t,r) \le |A| \cdot \pi(u)/2\right] &\le \exp\left(-\frac{\lambda^2}{2 \cdot \sum_{g \in A} \mathbb{E}\left[Y_g^2\right]}\right) \\
&\le \exp\left(-\frac{(|A| \cdot \pi(u))^2}{8 \cdot \sum_{g \in A} \pi(u) \cdot p^{2r}_{u,u}}\right) = \exp\left(-\frac{|A| \cdot \pi(u)}{8 \cdot p^{2r}_{u,u}}\right).
\end{aligned}$$ ◀

We will also need the following lemma.

▶ **Lemma 9.** *Let $X(t)$ be a simple random walk that starts at vertex u of a connected graph $G = (V,E)$. If $\deg(u)$ is the degree of u, and $d_{\min}$ is the smallest degree of G, then for any even $t \ge 0$, $\mathbb{P}[X(t) = u] \le \frac{\deg(u)}{|E|} + \frac{20 \cdot \deg(u)}{d_{\min} \cdot \sqrt{t+1}}$.*

Proof. The proof follows from an analogous result for lazy walks in [31]. Let $X'(t)$ be a lazy random walk on G starting at u, with a holding probability $1/2$. Let N_t be the number of times $X'(t)$ stays put in its first t rounds. Then,

$$\mathbb{P}\left[X'(t) = u\right] \ge \sum_{t'=0}^{t/2} \mathbb{P}\left[X'(t) = u \mid N_t = 2t'\right] \cdot \mathbb{P}\left[N_t = 2t'\right]$$

$$= \sum_{t'=0}^{t/2} \mathbb{P}\left[X(t-2t') = u\right] \cdot \binom{t}{2t'} \cdot 2^{-t}.$$

The two-step chain $X(2t'')$, for $t'' \geq 0$, is reversible and has non-negative eigenvalues, thus, $\mathbb{P}[X(2t'') = u]$ is a non-increasing function of t''. Thus, $\mathbb{P}[X(t-2t') = u] \geq \mathbb{P}[X(t) = u]$. By expanding $0 = (1 + (-1))^t$ using the binomial theorem, we get that $\sum_{t'=0}^{t/2} \binom{t}{2t'} = \sum_{t'=0}^{t/2} \binom{t}{2t'+1}$, which implies that both of these sums are equal to 2^{t-1}. Therefore,

$$\mathbb{P}\left[X'(t) = u\right] \geq \frac{1}{2} \cdot \mathbb{P}\left[X(t) = u\right].$$

On the other hand, by [31], for a lazy walk X',

$$\mathbb{P}\left[X'(t) = u\right] \leq \frac{\deg(u)}{2|E|} + \frac{10 \cdot \deg(u)}{d_{\min} \cdot \sqrt{t+1}}.$$

We combine these two inequalities to complete the proof of the lemma. ◀

4.2 Analysis

Suppose that $G = (V, E)$ is a d-regular graph with $d = O(\log n)$, thus $\pi(u) = 1/n$ for any $u \in V$. For a constant $\rho > 0$ define $r = r(\rho)$ as the smallest even integer such that

$$r \geq \max\{\rho \cdot \log^2 n,\ 256d \cdot \log n/\alpha\} = \Theta(\log^2 n). \tag{2}$$

We modify the VISIT-EXCHANGE process to create a new process called TWEAKED$_r$, as follows: At the end of each round $t \geq 0$, we add a minimal set of agents to the process to make sure that $\hat{N}_u(t, r) \geq |A| \cdot \pi(u)/2 = \alpha/2$, for every vertex u. Next we prove that, in the first polynomially many rounds TWEAKED$_r$ and VISIT-EXCHANGE are equivalent, w.h.p. Therefore, the results that we prove for TWEAKED$_r$, also hold for VISIT-EXCHANGE, w.h.p. This technique allows us to avoid dealing with dependencies of the random walks, which would arise if we directly analyzed VISIT-EXCHANGE conditioned on $\hat{N}_u(t, r) \geq \alpha/2$ for all u and t. (Similar tweaked processes are used in the proofs of Theorems 2 to 5 to circumvent some dependencies.)

▶ **Lemma 10.** *For any constant $c > 0$, there is a constant ρ such that* VISIT-EXCHANGE *and* TWEAKED$_r$ *are identical for the first T' rounds of their execution with probability at least $1 - T' \cdot n^{-(c+2)}$.*

Proof. By Lemma 9, $p_{u,u}^{2r} \leq \frac{2}{n} + \frac{20}{\sqrt{2r+1}} \leq \frac{20}{\sqrt{r}}$, since $r = O(\log^2 n)$. For $t < T'$, we substitute the above inequality into Lemma 8, and use the fact that $|A| \cdot \pi(u) = \alpha$, to get that

$$\mathbb{P}\left[\hat{N}_u(t, r) \leq \alpha/2\right] \leq \exp\left(-\frac{\alpha}{8 \cdot p_{u,u}^{2r}}\right) \leq \exp\left(-\frac{\alpha}{160} \cdot \sqrt{r}\right) \leq n^{-(c+3)},$$

for a sufficiently large constant ρ. By applying a union bound over all vertices u and rounds $t < T'$, we complete the proof. ◀

Consider two vertices u and v with distance $O(r/\max\{d, \log^2 \log n\})$, and assume u is informed at round t_0. The next key lemma provides a lower bound for the probability that v becomes informed $O(r)$ rounds after t_0. The lemma holds for any execution prefix of TWEAKED$_r$ up to round t_0, which means we can apply it repeatedly to prove Theorem 2. Let $\mathcal{K}_t$ be the σ-field that determines the execution of TWEAKED$_r$ until round t.

▶ **Lemma 11.** *Let $h = \max\{d, \log\log n\}$, and $k_{\max}(\gamma) = \frac{\gamma \cdot r}{\max\{d,(\log\log n)^2\}}$. There are constants $\gamma, \beta > 0$, such that for any round t_0 and any two vertices u, v with $\operatorname{dist}(u, v) \leq k_{max}(\gamma)$, given $\mathcal{K}_{t_0}$ and that u is informed at round t_0, vertex v is informed at round $t_0 + 2r$ with probability at least $1 - e^{-\beta \cdot h}$.*

Proof. Case $d = \omega(\log\log n)$. To simplify presentation, we assume $t_0 = 0$ and omit the conditional $\mathcal{K}_{t_0}$ throughout the proof. Fix the constant γ such that $k_{\max}(\gamma) \leq \frac{\alpha r}{256d}$. Consider two vertices u, v such that a shortest path between them is $u = u_0, \ldots, u_k = v$, where $k = \operatorname{dist}(u, v) \leq k_{\max}(\gamma)$. For a round $t \geq r$ and $i \in \{0, \ldots, k-1\}$, let $Z_{i,t}$ be the number of agents in the neighbourhood $\Gamma(u_i)$ of vertex u_i at round t. Then, by definition of $\textsc{tweaked}_r$,

$$\mathbb{E}\left[Z_{i,t}\right] = \sum_{w\in\Gamma(u_i)} \mathbb{E}\left[N_{u_i}(t)\right] = \sum_{w\in\Gamma(u_i)} \mathbb{E}\left[\hat{N}_{u_i}(t-r,r)\right] \geq \alpha \cdot d/2.$$

Since the agents make independent random walks, by a Chernoff bound we get that

$$\mathbb{P}\left[Z_{i,t} \geq \alpha \cdot d/4\right] \geq 1 - e^{-\alpha \cdot d/16}.$$

If $\mathcal{E}$ is the event that $Z_{i,t} \geq \alpha \cdot d/4$ for all $i \in \{0, \ldots k-1\}$ and $t \in \{r, \ldots 2r\}$ simultaneously, then, by a union bound,

$$\mathbb{P}\left[\mathcal{E}\right] \geq 1 - k \cdot r \cdot e^{-\alpha \cdot d/16} \geq 1 - e^{-\beta d}/2,$$

for a small enough constant β, because $kr = O(\operatorname{poly}(\log n))$ and $d = \omega(\log\log n)$.

We modify $\textsc{tweaked}_r$ as follows: If $\mathcal{E}$ does not hold, then we add a minimum number of agents to the process so that $\mathcal{E}$ holds. We call the new process $\textsc{r-tweaked}_r$, and observe that $\textsc{tweaked}_r$ and $\textsc{r-tweaked}_r$ are identical with probability at least $1 - e^{-\beta d}/2$.

We divide the rounds $r, \ldots, 2r-1$ of $\textsc{r-tweaked}_r$ into $r/2$ phases of 2 rounds each. For each $0 \leq i < r/2$, let $\mathcal{K}'_i$ be the σ-algebra which determines the execution prefix of $\textsc{r-tweaked}_r$ until round $r + 2i \leq 2r$. Let p_i be the largest integer, between 0 and k, such that vertex $w = u_{p_i}$ is informed at round $r + 2i$. If $p_i < k$, then each agent that is in the neighbourhood of w in round $r + 2i$, informs vertex u_{p_i+1} after two rounds, with probability $1/d^2$, by going through w. Define a Bernoulli random variable X_i, such that $X_i = 1$ if $p_i < k$ and u_{p_i+1} is informed in round $r + 2(i+1)$, i.e., the ith phase is successful. For technical convenience, we also define $X_i = 1$ if $p_i = k$, i.e., v is already informed in that phase. Then,

$$\mathbb{P}\left[X_i = 1 \mid \mathcal{K}'_i\right] \geq 1 - \left(1 - d^{-2}\right)^{\alpha \cdot d/4} \geq 1 - e^{-\alpha/(4d)} \geq \frac{\alpha}{8d}. \tag{3}$$

Define $Y = \sum_{i=0}^{r/2-1} Y_i$, where Y_i are independent Bernoulli random variables with success probability $\alpha/8d$. By our choice of γ and (2),

$$\mathbb{E}\left[Y\right] = \frac{\alpha r}{16d} \geq 8(k_{\max}(\gamma) + \log n) \geq 8(k + \log n),$$

and, by a Chernoff bound,

$$\mathbb{P}\left[Y \geq k\right] \geq \mathbb{P}\left[Y \geq \mathbb{E}\left[Y\right]/2\right] \geq 1 - e^{-\mathbb{E}[Y]/8} \geq 1 - 1/n \geq 1 - e^{-\beta d}/2,$$

since $d = O(\log n)$ and by choosing constant β smaller if necessary. On the other hand, for $X = \sum_{i=1}^{r/2-1} X_i$, (3) implies that X stochastically dominates Y, in particular,

$$\mathbb{P}\left[X \geq k\right] \geq \mathbb{P}\left[Y \geq k\right] \geq 1 - e^{-\beta d}/2.$$

Note, $X \geq k$ implies that v is informed in $\textsc{r-tweaked}_r$ at round $2r$. Since $\textsc{r-tweaked}_r$ and $\textsc{tweaked}_r$ are identical with probability $1 - e^{-\beta d}/2$, vertex v must be informed in $\textsc{tweaked}_r$ at round $2r$ with probability at least $1 - e^{-\beta d} = 1 - e^{-\beta h}$.

Case $d = O(\log\log n)$. As in the previous case, we assume $t_0 = 0$ and consider the spread of information along a shortest path from u to v, namely, $u = u_0, \ldots, u_k = v$. Fix a round $t \geq r$ and some $i \in \{0, \ldots, k-1\}$. Let $l = (\eta \log\log n)^2$ for some constant η that will be specified later. For an agent g define R_g as the number of times agent g visits u_i in rounds $t, \ldots, t+l-1$. If $X_g(t')$ is the position of the agent g at round t', then $R_g = \sum_{t'=t}^{t+l-1} \mathbb{1}_{X_g(t')=u_i}$, so by Lemma 9,

$$\mathbb{E}\left[R_g \mid X_g > 0\right] = \sum_{t'=t}^{t+l-1} \mathbb{P}\left[X_g(t') = u_i \mid R_g > 0\right] \leq 1 + \sum_{t'=t}^{t+l-1} \left(\frac{1}{n} + \frac{20}{\sqrt{t-t'+1}}\right) \leq 50 \cdot \sqrt{l}.$$

Let $Z_{i,t}$ be the number of unique agents that visit u_i in rounds $t, \ldots, t+l-1$.

$$\begin{aligned}
\mathbb{E}\left[Z_{i,t}\right] &= \sum_{g \in A} \mathbb{P}\left[R_g > 0\right] = \sum_{g \in A} \frac{\mathbb{E}\left[R_g\right]}{\mathbb{E}\left[R_g \mid R_g > 0\right]} \\
&\geq \frac{\sum_{g \in A} \mathbb{E}\left[R_g\right]}{50 \cdot \sqrt{l}} = \frac{\sum_{t'=t}^{t+l-1} \mathbb{E}\left[N_{u_i}(t')\right]}{50 \cdot \sqrt{l}} = \frac{\sum_{t'=t}^{t+l-1} \mathbb{E}\left[\hat{N}_{u_i}(t'-r, r)\right]}{50 \cdot \sqrt{l}} \\
&\geq \frac{l \cdot \alpha/2}{50 \cdot \sqrt{l}} = \frac{\alpha \cdot \sqrt{l}}{100}.
\end{aligned}$$

Since the agents are performing independent random walks, then by a Chernoff bound,

$$\mathbb{P}\left[Z_{i,t} \geq \alpha \cdot \sqrt{l}/200\right] \geq 1 - \exp\left(-\frac{\alpha\eta}{800} \cdot \log\log n\right) \geq 1 - 1/\log^5 n,$$

for a suitable choice of η. We now let $\mathcal{E}$ be the event $Z_{i,t} \geq \alpha \cdot \sqrt{l}/200$ for all $i \in \{0, \ldots, k-1\}$ and $t \in \{r, \ldots, 2r\}$, simultaneously. As before, we create R-TWEAKED$_r$ by adding minimum number of agents to TWEAKED$_r$ to ensure that $\mathcal{E}$ holds. Since $rk = O(\log^4 n)$, by a union bound, there is a constant β such that $\mathbb{P}\left[\mathcal{E}\right] \geq 1 - e^{-\beta h}/2$.

The rest of the proof follows the same line of logic as in the case of $d = \omega(\log\log n)$. The only difference is that instead of phases of 2 rounds, we consider phases of l rounds. $\mathcal{E}$ implies that after each phase R-TWEAKED$_r$ informs the next vertex on the path with a constant probability since $\sqrt{l} = \Omega(d)$. Therefore, as long as $k \leq \gamma \cdot r/l$ for a sufficiently small γ, vertex v becomes informed at round $2r$ of R-TWEAKED$_r$ w.h.p., which completes the proof. ◀

Proof of Theorem 2. First, we consider the TWEAKED$_r$ process for a constant ρ chosen by Lemma 10 such that TWEAKED$_r$ is identical to VISIT-EXCHANGE in the first n^2 rounds of its execution, with probability at least $1 - n^{-2}$. Consider a shortest path $s = u_0, \ldots, u_m = u$ from source vertex s to vertex u. Let $k = k_{max}(\gamma)$ be the upper bound on the distance from Lemma 11, and as before $h = \max\{d, \log\log n\}$. We divide the execution of TWEAKED$_r$ into phases of $2r$ rounds each. If vertex u_i is informed at the end of a phase, then by Lemma 11, the vertex $u_{\min\{m,i+k\}}$ will be informed in the next phase of $2r$ rounds with probability at least $1 - e^{-\beta h}$, independently from the past.

For some constant $\eta \in (0,1)$, let $l = \lceil m/k + \log n/h \rceil/(1-\eta)$. For $i \in \{1, \ldots, l\}$, let X_i be a Bernoulli random variable that is 0 if in the ith phase of TWEAKED$_r$ either k new vertices along the specified path become informed, or vertex u becomes informed, i.e., the phase is successful. For $X = \sum_{i=1}^{l} X_i$, if $X < l - \lceil m/k \rceil$ then vertex u is informed at the end of the lth phase, because at least $\lceil m/k \rceil$ phases were successful. By a stochastic dominance argument as in Lemma 11 we upper bound $\mathbb{P}\left[X < l - \lceil m/k \rceil\right]$.

Let $\{Y_i\}_{1 \leq i \leq l}$ be a collection of independent Bernoulli random variables $\mathbb{P}\left[Y_i = 1\right] = e^{-\beta h}$. By Lemma 11, $\mathbb{P}\left[X_i = 1 \mid X_1, \ldots, X_{i-1}\right] \leq \mathbb{P}\left[Y_i = 1\right]$, and therefore, for $Y = \sum_{i=1}^{l} Y_i$,

$$\mathbb{P}\left[X > l - \lceil m/k \rceil\right] \leq \mathbb{P}\left[Y > l - \lceil m/k \rceil\right] \leq \mathbb{P}\left[Y \geq l - \lceil m/k + \log n/h \rceil\right]$$

$$= \mathbb{P}[Y \ge \eta \cdot l] = \mathbb{P}\left[Y \ge \eta \cdot e^{\beta h} \cdot \mathbb{E}[Y]\right] \\ \le \left(\eta \cdot e^{\beta h - 1}\right)^{-\eta \cdot l} \le n^{-3},$$

by a Chernoff bound (Lemma 42) and by taking a value of η that is sufficiently close to 1. Thus, after $l \cdot 2r$ rounds of TWEAKED$_r$ vertex u is informed with probability $1 - n^{-3}$. By a union bound over all vertices, and the fact that TWEAKED$_r$ and VISIT-EXCHANGE are identical in the first n^2 rounds we get that $T \le l \cdot 2r$ w.h.p. Since $k = O(r/\max\{d, (\log\log n)^2\})$, and $m \le \mathrm{diam}(G)$, and $h = \max\{d, \log\log n\}$, we finally get that, w.h.p.,

$$T = O\left(\max\{d, (\log\log n)^2\} \cdot \mathrm{diam}(G) + \frac{\log^3 n}{h}\right) = \tilde{O}\left(d \cdot \mathrm{diam}(G) + \frac{\log^3 n}{d}\right). \quad ◀$$

5 Upper Bound with Average Degree

In this section we prove Theorem 3.

5.1 Preliminaries

Recall that $G = (V, E)$ is a graph with average degree d_{avg} and minimum degree $d_{\min} = \Omega(d_{\mathrm{avg}})$. The set of agents of VISIT-EXCHANGE is A, and $|A| = \alpha \cdot n$ for a constant $\alpha > 0$. The agents in A start their walks from the stationary distribution π. Let $\epsilon = d_{\min}/d_{\mathrm{avg}} = \Omega(1)$. Then, for every vertex $u \in V$,

$$\pi(u) = \frac{\deg(u)}{2|E|} = \frac{\deg(u)}{n \cdot d_{\mathrm{avg}}} \ge \frac{\epsilon}{n}. \tag{4}$$

We define $N_u(t), \hat{N}_u(t, r)$ and $p^t_{v,u}$ as in Section 4.

5.2 Analysis

As in Section 4, we modify the VISIT-EXCHANGE process to create a new process called TWEAKED$_r$, that depends on a parameter $r = \Theta(\log^2 n)$: For all rounds t and vertices u, we add a minimal set of agents to the process to make sure that $\hat{N}_u(t, r) \ge |A| \cdot \pi(u)/2$.

▶ **Lemma 12.** *For any constant $c > 0$, there is a parameter $r = O(\log^2 n)$ such that* VISIT-EXCHANGE *and* TWEAKED$_r$ *are identical for the first T' rounds of their execution with probability at least $1 - T' \cdot n^{-(c+2)}$.*

Proof. By Lemma 9 and condition (4),

$$p^{2r}_{u,u} \le 2\pi(u)\left(1 + \frac{20 \cdot |E|}{d_{\min} \cdot \sqrt{2r+1}}\right) \le 2\pi(u)\left(1 + \frac{40n}{\epsilon\sqrt{2r+1}}\right) \le \frac{100n \cdot \pi(u)}{\epsilon\sqrt{2r+1}},$$

where the last inequality holds assuming a large value of n and $r = O(\log^2 n)$. Substituting in Lemma 8, gives

$$\mathbb{P}\left[\hat{N}_u(t, r) \le |A| \cdot \pi(u)/2\right] \le \exp\left(-\frac{|A| \cdot \pi(u)}{8 \cdot p^{2r}_{u,u}}\right) \le \exp\left(-\frac{|A| \cdot \epsilon\sqrt{2r+1}}{100n}\right) \\ = \exp\left(-\frac{\alpha\epsilon}{100} \cdot \sqrt{2r+1}\right) \le n^{-(c+3)},$$

for $r = \eta \log^2 n$ for a sufficiently large constant η. By applying a union bound over all vertices u and rounds $t < T'$, we complete the proof. ◀

▶ **Lemma 13.** *Let k be the length of a shortest path from the source vertex s to vertex u. For any constant $c > 0$, vertex u becomes informed in at most $O(r \cdot d_{\text{avg}} \cdot (k + \log n))$ rounds of* TWEAKED$_r$*, with probability at least $1 - n^{-(c+1)}$.*

Proof. Let $s = u_0, u_1, \ldots, u_k = u$ be a shortest path from vertex s to u. We divide the execution of TWEAKED$_r$ into phases of $(r+1)$ rounds each. For each $i \geq 0$, let $\mathcal{K}_i$ be the σ-algebra fixing the execution prefix of TWEAKED$_r$ up to round $i(r+1)$. Let p_i be the largest integer, between 0 and k, such that $w = u_{p_i}$ is informed at round $i(r+1) - 1$. By the definition of TWEAKED,

$$\mathbb{E}\left[N_w\left((i+1)(r+1)-1\right) \mid \mathcal{K}_i\right] = \hat{N}_w(i(r+1), r) \geq \frac{|A| \cdot \pi(w)}{2}.$$

Since the agents move independently, by a Chernoff bound we have that for an event $\mathcal{E}_i = \left\{N_w\left((i+1)(r+1)-1\right) \geq \frac{|A| \cdot \pi(w)}{4}\right\}$,

$$\mathbb{P}\left[\mathcal{E}_i \mid \mathcal{K}_i\right] \geq 1 - \exp\left(-\frac{|A| \cdot \pi(w)}{16}\right) \geq 1 - e^{-\alpha\epsilon/16}.$$

Notice that, if $p_i < k$, then each agent that visits u in round $(i+1)(r+1)-1$, informs vertex u_{p_i+1} with probability $1/\deg(w)$ at the next round. Define $Y_i = 1$ if either u_{p_i+1} is informed in round $(i+1)(r+1)$, or $p_i = k$. Then,

$$\begin{aligned}\mathbb{P}\left[Y_i = 1 \mid \mathcal{E}_i; \mathcal{K}_i\right] &\geq 1 - \left(1 - \frac{1}{\deg(w)}\right)^{|A| \cdot \pi(w)/4} \geq 1 - \exp\left(-\frac{|A| \cdot \pi(w)}{4 \cdot \deg(w)}\right) \\ &\geq 1 - \exp\left(-\frac{\alpha}{4} \cdot \frac{n}{2|E|}\right) = 1 - e^{-\frac{\alpha}{4 d_{\text{avg}}}} \geq \min\left\{\frac{1}{2}, \frac{\alpha}{8 d_{\text{avg}}}\right\}.\end{aligned}$$

Then,

$$\begin{aligned}\mathbb{P}\left[Y_i = 1 \mid \mathcal{K}_i\right] &\geq \mathbb{P}\left[Y_i = 1 \mid \mathcal{K}_i; \mathcal{E}_i\right] \cdot \mathbb{P}\left[\mathcal{E}_i \mid \mathcal{K}_i\right] \\ &\geq \min\left\{\frac{1}{2}, \frac{\alpha}{8 d_{\text{avg}}}\right\} \cdot \left(1 - e^{-\alpha\epsilon/16}\right) \\ &= \eta / d_{\text{avg}}, \end{aligned} \tag{5}$$

where η is a constant that could depend on d_{avg} if $d_{\text{avg}} = O(1)$. As in Lemma 11, (5) implies that after at most $O(d_{\text{avg}}(k + \log n))$ phases, vertex u must become informed w.h.p. Since each phase lasts $r + 1$ rounds, we complete the proof. ◀

Proof of Theorem 3. Let $c > 0$ be any fixed constant and $r = O(\log^2 n)$ be as determined from Lemma 13. Let k be the length of a shortest path from the source vertex to a fixed vertex $u \in V$. By Lemma 13, in at most $O(d_{\text{avg}} \cdot r \cdot (k + \log n)) = O(d_{\text{avg}} \cdot \log^2 n \cdot (\text{diam}(G) + \log n))$ rounds u becomes informed with probability at least $1 - n^{-(c+1)}$. Applying a union bound for all n vertices, we prove that TWEAKED$_r$ informs all vertices in $T' = O(d_{\text{avg}} \cdot \log^2 n \cdot (\text{diam}(G) + \log n))$ rounds, with probability at least $1 - n^{-c}$. Finally, by Lemma 12, TWEAKED$_r$ and VISIT-EXCHANGE are identical in the first $T' \leq n^2$ rounds of their executions, with probability at least $1 - n^{-c}$, and therefore, VISIT-EXCHANGE informs all vertices of G in T' rounds with probability at least $1 - 2n^{-c}$. ◀

6 Constant Degree Expanders

In this section we prove Theorem 4, that bounds the spreading time of VISIT-EXCHANGE on a d-regular expander $G = (V, E)$, where $d \geq 3$ and $d = O(1)$. The process uses $\alpha \cdot n$ agents for a constant $\alpha > 0$. Let λ be the largest non-trivial eigenvalue of the normalised adjacency matrix $\frac{1}{d} \cdot A$ in absolute value, i.e., $\lambda := \max\{\lambda_2, |\lambda_n|\}$, where $1 = \lambda_1 > \lambda_2 \geq \cdots \geq \lambda_n \geq -1$ are the n real-valued eigenvalues. We will assume that G is a (spectral) expander, that is, λ is bounded away from 1 by a positive constant (independent of n).[4] For a set $S \subset V$, let $E(S, V \setminus S)$ be the set of edges (u, v) such that $u \in S$ and $v \notin S$. The *conductance* of G is

$$\phi = \min_{S \subset V; 1 \leq |S| \leq n/2} \frac{|E(S, V \setminus S)|}{d \cdot |S|}.$$

By Cheeger's inequality (e.g., [29, Inequality (13.6)]), $\phi \geq (1 - \lambda)/2 = \Omega(1)$. Let t^{∞}_{mix} be the uniform mixing time of a random walk in G, i.e., for a random walk $X(t)$ on G and a vertex $v \in V$,

$$|\mathbb{P}\left[X(t^{\infty}_{\text{mix}}) = v\right] - 1/n| \leq 1/(2n).$$

Since we assume $1 - \lambda = \Omega(1)$, by [29, Inequality (12.11)], $t^{\infty}_{\text{mix}} = O(\log n)$.

Throughout the proof we will denote the set of informed vertices at round t as I_t.

6.1 Preliminaries

The following result on probability amplification for expanders uses tools from [1].

▶ **Lemma 14.** *Let $X(t)$ be a non-lazy simple random walk on a d-regular graph $G = (V, E)$, that starts from the (uniform) stationary distribution π. Let λ be the largest non-trivial eigenvalue of the normalised adjacency matrix of G. For a set $S \subset V$ and integer $t \leq 2n/|S|$, if τ_S is the first time when the walk X visits any vertex in S, then*

$$\mathbb{P}\left[\tau_S \leq t\right] \geq \frac{t(1 - \lambda)|S|}{4n}.$$

Proof. Consider a lazy random walk $X'(t)$, with holding probability 1/2, which also starts from the stationary distribution. Let τ'_S be the first round when X' visits any vertex of S. If P is the transition matrix of the walk X', then its largest non-trivial eigenvalue is $\lambda' = (1 + \lambda)/2$. Let Q be the transition matrix P that is restricted to the set $\overline{S} = V \setminus S$, i.e., $Q_{u,v} = P_{u,v}$ for $u, v \notin S$, and 0 otherwise. If a walk does not reach the set S, it must start in $\overline{S}$ and only follow edges that are present in Q, up until round t. Thus,

$$\mathbb{P}\left[\tau'_S > t\right] = \sum_{u \notin S} \pi(u) \sum_{v \in \overline{S}} (Q^t)_{u,v} = \frac{1}{n} \sum_{u,v \in \overline{S}} (Q^t)_{u,v}.$$

The sum of the elements of the matrix Q^t above can be written as $||Q^t x^T||_1$, where x is a row vector taking values 1 on the set $\overline{S}$, and 0 otherwise. Let λ_Q be the largest non-trivial eigenvalue of Q. Then,

$$\mathbb{P}\left[\tau_S > t\right] = \frac{1}{n} \cdot ||Q^t x^T||_1 \leq \frac{1}{n} \cdot |\overline{S}| \cdot ||Q^t x^T||_\infty \leq \frac{|\overline{S}|}{n} \lambda_Q^t \leq \lambda_Q^t.$$

By [1, Corollary 3.34] and using the definition of relaxation time, we get that for the chain Q,

$$\lambda_Q \leq 1 - \frac{(1 - \lambda')|S|}{n} = 1 - \frac{(1 - \lambda)|S|}{2n},$$

---

[4] For a bipartite expander, we have $\lambda_n = -1$. In order to apply our result to bipartite graphs, simply add one self-loop to every vertex and note that for the resulting graph G' with self-loops, $\lambda(G')$ is still bounded away from 1 by a (different) positive constant.

and thus,

$$\begin{aligned}\mathbb{P}\left[\tau'_S \le t\right] &\ge 1 - \left(1 - \frac{(1-\lambda)|S|}{2n}\right)^t \\ &\ge 1 - \frac{1}{1 + t(1-\lambda)|S|/(2n)} \\ &\ge \frac{t(1-\lambda)|S|}{4n},\end{aligned}$$

where the last two inequalities hold because $t \le 2n/|S|$ and for any $x \in (0,1)$, $(1-x)^t \le 1/(1+tx)$. For the non-lazy walk $X(t)$, $\mathbb{P}\left[\tau_S \le t\right] \ge \mathbb{P}\left[\tau'_S \le t\right]$, which completes the proof. ◀

▶ **Lemma 15.** *Let u be any vertex of a graph H with largest degree Δ. For any integer $k > 0$, there are at most Δ^{2k} connected subgraphs of H with k vertices.*

Proof. Consider a (not necessarily simple) path in H that has length $2k$, starts at u and contains at most k unique vertices. The subgraph induced by the vertices on the path is connected and contains u. The number of such paths is at most Δ^{2k} because we can construct them starting from u, choosing one of the at most Δ neighbours at each step. Therefore, it suffices to show that each connected subgraph H' of H, that contains u can be traversed by a path of length at most $2k$. Consider a spanning tree R of H'. A depth first search traversal path of R starting from u uses at most $2(k-1)$ edges, therefore, it satisfies our requirement. ◀

6.2 Analysis

We divide the process into three phases, and we prove that each takes at most $O(\log n)$ rounds. The first phase corresponds to the rounds until $|I_t| \ge b \cdot \log n$ for some $b > 0$. The second one corresponds to the subsequent rounds after that until $|I_t| \ge n/2$. Finally, in the third phase all vertices of G become informed.

▶ **Lemma 16.** *For any $b, c > 0$, there is a round $\tau_1 = O(\log n)$ such that $|I_\tau| \ge b \log n$ with probability at least $1 - n^{-c}$.*

Proof. For an agent g, let X_g be the indicator variable that g visits the source vertex in the first $\tau' = \eta' \log n$ rounds. By considering the singleton set containing the source vertex and applying Lemma 14,

$$\mathbb{P}\left[X_g = 1\right] \ge \frac{(1-\lambda)\tau'}{4n}.$$

Thus, if A' is the set of agents that have visited the source in the first τ' rounds, then

$$\mathbb{E}\left[|A'|\right] \ge \frac{\alpha(1-\lambda)\tau'}{4}.$$

Furthermore, since the agents perform independent random walks, by a Chernoff bound,

$$\mathbb{P}\left[|A'| \ge \frac{\alpha(1-\lambda)\eta'}{8} \cdot \log n\right] \ge 1 - \exp\left(-\frac{\alpha(1-\lambda)\eta'}{32} \cdot \log n\right).$$

Thus, for any constant a', we can take a large enough η' such that

$$\mathbb{P}\left[|A'| \ge a' \log n\right] \ge 1 - n^{-c}/2.$$

We set $\tau = \tau' + t^{\infty}_{\mathrm{mix}} = O(\log n)$. Let N be the number of vertices that contain an agent from A' at round τ. By the property of mixing, for a vertex u, the probability that a given agent is at u is at least $1/(2n)$. Thus, $\mathbb{E}[N \mid A'] \geq |A'|/2$.

Furthermore, N is a function of the independent walks performed by the agents in A', and changing the walk of one of the agents can change N by at most 1. Therefore, by the Method of Bounded Differences (see [14, Corollary 5.2]),

$$\mathbb{P}\left[N \geq \frac{\mathbb{E}[N]}{2} \;\middle|\; A'\right] \geq 1 - e^{-\frac{2\cdot(\mathbb{E}[N]/2)^2}{|A'|}} = 1 - e^{-\frac{|A'|}{8}},$$

which implies that

$$\mathbb{P}\left[N \geq \frac{a' \cdot \log n}{2} \;\middle|\; |A'| \geq a' \cdot \log n\right] \geq 1 - n^{-\frac{a'}{8}}.$$

We take $a' \geq 2b$ and also such that $n^{-\frac{a'}{8}} \leq n^{-c}/2$. Then, by applying union bound, we get

$$\mathbb{P}[|I_\tau| \geq b \cdot \log n] \geq \mathbb{P}[N \geq b \cdot \log n] \geq 1 - n^{-c}. \qquad \blacktriangleleft$$

Next we describe a modification of the VISIT-EXCHANGE process, that depends on two constants b and r, which are fixed later. Recall that ϕ is the conductance of the graph and let $\epsilon = \phi/(2d)$. For a set $S \subset V$, let $G(S)$ be the subgraph of G induced by S. Define the following set of subsets of V:

$$\mathcal{S}(b) = \{S \subset V \mid s \in S,\ G(S) \text{ is connected, and } b \cdot \log n \leq |S| \leq n/2\}.$$

For a round $t \geq 0$, we say that a set $S \in \mathcal{S}(b)$ is (t,r)-good, if at least $(1-\epsilon)|S|$ vertices in S are visited by some agent in rounds $t, \ldots, t+r-1$.

If for some round t, some set $S \in \mathcal{S}(b)$ is not (t,r)-good, then we add a minimal set of agents in S at round $t+r-1$ to turn S into a (t,r)-good set. The modified process is called TWEAKED$_{b,r}$, for which we will use the same notation as for VISIT-EXCHANGE. We show that TWEAKED$_{b,r}$ and VISIT-EXCHANGE are identical in the first polynomial number of rounds, w.h.p.

▶ **Lemma 17.** *Let T' be any positive integer. For any constants $b, c > 0$, there is a positive integer r, such that every set $S \in \mathcal{S}(b)$ is (t,r)-good for every round $t \leq T'$ of* VISIT-EXCHANGE*, with probability at least $1 - T' \cdot n^{-(c+1)}$.*

Proof. Consider a fixed round $t \leq T'$ and a set $S' \subset S$ for some $S \in \mathcal{S}(b)$, such that $|S'| \geq \epsilon|S|$. For an integer r to be fixed later in the proof, let X_g be an indicator random variable that agent g visits some vertex in S' between rounds t and $t+r-1$. By Lemma 14, $\mathbb{P}[X_g = 1] \geq \frac{(1-\lambda)r|S'|}{4n}$ and thus, if $N_{S'}$ is the number of unique agents visiting S' between rounds t and $t+r-1$, then

$$\mathbb{E}[N_{S'}] \geq \frac{\alpha(1-\lambda)r}{4} \cdot |S'| \geq \frac{\epsilon\alpha(1-\lambda)r}{4} \cdot |S|.$$

By an application of a Chernoff bound, and setting $\eta = \frac{\epsilon\alpha(1-\lambda)}{32}$ for conciseness, we get

$$\mathbb{P}[N_{S'} \geq 1] \geq \mathbb{P}\left[N_{S'} \geq \frac{\mathbb{E}[N_{S'}]}{2}\right] \geq 1 - e^{-\eta r|S|}.$$

Next, notice that S is (t,r)-good if and only if for all $S' \subset S$ with $|S'| \geq \epsilon|S|$, $N_{S'} \geq 1$. Thus,

$$\mathbb{P}\left[S \text{ is } (t,r)\text{-good}\right] = \mathbb{P}\left[\bigcap_{\substack{S' \subset S \\ |S'| \geq \epsilon|S|}} \{N_{S'} \geq 1\}\right] \geq 1 - 2^{|S|} \cdot e^{-\eta r|S|} \geq 1 - e^{-(\eta r - 1)\cdot|S|}.$$

Next we apply another union bound for all sets S:

$$\begin{aligned}
\mathbb{P}\left[S \text{ is } (t,r)\text{-good for all } S \in \mathcal{S}(b)\right] &\geq 1 - \sum_{S \in \mathcal{S}(b)} e^{-(\eta r - 1)\cdot|S|} \\
&\geq 1 - \sum_{k \geq b \log n} \sum_{\substack{S \in \mathcal{S}(b) \\ |S| = k}} e^{-(\eta r - 1)\cdot k} \\
&\geq 1 - \sum_{k \geq b \log n} \left(d^2 \cdot e^{-\eta r + 1}\right)^k, \quad \text{by Lemma 15,} \\
&\geq 1 - 2\left(e^{-\eta r + 2\ln d + 1}\right)^{b \log n} \\
&\geq 1 - n^{-(c+1)},
\end{aligned}$$

for a sufficiently large constant r, that depends on b and c. Applying another union bound over all rounds $t \leq T'$ completes the proof. ◀

▶ **Lemma 18.** *For any constant $c > 0$, there is a constant b such that for any r, if $I_t \in \mathcal{S}(b)$ in* TWEAKED$_{b,r}$*, then $|I_{t+r}| \geq (1+\psi)|I_t|$, with probability at least $1 - n^{-(c+1)}$, where $\psi = \epsilon/(2d^2)$.*

Proof. Let S be the set of vertices $u \in I_t$ that have an uninformed neighbour. Then,

$$|E(I_t, V \setminus I_t)| = \sum_{u \in S} \deg_{V \setminus I_t}(u) \leq d \cdot |S|.$$

On the other hand, since G is an expander with conductance ϕ,

$$|E(I_t, V \setminus I_t)| \geq \phi \cdot |I_t|,$$

since $|I_t| \leq n/2$ by the condition of the lemma. Combining the two inequalities above gives

$$|S| \geq \phi \cdot |I_t|/d = 2\epsilon|I_t|.$$

Since we are considering the TWEAKED$_{b,r}$ process, the set I_t is (t,r)-good and therefore, at most $\epsilon|I_t|$ vertices in I_t are not visited by any agent in rounds $t, \ldots, t+r-1$. Since $S \subset I_t$, we conclude that at least $\epsilon|I_t|$ vertices in S are visited by some agent in rounds $t, \ldots, t+r-1$. Suppose the set of these vertices is S'.

For $u \in S'$, let X_u be an indicator random variable that the first agent that visits u in rounds $t, \ldots, t+r-1$ (there must be one), visits a vertex in $V \setminus I_t$. Then, $\mathbb{P}\left[X_u = 1 \mid I_t, S'\right] \geq 1/d$ and for $N = \sum_{u \in S'} X_u$,

$$\mathbb{E}\left[N \mid I_t, S'\right] \geq \frac{|S'|}{d} \geq \frac{\epsilon|I_t|}{d} = 2\psi d \cdot |I_t|.$$

Furthermore, since the variables X_u are independent, we can apply the Chernoff bound

$$\mathbb{P}\left[N \geq \psi d \cdot |I_t| \mid I_t, S'\right] \geq 1 - e^{-\epsilon|I_t|/(8d)}.$$

On the other hand, $|I_{t+r}| \geq |I_t| + N/d$, because every vertex in $I_{t+r} \setminus I_t$ has at most d neighbours in I_t. Therefore,

$$\mathbb{P}\left[|I_{t+r}| \geq (1+\psi) \cdot |I_t| \mid I_t\right] \geq 1 - e^{-\epsilon |I_t|/(8d)} \geq 1 - n^{-b\epsilon/(8d)}.$$

By taking $b \geq 8(c+1)d/\epsilon$, we complete the proof of the lemma. Note that b depends on c but not r. ◀

▶ **Lemma 19.** *For any constant $c > 0$, if $|I_{t_0}| \geq n/2$, then for some round $\tau_3 = t_0 + O(\log n)$, all vertices of the graph are informed with probability at least $1 - n^{-c}$.*

Proof. Let $t_1 = t_0 + t_{\text{mix}}^{\infty}$ and for an agent g, let $X_g(t)$ the location of agent g at round t. By definition of uniform mixing, $\mathbb{P}\left[X_g(t_1) \in I_t \mid I_t\right] \geq |I_t|/(2n) \geq 1/4$. Therefore, if A' is the set of informed agents at round t_1, by a Chernoff bound,

$$\mathbb{P}\left[|A'| \geq \alpha n/8 \mid I_t\right] \geq 1 - e^{-\alpha n/32}.$$

Set $t_2 = t_1 + t_{\text{mix}}^{\infty}$. Thus, for every vertex u and agent g, we have that

$$\mathbb{P}\left[X_g(t_2) = u \mid I_t, A'\right] \geq 1/(2n).$$

Finally, let $\tau_3 = t_2 + \eta \cdot \log n$, for some $\eta > 0$. For $g \in A'$, let $\tau_{u,g}$ be the first round when g visits u after t_2. Then, omitting the conditioning on I_t and A' for conciseness,

$$\begin{aligned}\mathbb{P}\left[\tau_{u,g} \leq \tau_3\right] &= \sum_{v \in V} \mathbb{P}\left[\tau_{u,g} \leq \tau_3 \mid X_g(t_2) = v\right] \cdot \mathbb{P}\left[X_g(t_2) = v\right] \\ &\geq \frac{1}{2} \sum_{v \in V} \mathbb{P}\left[\tau_{u,g} \leq \tau_3 \mid X_g(t_2) = v\right] \cdot \frac{1}{2n} \\ &\geq \frac{1}{2} \cdot \mathbb{P}\left[\tau'_u \leq \tau_3\right],\end{aligned}$$

where τ'_u is the hitting time of vertex u for a random walk that starts from the stationary distribution at round t_2. Thus, by Lemma 14,

$$\mathbb{P}\left[\tau_{u,g} \leq \tau_3 \mid I_t, A'\right] \geq \frac{1}{2} \cdot \mathbb{P}\left[\tau'_u \leq \tau_3\right] \geq \frac{(1-\lambda) \cdot (\tau_3 - t_2)}{8n} = \frac{\eta(1-\lambda)\log n}{8n}.$$

Let $\mathcal{E}_u$ be the event that u is informed at round τ_3, which will happen if at least one of the agents in A' will visit u at or before round τ_3. Since the walks are independent,

$$\begin{aligned}\mathbb{P}\left[\mathcal{E}_u \mid I_t\right] &\geq \mathbb{P}\left[\mathcal{E}_u \mid I_t, |A'| \geq \alpha n/8\right] \cdot \mathbb{P}\left[|A'| \geq \alpha n/8 \mid I_t\right] \\ &\geq 1 - \left(1 - \frac{\eta(1-\lambda)\log n}{8n}\right)^{\alpha n/8} - e^{-\alpha n/32} \\ &\geq 1 - n^{-\eta\alpha(1-\lambda)/64} - e^{-\alpha n/32} \\ &\geq 1 - n^{-(c+1)},\end{aligned}$$

for a sufficiently large constant η. Using a union bound for all $u \in V$ completes the proof. ◀

Proof of Theorem 4. For a constant $c > 0$, we choose b by using Lemma 18, and then r by Lemma 17. We consider the TWEAKED$_{b,r}$ process. Lemma 16 implies that for some $\tau_1 = O(\log n)$, $|I_{\tau_1}| \geq b \cdot \log n$, with probability at least $1 - n^{-c}$. Next, for $i \in \{1, \ldots, \lceil \log_{1+\psi} n \rceil\}$, consider rounds $\tau_1 + i \cdot r$. By Lemma 18, for each i, either $|I_{\tau_1 + i \cdot r}| \geq n/2$ or $|I_{\tau_1+(i+1)\cdot r}| \geq (1+\psi) \cdot |I_{\tau_1+i\cdot r}|$, with probability at least $1 - n^{-(c+1)}$. By a union bound over all values of i, we have that $|I_{\tau_1+\tau_2}| \geq n/2$, for $\tau_2 = O(\log n)$, with probability at least $1-2n^{-c}$. Next, we apply Lemma 19 and union bound again, thus, showing that TWEAKED$_{b,r}$ informs all vertices of the graph in $\tau = \tau_1 + \tau_2 + \tau_3 = O(\log n)$ rounds, with probability at least $1 - 3n^{-c}$. Finally, by Lemma 17, TWEAKED$_{b,r}$ and VISIT-EXCHANGE are identical in the first $\tau = O(\log n)$ rounds, with probability at least $1 - n^{-c}$. Therefore, $T = O(\log n)$, w.h.p. ◀

7 Upper Bound for Trees

In this section we prove the upper bound part of Theorem 5. Recall, $R_{b,h}$ is a rooted b-ary tree, where each vertex at distance less than h from the root has b children, and all leaves are at distance h from the root; thus h is the *height* of the tree. The total number of vertices is $n = (b^{h+1} - 1)/(b - 1)$. The set of children of vertex u is denoted C_u. The set of descendants of u is denoted D_u; precisely, D_u contains the vertices in the subtree rooted u, including u itself. The height of that subtree is denoted h_u. We define the set $B_{u,l} = \{v \in D_u \mid h_v = h_u - l\}$, which contains all descendants of v at distance l from u. Finally, $Z_u(t)$ denotes the set of agents at vertex u at round t, and $Z_S(t) = \bigcup_{u \in S} Z(t)$ is the set of agents in the set $S \subseteq V$ at that round.

7.1 The Lucky-Gambler Process

We define an auxiliary process, called LUCKY-GAMBLER, which will be used in the analysis. The process has three parameters: two integers $m, k > 0$, and a probability $p < 1/2$. Consider a path graph P_m of length m, with vertices 0 up to m. For every integer $s \geq 0$, at round s exactly k *gamblers* appear on vertex 1 and make a biased random walk: for $0 < i < m$, the probability of moving from vertex i to $(i+1)$ and $(i-1)$ is $p_{i,i+1} = p$ and $p_{i,i-1} = 1 - p = q$, respectively. When the gambler reaches vertex 0 or m, it stops, i.e., $p_{0,0} = p_{m,m} = 1$ (states $0, m$ are absorbing). We will write LUCKY-GAMBLER(m, p, k) to explicitly state the parameters of the process.

For a vertex v of $R_{b,h}$, where $h_v \geq m$, we are going to couple the movement of the agents in part of the subtree of v, with the gamblers in LUCKY-GAMBLER. Using the coupling and the next lemmas, we argue that v receives agents at a constant rate. By carefully selecting the agents that are coupled, we can claim that agents arrive at constant rate to every vertex v on a given path to the root, independently for each vertex.

▶ **Lemma 20.** *If $p = 1/(b+1)$ and $k \geq \epsilon \cdot b^{m-1}$, for some constant $\epsilon > 0$, then there is a constant $\beta < 1$ such that for any round $t \geq 4m$ and positive integer Δ the probability that no gambler reaches vertex m during any round in $\gamma_0 = \{t, \ldots, t + \Delta - 1\}$ is at most $(1 - \beta)^{\Delta}$.*

▶ **Lemma 21.** *If $p = 1/(b+1)$ and $k \geq \kappa \cdot b^{m-1}$, for some integer κ, then there is a constant γ, such that for any integer $\tau \geq 8m$, at least $\gamma\kappa\tau$ gamblers reach vertex m in the first τ rounds, with probability at least $1 - e^{-\gamma\kappa\tau/4}$.*

To prove Lemmas 20 and 21 we will use the next two results for a single gambler g making a biased random walk on P_m starting at round 0. Let $X_g(t)$ be the position of gambler g at round t and let $\tau_g(i) = \min\{t \mid X_g(t) = i\}$ be the hitting time of vertex i of g. We denote the event that $\tau_g(m) < \tau_g(0)$ as $\mathcal{L}_g$, and we will say that g is *lucky* if it occurs.

▶ **Lemma 22** ([18, Chapter 14]). *If $p \neq q$, then for $0 < i < m$, $\mathbb{P}[\mathcal{L}_g \mid X_g(0) = i] = \frac{(q/p)^i - 1}{(q/p)^m - 1}$.*

▶ **Lemma 23.** *If $p < q$, then for $0 < i < m$, $\mathbb{E}[\tau_g(m) \mid \mathcal{L}_g, X_g(0) = i] \leq \frac{m-i}{q-p}$.*

Proof. By Bayes' theorem and Lemma 22, we can explicitly find the transition probabilities of the Markov chain of g, conditioned on $\mathcal{L}_g$. For $0 < j < m$,

$$\begin{aligned}\mathbb{P}\left[X_g(1) = j+1 \mid \mathcal{L}_g, X_g(0) = j\right] &= \frac{\mathbb{P}\left[X_g(1) = j+1, \mathcal{L}_g \mid X(0) = j\right]}{\mathbb{P}\left[\mathcal{L}_g \mid X_g(0) = j\right]}\\ &= p_{j,j+1} \cdot \frac{\mathbb{P}\left[\mathcal{L}_g \mid X_g(0) = j+1\right]}{\mathbb{P}\left[\mathcal{L}_g \mid X_g(0) = j\right]}\\ &= p \cdot \frac{(q/p)^{j+1} - 1}{(q/p)^j - 1} = p \cdot \left(\frac{q}{p} + \frac{(q/p) - 1}{(q/p)^j - 1}\right) \geq q.\end{aligned}$$

Consider a new random walk $X'(t)$ on $\mathbb{Z}$ with transition probabilities $p'_{j,j+1} = q$ and $p'_{j+1,j} = p$, for any $j \in \mathbb{Z}$. Let $\tau'(j)$ be the hitting time of $X'(t)$ of vertex j. The inequality above implies that $\tau(j)$ conditioned on $\mathcal{L}_g$ is stochastically dominated by $\tau'(j)$, when both start from the same vertex i, and thus,

$$\mathbb{E}\left[\tau_g(m) \mid \mathcal{L}_g, X_g(0) = i\right] \leq \mathbb{E}\left[\tau'(m) \mid X'(0) = i\right] = \frac{m-i}{q-p}. \qquad \blacktriangleleft$$

Proof of Lemma 20. For $s \geq 0$ and $1 \leq i \leq k$, let $g_{s,i}$ be the ith gambler that starts its walk at round s at vertex 1. Let $\tau_{s,i} = \tau_{g_{s,i}}$ be defined as for the single gambler g above. Clearly, $\tau_{s,i}(j) - s$ and $\tau_g(j)$ are identically distributed, if $X_g(0) = 1$. We also extend the definition of γ_0, letting $\gamma_s = \{t - s, \ldots, t + \Delta - s - 1\}$.

We would like to study the number of lucky gamblers that reach m at rounds in γ_0. Consider first a "toy" example, which assumes that for each s, exactly one gambler is lucky among the k gamblers that start their walk at round s. Suppose that g'_s is that lucky gambler. We study the expected number of these agents that reach m during the rounds in γ_0:

$$\mathbb{E}\left[\sum_{s\geq 0} \mathbf{1}_{\{\tau_{g'_s}(m) \in \gamma_0\}} \mid \mathcal{L}_{g'_s} \text{ for } s \geq 0\right] = \sum_{s=0}^{t+\Delta} \mathbb{P}\left[\tau_{g'_s}(m) \in \gamma_0 \mid \mathcal{L}_{g'_s}\right] = \sum_{s=0}^{t+\Delta} \mathbb{P}\left[\tau_g(m) \in \gamma_s \mid \mathcal{L}_g\right].$$

The setup in the "toy" example is unlikely to occur, however, we use it as a motivation to lower bound the last quantity, which will be used in the main part of the proof.

$$\begin{aligned}\sum_{s=0}^{t+\Delta} \mathbb{P}\left[\tau_g(m) \in \gamma_s \mid \mathcal{L}_g\right] &= \sum_{l=0}^{\Delta-1} \sum_{\substack{0 \leq s \leq t+\Delta \\ s \equiv l \pmod{\Delta}}} \mathbb{P}\left[\tau_g(m) \in \gamma_s \mid \mathcal{L}_g\right],\\ &\qquad \text{the inner sum is over every } \Delta\text{th summand},\\ &\geq \sum_{l=0}^{\Delta-1} \mathbb{P}\left[\tau_g(m) < t \mid \mathcal{L}_g\right], \quad \text{by union of disjoint events},\\ &= \Delta \cdot \mathbb{P}\left[\tau_g(m) < t \mid \mathcal{L}_g\right]\\ &\geq \Delta \cdot \left(1 - \frac{\mathbb{E}\left[\tau_g(m) \mid \mathcal{L}_g\right]}{t}\right), \quad \text{by Markov's inequality},\\ &\geq \Delta \cdot \left(1 - \frac{m \cdot (b+1)}{t \cdot (b-1)}\right), \quad \text{by Lemma 23 as } q - p = \frac{b-1}{b+1},\\ &\geq \Delta \cdot \left(1 - \frac{b+1}{4(b-1)}\right), \quad \text{since } t \geq 4m,\\ &\geq \Delta/4.\end{aligned}$$

We can now bound the probability that no agent visits vertex m between rounds t and $t+\Delta$:

$$\begin{aligned}
\mathbb{P}\left[\bigcap_{\substack{0\le s\le t+\Delta\\ 1\le i\le k}} \{\tau_{s,i}(m)\notin\gamma_0\}\right] &= \prod_{s=0}^{t+\Delta} \left(\mathbb{P}\left[\tau_{s,i}(m)\notin\gamma_0\right]\right)^k, \quad \text{by independence of the walks,}\\
&= \prod_{s=0}^{t+\Delta} \left(\mathbb{P}\left[\tau_g(m)\notin\gamma_s\right]\right)^k\\
&= \prod_{s=0}^{t+\Delta} \left(1-\mathbb{P}\left[\mathcal{L}_g\right]\cdot\mathbb{P}\left[\tau_g(m)\in\gamma_s \mid \mathcal{L}_g\right]\right)^k\\
&= \prod_{s=0}^{t+\Delta} \left(1-\frac{b-1}{b^m-1}\cdot\mathbb{P}\left[\tau_g(m)\in\gamma_s \mid \mathcal{L}_g\right]\right)^k, \text{ by Lemma 22,}\\
&\le \prod_{s=0}^{t+\Delta} \exp\left(-\frac{k\cdot(b-1)}{b^m-1}\cdot\mathbb{P}\left[\tau_g(m)\in\gamma_s \mid \mathcal{L}_g\right]\right)\\
&\le \exp\left(-\epsilon\cdot\frac{b^{m-1}(b-1)}{b^m-1}\cdot\sum_{s=0}^{t+\Delta}\mathbb{P}\left[\tau_g(m)\in\gamma_s \mid \mathcal{L}_g\right]\right)\\
&\le \exp\left(-\frac{\epsilon\Delta}{8}\right), \quad \text{by the analysis of the toy example.}
\end{aligned}$$

◀

Proof of Lemma 21. For $i\in\{1,\ldots,k\}$, consider a gambler $g_{s,i}$ that starts its walk at round $s\le t/2$. Let $X_{s,i}=1$ if $g_{s,i}$ is lucky and reaches vertex m before round t, i.e., $\tau_{s,i}(m)\le t$. Since $\tau_{s,i}(m)-s$ and $\tau_g(m)$ are identically distributed,

$$\begin{aligned}
\mathbb{P}\left[X_{s,i}=1\right] &= \mathbb{P}\left[\tau_{s,i}(m)\le t \mid \mathcal{L}_{g_{s,i}}\right]\cdot\mathbb{P}\left[\mathcal{L}_{g_{s,i}}\right]\\
&= \mathbb{P}\left[\tau_g(m)+s\le t \mid \mathcal{L}_g\right]\cdot\mathbb{P}\left[\mathcal{L}_g\right]\\
&\ge \mathbb{P}\left[\tau_g(m)\le t/2 \mid \mathcal{L}_g\right]\cdot\mathbb{P}\left[\mathcal{L}_g\right], \quad \text{since } s\le t/2,\\
&\ge \left(1-\frac{\mathbb{E}\left[\tau_g(m)\mid\mathcal{L}_g\right]}{t/2}\right)\cdot\mathbb{P}\left[\mathcal{L}_g\right], \quad \text{by Markov's inequality,}\\
&\ge \left(1-\frac{2m(b+1)}{t(b-1)}\right)\cdot\frac{b-1}{b^m-1}, \quad \text{by Lemmas 22 and 23,}\\
&\ge \frac{1}{8\cdot b^{m-1}}.
\end{aligned}$$

If N is the number of gamblers that arrive at vertex m at before round t, then

$$\mathbb{E}\left[N\right]\ge\sum_{s=0}^{t/2}\sum_{i=1}^{k}\mathbb{E}\left[X_{s,i}\right]\ge\frac{\kappa b^{m-1}t}{2}\cdot\frac{1}{8b^{m-1}}=\frac{\kappa t}{16}.$$

Since the variables $X_{s,i}$ are independent, we can prove the lemma by an application of Chernoff bound. ◀

7.2 Analysis

We define another auxiliary process, called TWEAKED, which is a slight modification of the original VISIT-EXCHANGE process. Let m be the *smallest* integer such that $b^m\ge\mu\cdot\ln n$ for a constant μ to be defined later, and let $k=\lceil\alpha\cdot b^m/8\rceil$. Consider a vertex u of the tree, such that $h_u\ge m$, and recall that $B_{u,m}$ is the set of descendants of u at distance m. Let v be one of the children of u and define $Z'_{u,v}(t)$ be the set of agents that are in $B_{u,m-1}$ at round t and were in $B_{u,m} \setminus B_{v,m-1}$ the round before, i.e.,

$$Z'_{u,v}(t) = Z_{B_{u,m-1}}(t) \cap Z_{B_{u,m} \setminus B_{v,m-1}}(t-1).$$

For a round $t \geq 0$ let $q_{u,v}(t)$ be the smallest non-negative integer for which

$$|Z'_{u,v}(t)| + q_{u,v}(t) \geq \left\lceil \frac{\alpha}{8} \cdot |B_{u,m}| \right\rceil = \left\lceil \frac{\alpha}{8} \cdot b^m \right\rceil = k.$$

To construct TWEAKED we add exactly $q_{u,v}(t)$ agents in $B_{u,m-1}$ at round t (it is not important to which vertices in $B_{u,m-1}$ these agents are added).

To motivate the construction of TWEAKED, consider a vertex u and its child v, such that $m \leq h_u < h$. In round t of TWEAKED, there are at least k agents at vertices in $B_{u,m} \setminus B_{v,m-1}$ (of height $h_u - m$) that move closer to u in the next round. This allows us to couple these agents to that of gamblers in a LUCKY-GAMBLER$(m+1, 1/(b+1), k)$ process, and use our results from Section 7.1 to show that agents arrive at the parent of u at a constant rate. A key insight is that by not considering agents that are in descendants of v, the same argument can be made for vertex v, independently of u, if $h_v \geq m$ too. By repeating this argument, we show that in $O(\log n)$ rounds all vertices of height at least m are informed once one such vertex is informed. TWEAKED and LUCKY-GAMBLER are also used to analyse the spread of the message in the vertices of height at most m.

Using a Chernoff bound we can show that TWEAKED and VISIT-EXCHANGE are equivalent in the first polynomially many rounds, w.h.p.

▶ **Lemma 24.** *The probability that no agent is added in the* TWEAKED *process in the first r rounds is at least $1 - r \cdot n^{-\frac{\alpha \cdot \mu}{32}+1}$.*

Proof. Fix a vertex u with $h_u \geq m$ and a round $t \geq 0$ of VISIT-EXCHANGE. For an agent $g \in A$, let X_g be the indicator random variable that $g \in Z'_u(t)$. Since the agents are distributed by the stationary distribution, the probability that in any round, agent g traverses an edge in a particular direction is exactly $1/(2 \cdot |E|) = 1/(2 \cdot (n-1))$. Thus,

$$\mathbb{E}\left[|Z'_u(t)|\right] = \sum_{g \in A} \mathbb{P}\left[X_g = 1\right] = \frac{\alpha n \cdot |B_{u,m} \setminus B_{v,m-1}|}{2 \cdot (n-1)} \geq \frac{\alpha}{2} \cdot (b-1) \cdot b^{m-1} \geq \frac{\alpha}{4} \cdot b^m.$$

By an application of a Chernoff bound we get that

$$\mathbb{P}\left[b_{u,v}(t) > 0\right] = \mathbb{P}\left[|Z'_u(t)| < \frac{\alpha}{8} \cdot b^m\right] \leq \exp\left(-\frac{\alpha}{32} \cdot b^m\right) \leq n^{-\frac{\alpha \cdot \mu}{32}}.$$

By taking a union bound over all rounds t up until r and edges uv, we complete the proof. ◀

We will use the same notation for TWEAKED and VISIT-EXCHANGE processes.

▶ **Lemma 25.** *Let u be any vertex of the tree $R_{b,h}$ such that $h_u \geq m$. For any constant $c > 0$, if u is informed, then after $O(\log n)$ rounds of* TWEAKED *the root ρ of $R_{b,h}$ gets informed, with probability at least $1 - n^{-c}$.*

Proof. Consider the path $u = u_1, \ldots, u_l = \rho$ from u to the root of the tree. Due to the symmetry of the tree, we can assume that the path is the "leftmost" path of the tree, i.e., for any $i \geq 1$, u_{i-1} is the leftmost child of u_i (for consistency, we let u_0 be the leftmost child of u_1). Roughly speaking, we show that for any i, the number of rounds between two consecutive visits to u_i (by a certain subset of agent) follows a geometric distribution, independently of the other $u_{i'}$. To that end, we couple the movement of agents of TWEAKED to $l-1$ independent instances of process LUCKY-GAMBLER$(m+1, 1/(b+1), k)$, one corresponding to each of the vertices u_i for $1 \leq i < l$.

Next we give some definitions and describe the coupling for a fixed i. For simplicity, define $B_i = B_{u_i,m}$ and $B'_i = B_i \setminus B_{u_{i-1},m-1} = \bigcup_{v \in C_{u_i} \setminus \{u_{i-1}\}} B_{v,m-1}$. I.e., B_i is the set of descendants of u_i at distance m from it, and to get B'_i we remove the descendants of u_{i-1} from B_i. Let $g_1, \ldots, g_{z_{i,t}}$ be the agents in TWEAKED that were at B'_i in round $t-1$ and moved closer to the root in the next round. By definition of TWEAKED, there are at least $k = \lceil \alpha \cdot b^m/8 \rceil$ such agents.

In the LUCKY-GAMBLER$(m+1, 1/(b+1), k)$ process that corresponds to vertex u_i, we start k gamblers in round t, denoted $g'_1, \ldots, g'_k$. For each $1 \leq j \leq k$, and for each round $t' \geq t$ until g_j reaches u_{i+1} or any vertex in B_i, the walks g_j and g'_j are coupled: if g_j moves closer to the root then g'_j moves to the right on the path, and if g_j moves away from the root, g'_j moves left. If g_j is at u_{i+1} or in $B_{u_i,m}$, then by the coupling, g'_j has finished its walk at one of the endpoints of the path. Before this happens we say that g_j is *i-coupled*.

Let $t_1 = 4 \cdot (m+1)$, and let t_{i+1} be the first round after t_i when u_{i+1} receives an i-coupled agent from u_i. Now, notice that by construction no agent can be i-coupled and i'-coupled at the same time for $i' \neq i$. It implies that the rounds when u_{i+1} receives i-coupled agents are independent from the walks of i'-coupled agents. On the other hand the walks of i-coupled agents are coupled with an independent LUCKY-GAMBLER process thus, Lemma 20 implies

$$\mathbb{P}\left[t_{i+1} - t_i \leq s \mid t_1, \ldots, t_i\right] = (1-\beta)^s = \mathbb{P}\left[F_i \geq s\right],$$

where $F_i \sim \mathit{Geom}(\beta)$, $1 \leq i < l$, are a collection of independent geometric random variables with success probability β. If τ_ρ is the round when the root is informed then $\tau_\rho \leq t_l = t_1 + \sum_{i=1}^{l-1}(t_{i+1} - t_i)$. It follows that $(\tau_\rho - t_1)$ is stochastically dominated by $F = \sum_{i=1}^{l-1} F_i$, and from a Chernoff bound for the sum of independent geometric random variables (Lemma 43),

$$\mathbb{P}\left[\tau_\rho \geq f + t_1\right] \leq \mathbb{P}\left[F \geq f\right] \leq e^{-f \cdot \beta/8},$$

for any $f \geq 2h/\beta$. Since $t_1 = O(h)$, we can take a large enough $f = O(\log n)$, completing the proof. ◀

Next we prove that if vertex u of height $h_u = m$ is informed, then after at most $O(m \ln n)$ rounds a given leaf v in u's subtree becomes informed, w.h.p. For that, we first show that there are at least $\Theta(m \ln n)$ visits to u in those rounds (possibly multiple times by the same agent). Using a lower bound on the probability that an agent that is at u visits v before returning to u, we can show that one of these agents will visit v in $O(m \ln n)$ rounds, w.h.p.

▶ **Lemma 26.** *Let u be such that $h_u = m$. For any constant $c > 0$, there is a round $\tau = O(m \ln n)$ such that in the first τ rounds of* TWEAKED*, u is visited at least $c \cdot mb \cdot \ln n$ times, with probability at least $1 - n^{-cmb}$.*

Proof. For a round t, let $g_1, \ldots, g_{z_{u,t}}$ be the agents that are in $B_{u,m-1}$ at round t, and have also been at the leaf vertices $B_{u,m}$ in the previous round. By the definition of TWEAKED, $z_{u,t} \geq k$, where $k = \lceil \alpha b^m/8 \rceil$. We construct an instance of LUCKY-GAMBLER$(m, 1/(b+1), k)$ as follows. If $g'_1, \ldots, g'_k$ are the gamblers that started their walk at round t, then for each $1 \leq j \leq k$, the walk of agent g_j is coupled with the walk of the gambler g'_k: If g_j moves closer to the root of the tree, then g'_j moves right on the path and left otherwise. The coupling ends when g'_j arrives at either vertex 0 or m of its path. That corresponds to g_j either visiting a leaf vertex in $B_{u,m}$ or visiting vertex u.

Consider the first τ rounds of TWEAKED. Since $k \geq \alpha b^m/8$, we can apply Lemma 21 with parameter $\kappa = \alpha b/8$ to the coupled LUCKY-GAMBLER process. Let γ be the constant guaranteed by the lemma and let $\tau = \frac{8c}{\alpha\gamma} \cdot m \ln n$. Lemma 21 implies that in the first τ rounds of LUCKY-GAMBLER there are at least $\gamma\kappa\tau = c \cdot mb \cdot \ln n$ lucky gamblers, with probability at least $1 - e^{-\gamma\kappa\tau/4} = 1 - e^{-cmb\ln n} = 1 - n^{-cmb}$. Since each lucky gambler corresponds to a single visit to u by some agent, we complete the proof. ◀

▶ **Lemma 27.** *Let u be such that $h_u = m$ and let v be a leaf in the subtree of u. For any constant $c_l > 0$, if vertex u is informed then after at most $O(m \ln n)$ rounds of* TWEAKED*, vertex v is informed with probability at least $1 - n^{-c_l}$.*

Proof. Let τ be the round guaranteed by Lemma 26 for a constant $c > 0$. If after the first τ rounds of TWEAKED, there have been fewer than $cmb \ln n$ visits to u, then we add a minimal number of agents to u at round τ to have at least $cmb \ln n$ agents there. We call the resulting process TWEAKED$_u$. By Lemma 26 and an application of union bound over the first $\log^2 n = \omega(m \ln n)$ rounds, TWEAKED$_u$ and TWEAKED are identical in the first $\Theta(m \ln n)$ rounds of execution with probability at least $1-n^{-cmb} \log^2 n$. We therefore analyse TWEAKED$_u$.

For a round $t \leq \tau$, consider an agent g that visits u at round t. Let $\mathcal{D}_{g,t}$ be the event that g moves to one of u's children at round $t + 1$. Let also $\mathcal{E}_{g,t}$ be the event that g visits v before returning to u, and before round $\tau' = \tau + 8mb^{m-1}$. Clearly, $\mathcal{E}_{g,t}$ implies $\mathcal{D}_{g,t}$, and $\mathbb{P}[\mathcal{D}_{g,t}] = \frac{b}{b+1}$. Also, we can show that $\mathbb{P}[\mathcal{E}_{g,t} \mid \mathcal{D}_{g,t}] \geq 1/(12mb)$, by analysing a single random walk in $R_{b,m}$ that starts in the root of the tree (Lemma 36). Therefore,

$$\mathbb{P}[\mathcal{E}_{g,t}] = \mathbb{P}[\mathcal{E}_{g,t} \cap \mathcal{D}_{g,t}] = \mathbb{P}[\mathcal{D}_{g,t}] \cdot \mathbb{P}[\mathcal{E}_{g,t} \mid \mathcal{D}_{g,t}] \geq \frac{b}{b+1} \cdot \frac{1}{12mb} \geq \frac{1}{18mb}.$$

The probability that v is not visited by any informed agent before round τ' is at most

$$\mathbb{P}\left[\bigcap_{t\leq\tau,\, g\in Z_u(t)} \neg\mathcal{E}_{g,t}\right] \leq \left(1 - \frac{1}{18mb}\right)^{cmb\ln n} \leq e^{-c\ln n/18} \leq n^{-c/18} \leq n^{-c_l-1},$$

for a large enough constant c. Notice that $\tau' = \tau + 8mb^{m-1} = O(m \ln n)$ by the definition of m. Since TWEAKED and TWEAKED$_u$ are identical in the first $\log^2 n$ rounds with probability at least $1 - n^{-cmb} \log^2 n$, v will be informed in $O(m \ln n)$ rounds in TWEAKED, with probability at least $1 - n^{-c_l-1} - n^{-cmb} \log^2 n \geq 1 - n^{-c_l}$. ◀

Proof of the Upper Bound of Theorem 5. We will use the following simple symmetry lemma, which holds for any graph (Lemma 39): If $T_{u,v}$ is the number of rounds of VISIT-EXCHANGE until vertex v is informed when the information originates at u, then the random variables $T_{u,v}$ and $T_{v,u}$ have the same distribution.

Consider the TWEAKED process, and suppose that the source of the information is vertex u with $h_u = m$, for m as defined at the beginning of Section 7.2. By Lemma 25, for an arbitrary constant c, there is $T_1 = O(\log n)$ such that the root ρ is informed by time T_1, with probability at least $1 - n^{-c}$. Lemma 24 then implies that the same bound T_1 holds for the VISIT-EXCHANGE process, with probability $p \geq 1 - n^{-c} - n^{-\alpha\mu/32}$, for an arbitrary large μ. From the symmetry lemma above, it follows that if ρ is the initial source of the information instead, then u becomes informed within T_1 rounds of VISIT-EXCHANGE with the same probability $p \geq 1 - n^{-c} - n^{-\alpha\mu/32}$.

Suppose again that information originates at some u with $h_u = m$, and let v be any leaf that is a descendant of u. From Lemma 27 and Lemma 24, for an arbitrary constant

c, there is some $T_2 = O(m \log n)$, such that v gets informed after at most T_2 rounds of VISIT-EXCHANGE, with probability at least $1 - n^{-c} - n^{-\alpha\mu/32}$.

Combining the above we obtain that if ρ is the source of the information, then any given leaf v is informed after at most $T_1 + T_2$ rounds of VISIT-EXCHANGE, with probability at least $1 - 2n^{-c} - 2n^{-\alpha\mu/32}$. And by a union bound, all leaves (and thus all vertices) are informed within $T_1 + T_2$ rounds with probability at least $1 - 2n^{-c+1} - 2n^{-\alpha\mu/32+1}$.

Finally, by employing the symmetry argument above again, we obtain that for any source vertex (not just ρ), all vertices are informed within $2(T_1+T_2)$ rounds with probability at least $1-4n^{-c+1} - 4n^{-\alpha\mu/32+1}$. Since $T_1+T_2 = O(\log n + m \log n) = O(\log n + \log_b \log n \cdot \log n) = O(\log n + h \log h)$, the theorem follows. ◀

8 Lower Bound for Binary Trees

We prove the lower bound part of Theorem 5, i.e., the spreading time of VISIT-EXCHANGE on a binary tree is $\Omega(\log n \cdot \log\log n)$, w.h.p.

Proof of the Lower Bound of Theorem 5. We show that there is a leaf vertex that is not visited by any agent in the first $\tau = c \ln n \cdot \ln\ln n$ rounds of VISIT-EXCHANGE, w.h.p., where c is a small enough constant, to be determined later. For convenience, we assume that τ is even. For a fixed leaf vertex v and an agent g, let $N_g(v)$ be the number of times g visits v in the rounds $0, \ldots, \tau - 1$ of its walk. Since g starts from a stationary distribution,

$$\mathbb{E}\left[N_g(v)\right] = \pi(v) \cdot \tau = \tau/(2(n-1)).$$

Let τ_v be the first time when g visits v. Then,

$$\begin{aligned}\mathbb{E}\left[N_g(v) \mid \tau_v < \tau\right] &\geq \mathbb{E}\left[N_g(v) \mid \tau_v < \tau/2\right] \cdot \mathbb{P}\left[\tau_v < \tau/2 \mid \tau_v < \tau\right] \\ &\geq \mathbb{E}\left[N_g(v) \mid \tau_v < \tau/2\right] \cdot (1/2),\end{aligned}$$

where the second inequality holds because $\mathbb{P}\left[\tau_v < \tau\right] \leq 2\mathbb{P}\left[\tau_v < \tau/2\right]$, as g is equally likely to visit v in the intervals $0, \ldots, \tau/2 - 1$ and $\tau/2, \ldots, \tau - 1$ for g starts its walk from stationarity. From Lemma 38, if $X_g(t)$ denotes the position of the random walk of g at t,

$$\begin{aligned}\mathbb{E}\left[N_g(v) \mid \tau_v < \tau/2\right] &\geq \mathbb{E}\left[N_g(v) \mid \tau_v = \tau/2 - 1\right] \geq \sum_{t=\tau/2}^{\tau-1} \mathbb{P}\left[X_g(t) = v \mid \tau_v = \tau/2 - 1\right] \\ &\geq \sum_{t=\tau/2}^{\tau-1} \frac{1}{32 \cdot (t - (\tau/2 - 1))} \geq \frac{\ln(\tau/2)}{32} \geq \frac{\ln\ln n}{32},\end{aligned}$$

for n sufficiently large. It follows

$$\mathbb{P}\left[\tau_v < \tau\right] = \mathbb{P}\left[N_g(v) \geq 1\right] = \frac{\mathbb{E}\left[N_g(v)\right]}{\mathbb{E}\left[N_g(v) \mid N_g(v) \geq 1\right]} \leq \frac{\tau/(2(n-1))}{(1/2)\cdot(\ln\ln n/32)} = \frac{32c \cdot \ln n}{n-1}.$$

Returning to the case of n agents, for the leaf v,

$$\mathbb{P}\left[v \text{ not visited by any agent}\right] \geq \left(1 - \frac{32c \cdot \ln n}{n-1}\right)^n \geq \frac{1}{2} \cdot e^{-32c \cdot \ln n} = \frac{1}{2} \cdot n^{-32c}.$$

Thus, if X is the number of leaves that are *not* visited by any agent, then $\mathbb{E}\left[X\right] \geq n^{1-32c}/4$, as there are at least $n/2$ leaves. We can now use the method of bounded differences [14, Sec. 5.4] to give a lower bound on the probability that at least one vertex is not visited by any agent.

If L_g is the set of leaves that g visits, then X can be written as a function of independent variables L_g. Notice that changing L_g can change X by at most τ. Thus,

$$\mathbb{P}\left[X \leq \frac{\mathbb{E}\left[X\right]}{2}\right] \leq \exp\left(-\frac{\mathbb{E}\left[X\right]^2}{2 \cdot n \cdot \tau^2}\right) \leq \exp\left(-\frac{n^{1-64c}}{64 \cdot \tau^2}\right).$$

By taking $c < 1/64$ we get that, w.h.p., there is at least one leaf that has not been visited by any agent. ◀

9 Grids

In this section we prove Theorem 6. To motivate our analysis technique, which is closely related to that of [24, 25], we consider a path graph and assume that the source vertex is at one end. At any round, the agents are located according to the stationary distribution, therefore, w.h.p., there is a sub-path of logarithmic length without any agents. When the most recently informed vertex belongs to such a sub-path, the progress of informing new vertices is delayed, hence we have to argue that such situations are rare. We will show that for the majority of the rounds up to $T^* = c\ell$, there is an agent at most constant steps away from the most recently informed vertex. We tessellate the space-time into square blocks of constant length Δ_1, both in space and time, but due to dependencies, we cannot argue directly that most of them are densely populated, or *good* as we call them. We build further tessellations of the space-time into square blocks of increasing sizes $\Delta_r = O(\text{poly}\log n)$, for a scale parameter $r \in \{1, \ldots, R\}$. It is easy to argue that the blocks of size Δ_R in the coarsest tessellation are all good, w.h.p. Then we show that with a sufficiently high probability a good block does not contain any bad blocks in a finer tessellation. Repeating this argument recursively, from the largest scale to the smallest one, we prove our desired result.

9.1 Definitions

Recall that $G_{k,n}$ is a k-dimensional n-vertex grid, and let $\ell = n^{1/k} - 1$ denote the length of each side of the grid. In the analysis we assume $k \geq 1$ is a constant, thus, $\text{diam}(G_{k,n}) = k\ell = \Theta(\ell)$. Recall, that the total number of agents is αn, for a constant $\alpha > 0$. Our goal will be to prove that $T \leq T^* = c\ell$, w.h.p., for a large enough constant c that will be specified in the proof. We represent V as a collection of k-dimensional vectors $\mathbf{x} = (x_1, \ldots, x_k)$ where $x_j \in \{0, \ldots, \ell\}$ for $1 \leq j \leq k$. For $\mathbf{x} = (x_1, ..., x_k) \in \mathbb{Z}^k$ and an integer z, we will write $\mathbf{x} + z = (x_1 + z, \ldots, x_k + z)$. Also, for $\mathbf{y} = (y_1, ..., y_k) \in \mathbb{Z}^k$, will write $\mathbf{x} < \mathbf{y}$ to denote that $x_i < y_i$, for all $1 \leq i \leq k$; and similarly for $\mathbf{x} \leq \mathbf{y}$.

Partitioning Space-Time. We prove our result using a multi-scale argument, with a scaling parameter $r \in \{1, 2, \ldots, R\}$, for some $R = \Theta(\log\log n)$ to be fixed later. For each r, let $\Delta_r = C^{4kr}$ for a constant even integer C to be specified later. During our analysis, for each scale r, we only consider rounds $s \cdot \Delta_r$ for an integer $s \geq 0$. For a vector $\mathbf{i} \in \mathbb{Z}^k$ and an integer s, we define the following sets in space and space-time, respectively:

$$S_r(\mathbf{i}) = \{\mathbf{x} \in \mathbb{Z}^k \mid \mathbf{i} \cdot \Delta_r \leq \mathbf{x} < (\mathbf{i} + \mathbf{1}) \cdot \Delta_r\}, \quad \tilde{S}_r(\mathbf{i}, s) = S_r(\mathbf{i}) \times \{s \cdot \Delta_r\}.$$

The collection of sets $\{S_r(\mathbf{i})\}_{\mathbf{i} \in \mathbb{Z}^k}$ partitions $\mathbb{Z}^k$. Additionally, we define extended versions of these sets:

$$B_r(\mathbf{i}) = \{\mathbf{x} \in \mathbb{Z}^k \mid (\mathbf{i} - \mathbf{3}) \cdot \Delta_r \leq \mathbf{x} < (\mathbf{i} + \mathbf{4}) \cdot \Delta_r\}, \quad \tilde{B}_r(\mathbf{i}, s) = B_r(\mathbf{i}) \times \{s \cdot \Delta_r\}.$$

The second definitions, in space-time, are used as a shorthand, so instead of saying “agents in $S_r(\mathbf{i})$ at round $s \cdot \Delta_r$,” we can say “agents at $\tilde{S}_r(\mathbf{i}, s)$.”

We call $\tilde{B}_r(\mathbf{i}, s-1)$ the *base* of $\tilde{S}_r(\mathbf{i}, s)$. If an agent is at $\tilde{S}_r(\mathbf{i}, s)$, then it must have also been at its base. The *parent* of $\tilde{S}_r(\mathbf{i}, s)$ is the set $\tilde{S}_{r+1}(\mathbf{j}, l)$ corresponding to the unique pair $(\mathbf{j}, l)$ such that $S_r(\mathbf{i}) \subset S_{r+1}(\mathbf{j})$ and $s \cdot \Delta_r \in [l \cdot \Delta_{r+1}, (l+1) \cdot \Delta_{r+1})$. Correspondingly, $\tilde{S}_r(\mathbf{i}, s)$ is one of the *children* of $\tilde{S}_{r+1}(\mathbf{j}, l)$.

Let $V' = \prod_{j=1}^{s}\{6\Delta_R, \ldots, \ell - 6\Delta_R\}$ be the “central” part of V. We will only consider sets $\tilde{S}_r(\mathbf{i}, s)$ for which it holds $\tilde{S}_r(\mathbf{i}, s) \subset V' \times [\Delta_R, T^* - \Delta_R]$. If P is the set of pairs $(\mathbf{i}, s)$ that satisfy the last relation, then for any $(\mathbf{i}, s) \in P$, we have $\tilde{B}_r(\mathbf{i}, s) \subset V \times [0, T^*]$, and moreover if $r < R$, then $\tilde{S}_r(\mathbf{i}, s)$ has a parent $\tilde{S}_{r+1}(\mathbf{j}, l)$ with $(\mathbf{j}, l) \in P$.

Good and Bad Sets. Let $\gamma_R = \alpha/2$. For $r < R$, define $\gamma_r = \gamma_{r+1} \cdot (1 - C^{-(r+1)/8})$. Then,

$$\begin{aligned}
\gamma_1 &= \gamma_R \cdot \prod_{j=1}^{R-1}(1 - C^{-(j+1)/8}) \\
&\geq \gamma_R \cdot \left(1 - \sum_{j=1}^{R-1} C^{-(j+1)/8}\right), \text{ by Weierstrass' inequality,} \\
&\geq \gamma_R \cdot \left(1 - \frac{C^{-1/4}}{1 - C^{-1/8}}\right) \\
&\geq \gamma_R/2,
\end{aligned}$$

for $C \geq 256$. Since $\gamma_r \geq \gamma_{r-1}$ for any $r \geq 2$, we have that for any r, $\gamma_r \in [\alpha/4, \alpha/2]$.

For any set of vertices S, let $N(S, t)$ be the number of agents in S at round t, and for $\mathbf{x} \in V$, we write $N(\mathbf{x}, t) = N(\{\mathbf{x}\}, t)$. For a space-time set $\tilde{S}$, $N(\tilde{S}) = \sum_{(\mathbf{x},t)\in\tilde{S}} N(\mathbf{x}, t)$. Next, for any $\mathbf{x} \in \mathbb{Z}^k$ and integer s define

$$Q_r(\mathbf{x}) = \{\mathbf{y} \in \mathbb{Z}^k \mid \mathbf{x} \leq \mathbf{y} < \mathbf{x} + C^r\}, \quad \tilde{Q}_r(\mathbf{x}, s) = Q_r(\mathbf{x}) \times \{s \cdot \Delta_r\}.$$

We say that $\tilde{S}_r(\mathbf{i}, s)$ is *good*, if for every $\mathbf{x}$ such that $\tilde{Q}_r(\mathbf{x}, s-1) \subset \tilde{B}_r(\mathbf{i}, s-1)$,

$$N(\tilde{Q}_r(\mathbf{x}, s-1)) \geq \gamma_r \cdot |Q_r(\mathbf{x})| = \gamma_r \cdot C^{kr}.$$

Otherwise, $\tilde{S}_r(\mathbf{i}, s)$ is *bad*.

9.2 Bounding the Probability of Having a Bad Child

In this section we consider a set $\tilde{S}_{r+1}(\mathbf{i}, s)$ for a scale $r \geq 1$ and an integer $s \geq 1$. Our goal is to show that if a set $\tilde{S}_{r+1}(\mathbf{i}, s)$ is a good set, then all its children are also good with high enough probability. To achieve that, we first show that in expectation there are sufficiently many agents in each set $Q_r(\mathbf{y}, t')$ that is contained in a base of a child. Then, by the independence of the walks and an application of a Chernoff bound, we can lower bound the desired probability. This result holds given any execution of the walks until round $(s-1)\Delta_{r+1}$, given by a σ-field $\mathcal{K}_{r+1}(s-1)$, which allows us to apply it for a number of (space- and time-separated) sets at once.

▶ **Lemma 28.** *There is a constant $C_1 > 0$, such that for any $C \geq C_1$ and scale $r \geq 1$, if $\tilde{S}_{r+1}(\mathbf{i}, s)$ is good, then for any even integer $u \in [\Delta_{r+1} - \Delta_r, 2\Delta_{r+1})$ and any vertex $\mathbf{y}$ with $Q_r(\mathbf{y}) \subset \{\mathbf{x} \mid \mathbf{i} \cdot \Delta_{r+1} - 3 \cdot \Delta_r \leq \mathbf{x} < (\mathbf{i}+\mathbf{1}) \cdot \Delta_{r+1} + 3 \cdot \Delta_r\}$,*

$$\mathbb{E}\left[N(Q_r(\mathbf{y}), t+u)\right] \geq \gamma_{r+1} \cdot C^{kr} \cdot \left(1 - C^{-(r+1)/2}\right),$$

where $t = (s-1)\Delta_{r+1}$.

Proof. Notice that only agents that are at $\tilde{B}_{r+1}(\mathbf{i}, s-1)$ can be at $Q_r(\mathbf{y})$ at round $t+u$. For $\mathbf{j} \in \mathbb{Z}^k$, define $\mathbf{x_j} = \mathbf{y} + \mathbf{j} \cdot C^{r+1}$. Construct a partition of $\mathbb{Z}^k$ into a grid of blocks $M_\mathbf{j}$ which have corners at vertices $\mathbf{x_j}$:

$$M_\mathbf{j} = \{\mathbf{x} \mid \mathbf{x_j} \leq \mathbf{x} < \mathbf{x_{j+1}}\}.$$

Each set $M_\mathbf{j}$ contains $C^{k(r+1)}$ vertices. Notice that $\mathbf{y}$ is also a corner for some of the blocks ($2k$ of them). We will only consider the set J of indices $\mathbf{j}$ such that $M_\mathbf{j} \subset B_{r+1}(\mathbf{i})$. Since $\tilde{S}_{r+1}(\mathbf{i}, s)$ is good, $N(M_\mathbf{j}, t) \geq \gamma_{r+1} C^{k(r+1)}$ by definition.

Let $W(u)$ be the position at round u of a random walk in $\mathbb{Z}^k$ starting at $\mathbf{0}$. For a vertex $\mathbf{x} \in V$, let $W_\mathbf{x}(u)$ be the position at round u of a random walk in G that starts at $\mathbf{x}$. Then,

$$\begin{aligned}
\mathbb{E}\left[N(Q_r(\mathbf{y}), t+u)\right] &\geq \sum_{\mathbf{j}\in J} \sum_{\mathbf{x}\in M_\mathbf{j}} N(\mathbf{x}, t) \cdot \mathbb{P}\left[W_\mathbf{x}(u) \in Q_r(\mathbf{y})\right] \\
&\geq \sum_{\mathbf{j}\in J} \sum_{\mathbf{x}\in M_\mathbf{j}} N(\mathbf{x}, t) \cdot \min_{\mathbf{x}'\in M_\mathbf{j}} \mathbb{P}\left[W_{\mathbf{x}'}(u) \in Q_r(\mathbf{y})\right] \\
&= \sum_{\mathbf{j}\in J} N(M_\mathbf{j}, t) \cdot \min_{\mathbf{x}'\in M_\mathbf{j}} \mathbb{P}\left[W_{\mathbf{x}'}(u) \in Q_r(\mathbf{y})\right] \\
&\geq \gamma_{r+1} \cdot C^{k(r+1)} \cdot \sum_{\mathbf{j}\in J} \min_{\mathbf{x}'\in M_\mathbf{j}} \mathbb{P}\left[W_{\mathbf{x}'}(u) \in Q_r(\mathbf{y})\right].
\end{aligned}$$

The probability that $W_{\mathbf{x}'}(u) \in Q_r(\mathbf{y})$ is minimised when $\mathbf{x}'$ is the farthest possible vertex from $Q_r(\mathbf{y})$ in $M_\mathbf{j}$. Thus, it will be minimised at one of the corners of $M_\mathbf{j}$, suppose $\mathbf{x}'_\mathbf{j}$. First, notice that $\mathbf{x}'_\mathbf{j}$ cannot share a coordinate with $\mathbf{y}$ because, otherwise, we could change that coordinate and get farther from $Q_r(\mathbf{y})$. Additionally, for $\mathbf{j_1} \neq \mathbf{j_2}$, $\mathbf{x}'_{\mathbf{j}_1} \neq \mathbf{x}'_{\mathbf{j}_2}$. Thus, the collection of $\mathbf{x}'_\mathbf{j}$ is precisely the set of corners $\mathbf{x_{j'}}$ which do not share a coordinate with $\mathbf{y}$, i.e., when $\mathbf{j}'$ does not have a 0 coordinate. So we define

$$J' = \{\mathbf{j}' \in J \mid \mathbf{j}' \text{ does not have coordinate } 0\},$$

and, continuing our bound,

$$\begin{aligned}
\mathbb{E}\left[N(Q_r(\mathbf{y}), t+u)\right] &\geq \gamma_{r+1} \cdot C^{k(r+1)} \cdot \sum_{\mathbf{j}\in J} \mathbb{P}\left[W_{\mathbf{x}'_\mathbf{j}}(u) \in Q_r(\mathbf{y})\right] \\
&= \gamma_{r+1} \cdot C^{k(r+1)} \cdot \sum_{\mathbf{j}'\in J'} \mathbb{P}\left[W_{\mathbf{x_{j'}}}(u) \in Q_r(\mathbf{y})\right] \\
&\geq \gamma_{r+1} \cdot C^{k(r+1)} \cdot \sum_{\mathbf{j}'\in J'} \mathbb{P}\left[W(u) + \mathbf{x_{j'}} \in Q_r(\mathbf{y})\right] \\
&= \gamma_{r+1} \cdot C^{k(r+1)} \cdot \left(\sum_{\mathbf{j}\in J} \mathbb{P}\left[W(u) + \mathbf{j} \cdot C^{r+1} \in Q_r(\mathbf{0})\right] - \right. \\
&\qquad \left. \sum_{\mathbf{j}\in J\setminus J'} \mathbb{P}\left[W(u) + \mathbf{j} \cdot C^{r+1} \in Q_r(\mathbf{0})\right] \right). \qquad (6)
\end{aligned}$$

Next, we bound the sums above separately. Let $W'(u)$ be a random walk on a k-dimensional torus H_k, with vertex set $V(H_k) = \{0, \ldots, C^{r+1}-1\}^k$, starting from vertex $\mathbf{0}$. Then,

$$\sum_{\mathbf{j}\in J} \mathbb{P}\left[W(u) + \mathbf{j} \cdot C^{r+1} \in Q_r(\mathbf{0})\right] = \mathbb{P}\left[W(u) + \mathbf{j} \cdot C^{r+1} \in Q_r(\mathbf{0}) \text{ for some } \mathbf{j} \in \mathbb{Z}^k\right]$$

$$= \mathbb{P}\left[W'(u) \in Q_r(\mathbf{0})\right],$$

where the first equality holds because in $u \leq \Delta_{r+1} - \Delta_r$ steps, the walk $W(u)$ can not reach a point $\mathbf{z} - \mathbf{j} \cdot C^{r+1}$ for $\mathbf{z} \in Q_r(\mathbf{0})$ but $j \notin J$. We have that $W'(0) \in Q_r(\mathbf{0})$, so as $u' \geq 0$ increases, $\mathbb{P}\left[W'(2u') \in Q_r(\mathbf{0})\right]$ decreases monotonically to its stationary value. Thus, $\mathbb{P}\left[W'(u) \in Q_r(\mathbf{0})\right] \geq |Q_r(\mathbf{0})/|V(H_k)| = C^{-k}$, which bounds the first sum in (6).

For the second sum in (6), we consider the cases when each component l is 0 separately. For $l \in \{1, \ldots, k\}$, let $J_l = \{(j_1, \ldots, j_k) \in J \mid j_l = 0\}$. For an integer $h \geq 0$, let $L_l^h = \{\mathbf{x} \in Q_r(\mathbf{0}) \mid \mathbf{x}_l = h\}$, which partition $Q_r(\mathbf{0})$ into disjoint sets of size $C^{(k-1)r}$. For $\mathbf{j} \in J_l$ the probability that a walk starting at $\mathbf{j} \cdot C^{r+1}$ is in L_l^h is greatest for $h = 0$. Therefore,

$$\begin{aligned}
\sum_{\mathbf{j} \in J \setminus J'} \mathbb{P}\left[W(u) + \mathbf{j} \cdot C^{r+1} \in Q_r(\mathbf{0})\right] &\leq \sum_{l=1}^{k} \sum_{\mathbf{j} \in J_l} \mathbb{P}\left[W(u) + \mathbf{j} \cdot C^{r+1} \in Q_r(\mathbf{0})\right] \\
&= \sum_{l=1}^{k} \sum_{\mathbf{j} \in J_l} \sum_{h=0}^{C^r - 1} \mathbb{P}\left[W(u) + \mathbf{j} \cdot C^{r+1} \in L_l^h\right] \\
&\leq C^r \cdot \sum_{l=1}^{k} \sum_{\mathbf{j} \in J_l} \mathbb{P}\left[W(u) + \mathbf{j} \cdot C^{r+1} \in L_l^0\right] \\
&= C^r \cdot \sum_{l=1}^{k} P_l, \quad \text{where } P_l \text{ is the internal sum above.}
\end{aligned}$$

Consider a fixed $l \in \{1, \ldots, k\}$. For any $\mathbf{j} \in J_l$ and $\mathbf{x} \in L_l^0$, the $\mathbf{x} - \mathbf{j} \cdot C^{r+1}$ are all unique vectors, and have 0 as their lth coordinate. Therefore, if $x(u)$ is $W(u)$'s lth coordinate, then $P_l \leq \mathbb{P}\left[x(u) = 0\right]$. Notice, however, that $x(u)$ is a lazy random walk on $\mathbb{Z}$, starting at 0, with holding probability $1 - 1/k$. Therefore, by Lemma 41 there is a constant η, such that

$$P_l \leq \mathbb{P}\left[x(u) = 0\right] \leq \frac{\eta}{\sqrt{u}} \leq \frac{\eta}{\sqrt{\Delta_{r+1} - \Delta_r}} \leq \frac{\eta}{\sqrt{\Delta_{r+1}/2}} \leq \eta \cdot C^{-2k(r+1)}.$$

We substitute these bounds in (6):

$$\begin{aligned}
\mathbb{E}\left[N(Q_r(\mathbf{y}), t + u)\right] &\geq \gamma_{r+1} \cdot C^{k(r+1)} \cdot \left(C^{-k} - \sqrt{2}\eta k \cdot C^r \cdot C^{-2k(r+1)}\right) \\
&\geq \gamma_{r+1} \cdot C^{kr}\left(1 - \sqrt{2}\eta k \cdot C^{-(r+1)}\right) \geq \gamma_{r+1} \cdot C^{kr}\left(1 - C^{-(r+1)/2}\right),
\end{aligned}$$

for $C \geq C_1 = \sqrt{2}\eta k$. ◀

▶ **Lemma 29.** *There is a constant $C_2 > 0$, such that if $C \geq C_2$, then given $\mathcal{K}_{r+1}(s-1)$ and the event that $\tilde{S}_{r+1}(\mathbf{i}, s)$ is good, the probability that all the children of $\tilde{S}_{r+1}(\mathbf{i}, s)$ are good is at least $1 - \rho_{r+1}$, where*

$$\rho_{r+1} = C^{8k^2(r+1)} \cdot \exp\left(-\frac{\gamma_{r+1}}{4} \cdot C^{k(r+1)/4}\right).$$

Proof. For convenience denote $t = (s-1) \cdot \Delta_{r+1}$. Suppose the child $\tilde{S}_r(\mathbf{j}, m)$ of $\tilde{S}_{r+1}(\mathbf{i}, s)$ is bad for some $\mathbf{j}$ and m. Then there is a set $Q_r(\mathbf{y}) \subset B_r(\mathbf{j})$ such that $N(\tilde{Q}_r(\mathbf{y}, m-1)) < \gamma_r \cdot C^{kr}$. We fix such $\mathbf{y}$ and m, then bound the probability of $N(\tilde{Q}_r(\mathbf{y}, m-1)) < \gamma_r \cdot C^{kr}$, and take a union bound over all such pairs.

For an agent g in $\tilde{B}_{r+1}(\mathbf{i}, s-1)$, let X_g be the indicator random variable that g is at $\tilde{Q}_r(\mathbf{y}, m-1)$. Note that g has to travel for $u = (m-1) \cdot \Delta_r - t$ rounds to be at $\tilde{Q}_r(\mathbf{y}, m-1)$.

By definition of being a child, $u \in [\Delta_{r+1} - \Delta_r, 2\Delta_{r+1})$. Then

$$N(\tilde{Q}_r(\mathbf{y}, m-1)) = N(Q_r(\mathbf{y}), t+u) = \sum_{g \text{ at } \tilde{B}_{r+1}(\mathbf{i}, s-1)} X_g.$$

We take $C_2 = \max\{6, C_1\}$, where C_1 is determined in Lemma 28. Since the agents move independently after round t, we can apply Chernoff bound:

$$\begin{aligned}
\mathbb{P}\left[N(Q_r(\mathbf{y}), t+u) < \gamma_r \cdot C^{kr}\right] &= \mathbb{P}\left[N(Q_r(\mathbf{y}), t+u) < \gamma_{r+1} \cdot C^{kr} \cdot (1 - C^{-(r+1)/8})\right] \\
&\le \mathbb{P}\left[N(Q_r(\mathbf{y}), t+u) < \mathbb{E}\left[N(Q_r(\mathbf{y}), t+u)\right] \cdot \frac{1 - C^{-(r+1)/8}}{1 - C^{-(r+1)/2}}\right] \\
&= \mathbb{P}\left[N(Q_r(\mathbf{y}), t+u) < \mathbb{E}\left[N(Q_r(\mathbf{y}), t+u)\right] \cdot \left(1 - \frac{C^{-(r+1)/8} - C^{-(r+1)/2}}{1 - C^{-(r+1)/2}}\right)\right] \\
&\le \exp\left(-\frac{1}{2} \cdot \left(\frac{C^{-(r+1)/8} - C^{-(r+1)/2}}{1 - C^{-(r+1)/2}}\right)^2 \cdot \mathbb{E}\left[N(Q_r(\mathbf{y}), t+u)\right]\right) \\
&\le \exp\left(-\frac{1}{2} \cdot \left(\frac{C^{-(r+1)/8} - C^{-(r+1)/2}}{1 - C^{-(r+1)/2}}\right)^2 \cdot \gamma_{r+1} \cdot C^{kr} \cdot \left(1 - C^{-(r+1)/2}\right)\right) \\
&= \exp\left(-\frac{\gamma_{r+1} \cdot C^{kr}}{2} \cdot C^{-(r+1)/4} \cdot \frac{\left(1 - C^{-3(r+1)/8}\right)^2}{1 - C^{-(r+1)/2}}\right) \\
&\le \exp\left(-\frac{\gamma_{r+1}}{4} \cdot C^{k(r+1)/4}\right),
\end{aligned}$$

where it is easy to verify that the last inequality holds since $C \ge 6$. The number of pairs $(\mathbf{y}, m)$, which may contain less than $\gamma_r \cdot C^{kr}$ agents and render the child $\tilde{S}_r(\mathbf{j}, m)$ of $\tilde{S}_{r+1}(\mathbf{i}, s)$ bad, is at most $(3 \cdot \Delta_{r+1})^k \cdot (\Delta_{r+1}/\Delta_r) = 3^k \cdot C^{4k^2(r+1)+4k} \le C^{8k^2(r+1)}$. Thus, the proof is complete after an application of a union bound. ◀

9.3 Information Paths

An *information path* $\overline{\mathbf{x}}$ is defined as a sequence $\mathbf{x}_t$ of vertices in V', such that for any $t \ge 0$, either $\mathbf{x}_{t+1} = \mathbf{x}_t$ or $\operatorname{dist}(\mathbf{x}_{t+1}, \mathbf{x}_t) = 1$. Let Θ be the set of all information paths of length exactly $T^* = c\ell$, for a constant $c > 0$. We say that an information path $\overline{\mathbf{x}} \in \Theta$ intersects set $\tilde{S}_r(\mathbf{i}, s)$ if $\mathbf{x}_{s \cdot \Delta_r} \in S_r(\mathbf{i})$. Let $\phi_r(\overline{\mathbf{x}})$ be the number of bad sets $\tilde{S}_r(\mathbf{i}, s)$ that $\overline{\mathbf{x}}$ intersects, and

$$\Phi_r = \max_{\overline{\mathbf{x}} \in \Theta} \phi_r(\overline{\mathbf{x}}).$$

For $r > 1$, we also define $\psi_r(\overline{\mathbf{x}})$ as the number of good sets $\tilde{S}_r(\mathbf{i}, s)$ that have a bad child and intersect $\overline{\mathbf{x}}$. We define

$$\Psi_r = \max_{\overline{\mathbf{x}} \in \Theta} \psi_r(\overline{\mathbf{x}}).$$

In this section, we prove an upper bound on the maximum number of bad sets at scale r that intersect an information path in Θ. In particular, our final lemma, bounding Φ_1, argues that at least a constant fraction of sets $\tilde{S}_1(\mathbf{i}, s)$ that intersect any information path are good. This allows us to argue, roughly, that if we split time into phases of $\Delta_1 = O(1)$, then in a constant fraction of those phases, progress is made toward a target vertex with constant probability.

▶ **Lemma 30.** *For any scale* $r \in \{1, \ldots, R-1\}$,

$$\Phi_r \leq (\Phi_{r+1} + \Psi_{r+1}) \cdot (\Delta_{r+1}/\Delta_r) = (\Phi_{r+1} + \Psi_{r+1}) \cdot C^{4k}.$$

Proof. If $\overline{\mathbf{x}} \in \Theta$ intersects the set $\tilde{S}_{r+1}(\mathbf{i}, s)$, then it can also intersect at most Δ_{r+1}/Δ_r of its children, since it can only intersect a child at rounds $s \cdot \Delta_{r+1} + i \cdot \Delta_r$ for an integer $0 \leq i < \Delta_{r+1}/\Delta_r = C^{4k}$. If $\tilde{S}_{r+1}(\mathbf{i}, s)$ is either bad, or has a bad child, we assume that all its children that $\overline{\mathbf{x}}$ intersects are bad. This gives us an upper bound:

$$\phi_r(\overline{\mathbf{x}}) \leq \phi_{r+1}(\overline{\mathbf{x}}) \cdot C^{4k} + \psi_{r+1}(\overline{\mathbf{x}}) \cdot C^{4k}.$$

The proof is completed by taking a maximum on both sides of the inequality with respect to all paths $\overline{\mathbf{x}} \in \Theta$. ◀

▶ **Lemma 31.** *For any constants* $c, C > 0$ *and* $\kappa > 0$*, there is a value* $R = \Theta(\log \log n)$ *such that* $\mathbb{P}[\Phi_R = 0] \geq 1 - n^{-\kappa}$.

Proof. Let $R = \lceil \log_C(\eta \ln n)/k \rceil$ for a constant $\eta > 0$. Consider some space-time set $\tilde{Q}_R(\mathbf{x}, s) \subset V \times [0, T^*)$. By definition $|Q_R(\mathbf{x})| = C^{kR} \geq \eta \cdot \ln n$. By a Chernoff bound,

$$\mathbb{P}\left[N(\tilde{Q}_R(\mathbf{x}, s) \geq \frac{\alpha}{2} \cdot |Q_R(\mathbf{x})|\right] \geq 1 - e^{-\frac{\alpha}{8} \cdot |Q_R(\mathbf{x})|} \geq 1 - n^{-\frac{\alpha \cdot \eta}{8}}.$$

If $\Phi_R > 0$, then $N(\tilde{Q}_R(\mathbf{x}, s)) < \gamma_R \cdot C^{kR} = \alpha \cdot C^{kR}/2$ for some vertex $\mathbf{x}$ and integer s. Since the number of such sets $\tilde{Q}_R(\mathbf{x}, s)$ is at most $(T^*/\Delta_R) \cdot \ell^k < c \cdot \ell^{k+1}$, by a union bound,

$$\mathbb{P}[\Phi_R = 0] \geq 1 - c \cdot \ell^{k+1} \cdot n^{-\frac{\alpha \cdot \eta}{8}} \geq 1 - n^{-\kappa},$$

for a constant η large enough. ◀

▶ **Lemma 32.** *There is a constant* $C_3 > 0$*, such that for any* $C \geq C_3$*, if* $r \geq 2$*, then*

$$\mathbb{P}\left[\Psi_r \geq e^{-r} \cdot \frac{T^*}{\Delta_r}\right] \leq e^{-T^*/\Delta_r}.$$

Proof. Let P_r denote the set of pairs $(\mathbf{i}, s)$ that satisfy $\tilde{S}_r(\mathbf{i}, s) \subset V' \times [\Delta_R, T^* - \Delta_R]$. We partition P_r into $m = 2 \cdot 7^k$ disjoint sets, which are defined using an integer $v \in \{0, 1\}$ and a vector $\mathbf{h} \in H = \{0, \ldots, 6\}^k$, as follows:

$$P_r(\mathbf{h}, v) = \{(\mathbf{i}, s) \in P_r \mid \mathbf{i} \equiv \mathbf{h} \pmod 7,\ s \equiv v \pmod 2\}.$$

Let $\psi_{r,\mathbf{h},v}(\overline{\mathbf{x}})$ be the number of good segments $\tilde{S}_r(\mathbf{i}, s)$, for $(\mathbf{i}, s) \in P_r(\mathbf{h}, v)$, that have a bad child and are intersected by $\overline{\mathbf{x}} \in \Theta$. Let also

$$\Psi_r(\mathbf{h}, v) = \max_{\overline{\mathbf{x}} \in \Theta} \psi_{r,\mathbf{h},v}(\overline{\mathbf{x}}).$$

To prove the lemma we first bound $\psi_{r,\mathbf{h},v}(\overline{\mathbf{x}})$ for a fixed pair $\mathbf{h}, v$ and path $\overline{\mathbf{x}}$. We then use union bound twice: first to bound $\Psi_r(\mathbf{h}, v)$, and then to bound Ψ_r, as the latter is the sum of all $\Psi_r(\mathbf{h}, v)$.

Consider a fixed pair $\mathbf{h}, v$. For $(\mathbf{i}, s) \in P_r(\mathbf{h}, v)$, define $Y(\mathbf{i}, s)$ as the indicator random variable that is 1 if $\tilde{S}_r(\mathbf{i}, s)$ is good, but has a bad child. From Lemma 29,

$$\mathbb{P}[Y(\mathbf{i}, s) = 1 \mid \mathcal{K}_r(\mathbf{i}, s-1)] = \mathbb{P}\left[\tilde{S}_r(\mathbf{i}, s) \text{ has a bad child} \mid \mathcal{K}_r(\mathbf{i}, s-1); \tilde{S}_r(\mathbf{i}, s) \text{ is good}\right] \cdot \mathbb{P}\left[\tilde{S}_r(\mathbf{i}, s) \text{ is good} \mid \mathcal{K}_r(\mathbf{i}, s-1)\right]$$

$$\leq \rho_r.$$

Consider the following ordering of elements of $P_r(\mathbf{h}, v)$: $(\mathbf{i}', s') \prec (\mathbf{i}, s)$ if $s' < s$, or $s' = s$ and $\mathbf{i}'$ is lexicographically smaller than $\mathbf{i}$ (this decision is arbitrary). Due to the space and time separation of the sets $\tilde{S}_r(\mathbf{i}, s)$ for $(\mathbf{i}, s) \in P_r(\mathbf{h}, v)$,

$$\begin{aligned}&\mathbb{P}\left[Y(\mathbf{i}, s) = 1 \mid \mathcal{K}_r(\mathbf{i}, s-1); Y(\mathbf{i}', s') \text{ for all } (\mathbf{i}', s') \prec (\mathbf{i}, s)\right] \\ &\qquad\qquad = \mathbb{P}\left[Y(\mathbf{i}, s) = 1 \mid \mathcal{K}_r(\mathbf{i}, s-1)\right] \leq \rho_r.\end{aligned}$$

For pairs $(\mathbf{i}, s) \in P_r(\mathbf{h}, v)$, let $\{Z(\mathbf{i}, s)\}$ be a collection of independent Bernoulli random variables with success probability ρ_r. From the above it follows that

$$\mathbb{P}\left[Y(\mathbf{i}, s) = 1 \mid \mathcal{K}_r(\mathbf{i}, s-1);\ Y(\mathbf{i}', s') \text{ for all } (\mathbf{i}', s') \prec (\mathbf{i}, s)\right] \leq \mathbb{P}\left[Z(\mathbf{i}, s) = 1\right]. \tag{7}$$

Notice that to bound $\Psi_r(\mathbf{h}, v)$ it is wasteful to take a union bound over all paths in Θ, because many information paths intersect exactly the same collection of sets $\tilde{S}_r(\mathbf{i}, s)$, for $(\mathbf{i}, s) \in P_r(\mathbf{h}, v)$. Thus, we can group them into such equivalence classes, reducing the number of objects we need to take a union bound over. For an information path $\overline{\mathbf{x}}$, define

$$I_{r,\mathbf{h},v}(\overline{\mathbf{x}}) = \{(\mathbf{i}, s) \in P_r(\mathbf{h}, v) \mid \overline{\mathbf{x}} \text{ intersects } \tilde{S}_r(\mathbf{i}, s)\}.$$

Then,

$$\psi_{r,\mathbf{h},v}(\overline{\mathbf{x}}) = \sum_{(\mathbf{i},s)\in I_{r,\mathbf{h},v}(\overline{\mathbf{x}})} Y(\mathbf{i}, s) \leq |I_{r,\mathbf{h},v}|.$$

Next, we bound the probability that $\psi_{r,\mathbf{h},v}(\overline{\mathbf{x}}) \geq \frac{e^{-r}}{m} \cdot \frac{T^*}{\Delta_r}$, where $m = 2 \cdot 7^k$ (recall, m is the number of sets $P_r(\mathbf{h}, v)$ to which P_r is partitioned). Let $Z = \sum_{(\mathbf{i},s)\in I_{r,\mathbf{h},v}(\overline{\mathbf{x}})} Z(\mathbf{i}, s)$, and $b = \frac{2e^{-r}}{m \cdot \rho_r}$. Then,

$$b \cdot \mathbb{E}\left[Z\right] = b \cdot |I_{r,\mathbf{h},v}(\overline{\mathbf{x}})| \cdot \rho_r \leq b \cdot \frac{T^*}{2\Delta_r} \cdot \rho_r = \frac{e^{-r}}{m} \cdot \frac{T^*}{\Delta_r}. \tag{8}$$

Let C_2 be as in Lemma 29, and let C_3' be the smallest constant such that for any $C \geq C_3'$,

$$\frac{\alpha}{32} \cdot C^{kr/4} \geq r(1 + 8k^2 \ln C), \quad \text{and} \quad C^{kr/8} \cdot \frac{e^{-r}}{m} \cdot \frac{\alpha}{64} \geq 1, \text{ for any } r \geq 1.$$

For $C \geq \max\{256, C_2, C_3'\}$, by substituting for ρ_r and using $\gamma_r \geq \alpha/4$,

$$\ln b \geq \ln \frac{2}{m} - r(1 + 8k^2 \ln C) + \frac{\gamma_r}{4} \cdot C^{kr/4} \geq \frac{\alpha}{32} \cdot C^{kr/4}. \tag{9}$$

Furthermore, (7) implies that we can couple the sets of variables $\{Y(\mathbf{i}, s)\}$ and $\{Z(\mathbf{i}, s)\}$, for $(\mathbf{i}, s) \in I_{r,\mathbf{h},v}(\overline{\mathbf{x}})$, such that $Y(\mathbf{i}, s) \leq Z(\mathbf{i}, s)$, and thus, $\psi_{r,\mathbf{h},v}(\overline{\mathbf{x}})$ is stochastically dominated by Z. Then,

$$\begin{aligned}\mathbb{P}\left[\psi_{r,\mathbf{h},v}(\overline{\mathbf{x}}) \geq \frac{e^{-r}}{m} \cdot \frac{T^*}{\Delta_r}\right] &\leq \mathbb{P}\left[Z \geq \frac{e^{-r}}{m} \cdot \frac{T^*}{\Delta_r}\right], \quad \text{by the coupling,} \\ &\leq \left(\frac{b}{e}\right)^{-\frac{e^{-r}}{m} \cdot \frac{T^*}{\Delta_r}}, \quad \text{by Lemma 42 and (8),} \\ &\leq \exp\left(-\frac{T^*}{\Delta_r} \cdot \frac{e^{-r}}{m} \cdot \left(\frac{\alpha}{32} \cdot C^{kr/4} - 1\right)\right), \quad \text{by (9),}\end{aligned}$$

$$\leq \exp\left(-\frac{T^*}{\Delta_r}\cdot\frac{e^{-r}}{m}\cdot\frac{\alpha}{64}\cdot C^{kr/4}\right)$$
$$\leq \exp\left(-\frac{T^*}{\Delta_r}\cdot C^{kr/8}\right),$$

where the last two inequalities hold since $C \geq C_3'$.

Next we upper bound the number of distinct values that the set $I_{r,\mathbf{h},v}(\overline{\mathbf{x}})$ takes, for all $\overline{\mathbf{x}} \in \Theta$, i.e., the cardinality of set $\bigcup_{\overline{\mathbf{x}}\in\Theta}\{I_{r,\mathbf{h},v}(\overline{\mathbf{x}})\}$. It suffices to bound instead the cardinality of $I_r = \bigcup_{\overline{\mathbf{x}}\in\Theta}\{I_r(\overline{\mathbf{x}})\}$, where $I_r(\overline{\mathbf{x}}) = \bigcup_{\mathbf{h}\in H, v\in\{0,1\}} I_{r,\mathbf{h},v}(\overline{\mathbf{x}})$ is the set of all $\tilde{S}_r(\mathbf{i}, s)$ that $\overline{\mathbf{x}}$ intersects. We do that by looking at how many possible sets $\tilde{S}_r(\mathbf{i}, s)$ can be in $I_r(\overline{\mathbf{x}})$ for each $s \in \{1, \ldots, \lfloor T^*/\Delta_r\rfloor\}$, given the previous elements in $I_r(\overline{\mathbf{x}})$: For $s = 1$, there are at most $(\ell/\Delta_r)^k$ possible choices. For $s \geq 2$, if $\overline{\mathbf{x}}$ intersects both $\tilde{S}_r(\mathbf{i}, s-1)$ and $\tilde{S}_r(\mathbf{j}, s)$, then $\mathbf{i}$ and $\mathbf{j}$ differ by at most 1 in each coordinate. Therefore, given the first elements in $I_r(\overline{\mathbf{x}})$ up to s, there are at most 3^k possible choices of the next element. Therefore,

$$|I_r| \leq (\ell/\Delta_r)^k \cdot (3^k)^{T^*/\Delta_r} \leq \exp\left(k\ln\ell + k\ln 3\cdot\frac{T^*}{\Delta_r}\right) \leq \exp\left(2k\cdot\frac{T^*}{\Delta_r}\right).$$

Using a union bound we get that,

$$\mathbb{P}\left[\Psi_r(\mathbf{h}, v) \geq \frac{e^{-r}}{m}\cdot\frac{T^*}{\Delta_r}\right] \leq \exp\left(2k\cdot\frac{T^*}{\Delta_r}\right)\cdot\exp\left(-\frac{T^*}{\Delta_r}\cdot C^{kr/8}\right) \leq e^{-T^*/\Delta_r}/m,$$

for $C \geq C_3'' = (2k + \ln m + 1)^8$. By another union bound over all m values of pair $\mathbf{h}, v$, we prove the desired result for $C_3 = \max\{C_2, C_3', C_3''\}$. ◀

▶ **Lemma 33.** *For any constant $\kappa > 0$, there are constants c and C, such that*

$$\mathbb{P}\left[\Phi_1 \leq \frac{T^*}{4\Delta_1}\right] \geq 1 - 2n^{-\kappa}.$$

Proof. Let C be the smallest even integer that is at least C_3 (defined in Lemma 32). Let $\mathcal{E}_1$ be the event that $\Psi_{r+1} < e^{-(r+1)}\cdot\frac{T^*}{\Delta_{r+1}}$ for all $r \in \{1, \ldots, R-1\}$ at the same time, and $\mathcal{E}_2$ be the event that $\Phi_R = 0$. If $\mathcal{E}_1$ and $\mathcal{E}_2$ hold, then, by a recursive application of Lemma 30,

$$\Phi_1 \leq \Phi_R\cdot C^{4k(R-1)} + \sum_{r=1}^{R-1}\Psi_{r+1}\cdot C^{4kr}$$
$$\leq T^*\cdot\sum_{r=1}^{R-1}\frac{e^{-(r+1)}}{\Delta_{r+1}}\cdot C^{4kr}$$
$$= \frac{T^*}{C^{4k}}\cdot\sum_{r=1}^{R-1}e^{-(r+1)}$$
$$\leq \frac{T^*}{e^2(1-1/e)\cdot\Delta_1}$$
$$\leq \frac{T^*}{4\Delta_1}.$$

By a union bound and Lemma 32,

$$\mathbb{P}[\mathcal{E}_1] \geq 1 - \sum_{r=1}^{R-1}\mathbb{P}\left[\Psi_{r+1} \geq e^{-(r+1)}\cdot\frac{T^*}{\Delta_{r+1}}\right] \geq 1 - \sum_{r=1}^{R-1}e^{-T^*/\Delta_{r+1}} \geq 1 - R\cdot e^{-T^*/\Delta_R} \geq 1 - n^{-\kappa},$$

where the last inequality holds for $c > 0$ large enough and the corresponding value of R as determined by Lemma 31. Since $\mathbb{P}[\mathcal{E}_2] \geq 1 - n^{-\kappa}$, we have that $\mathbb{P}\left[\Phi_1 \leq \frac{T^*}{4\Delta_1}\right] \geq 1 - 2n^{-\kappa}$. ◀

9.4 Putting the Pieces Together

Recall that $V' = \prod_{j=1}^{s}\{6\Delta_R, \dots, \ell - 6\Delta_R\}$ is the "central" part of V. We first prove, using the bound on Φ_1, that if a vertex in V' is informed, then in at most $O(\ell)$ rounds, any fixed vertex of V' becomes informed, w.h.p. Then, using standard random walk techniques, we prove that if the source vertex is an "edge" vertex in $V \setminus V'$, then some vertex in V' becomes informed in $O(\ell)$ rounds. We combine these facts to prove Theorem 6.

▶ **Lemma 34.** *Let $\mathbf{x}, \mathbf{y} \in V'$. For any constant $\kappa > 0$, there is a large enough constant $c > 0$ such that if $\mathbf{x}$ is informed, then after at most $T^* = c\ell$ rounds, $\mathbf{y}$ also becomes informed, with probability at least $1 - 3n^{-\kappa}$.*

Proof. Fix any shortest path $\mathbf{x}_0, \mathbf{x}_1, \dots, \mathbf{x}_\lambda$ between $\mathbf{x}$ and $\mathbf{y}$, where $\lambda = \operatorname{dist}(\mathbf{x}, \mathbf{y})$. We consider a process $\hat{\mathbf{x}}(t)$, $t \geq 0$, on the vertices of that path, such that $\hat{\mathbf{x}}(0) = \mathbf{x}_0 = \mathbf{x}$, and for $t \geq 1$, $\hat{\mathbf{x}}(t)$ is defined as follows: if $\hat{\mathbf{x}}(t-1) = \mathbf{x}_i$ for some $i < \lambda$, and some agent moves from $\mathbf{x}_i$ to $\mathbf{x}_{i+1}$ in round t, then $\hat{\mathbf{x}}(t) = \mathbf{x}_{i+1}$; otherwise, $\hat{\mathbf{x}}(t) = \hat{\mathbf{x}}(t-1)$. At round t, $\hat{\mathbf{x}}(t)$ is the closest informed vertex to $\mathbf{y}$ along the designated path from $\mathbf{x}$ to $\mathbf{y}$. Let $\tau = \min\{t \mid \hat{\mathbf{x}}(t) = \mathbf{y}\}$. Since τ is an upper bound on the number of rounds until $\mathbf{y}$ is informed, it suffices to show that $\mathbb{P}[\tau \leq T^*] \geq 1 - 3n^{-\kappa}$.

Since $\mathbf{x}, \mathbf{y} \in V'$, clearly $\mathbf{x}_i \in V'$, for all $0 \leq i \leq \lambda$, and thus $\hat{\mathbf{x}}(t) \in V'$ for any $t \geq 0$. Then, for any integer $s \geq 0$, and for $t_s = s\Delta_1$, there is some index $\mathbf{i}_s$ such that $\hat{\mathbf{x}}(t_s) \in S_1(\mathbf{i}_s)$.

Suppose that the space-time set $\tilde{S}_1(\mathbf{i}_s, s+1)$ is good. Then by definition of goodness, every set $Q_1(\mathbf{z})$ that is a subset of $B_1(\mathbf{i}_s)$ contains at least $\gamma_1 \cdot C^k$ agents at round $t_s = s\Delta_1$. Also, by construction $\hat{\mathbf{x}}(t_s) \in B_1(\mathbf{i}_s)$, and therefore, it is also the case that $\hat{\mathbf{x}}(t_s) \in Q_1(\mathbf{z}) \subset B_1(\mathbf{i}_s)$, for some vertex $\mathbf{z}$. Since $Q_1(\mathbf{z})$ is a k-dimensional cube of side C, it follows that there some agent g at a vertex $\mathbf{w}$ with $\operatorname{dist}(\hat{\mathbf{x}}(t_s), \mathbf{w}) < kC$, at time t_s. Suppose that $\hat{\mathbf{x}}(t_s) = \mathbf{x}_i$, for some $i < \lambda$. Then, with probability at least $\epsilon = (2k)^{-kC}$, agent g visits vertex $\mathbf{x}_i$ at some round $t \in \{t_s, \dots, t_s + kC\}$, followed by a visit to $\mathbf{x}_{i+1}$ at round $t + 1$ (note that, this probability bounds is very crude, but we only need a constant ϵ).

Let $m = \lfloor T^*/\Delta_1 \rfloor$. For $s \in \{0, \dots, m-1\}$, let Z_s be the indicator random variable of the event that the space-time set $\tilde{S}_1(\mathbf{i}_s, s+1)$ is good, and let $Z = \sum_{s=0}^{m-1} Z_s$. Then $Z \geq m - \Phi_1$, and choosing C and c as in Lemma 33, we have that

$$\mathbb{P}\left[Z \geq \frac{3}{4} \cdot m\right] \geq 1 - 2n^{-\kappa}. \tag{10}$$

For $s \in \{0, \dots, m-1\}$, let Y_s be the indicator random variable of the event that $\hat{\mathbf{x}}(t_s) = \mathbf{y}$ or $\hat{\mathbf{x}}(t_{s+1}) \neq \hat{\mathbf{x}}(t_s)$, i.e., the process makes progress towards $\mathbf{y}$ between rounds t_s and t_{s+1}, if it has not already reached $\mathbf{y}$ at time t_s. Recall that $\mathcal{K}_1(s)$ is the σ-algebra generated by the positions of all agents up to round t_s. Then, as we argued above,

$$\mathbb{P}[Y_s = 1 \mid \mathcal{K}_1(s); Z_s = 1] \geq \epsilon. \tag{11}$$

Note also that if $Y = \sum_{s=0}^{m-1} Y_s$, then $Y \geq k \cdot \ell$ implies $\tau \leq T^*$. Thus it suffices to show $\mathbb{P}[Y \geq k \cdot \ell] \geq 1 - n^{-\kappa}$.

Let $p_1, \dots, p_\mu$ denote the sequence of all $s \in \{0, \dots, m-1\}$, for which $Z_s = 1$. Define $X_j = Y_{p_j}$ for $1 \leq j \leq \mu$, and $X_j = 1$ for $j > \mu$. It follows from (11) that, for any $j \geq 1$,

$$\mathbb{P}[X_j = 1 \mid X_1, \dots, X_{j-1}] \geq \epsilon.$$

Then for $X = \sum_{s=1}^{3m/4} X_s$, we can apply a standard Chernoff bound[5] to obtain

$$\mathbb{P}\left[X \geq \frac{3\epsilon m}{8}\right] \geq 1 - e^{-3\epsilon m/32}.$$

Choosing c large enough that $\frac{3\epsilon m}{8} \geq k\ell$ and $e^{-3\epsilon m/32} \leq n^{-\kappa}$, we obtain

$$\mathbb{P}\left[X \geq k\ell\right] = 1 - n^{-\kappa}. \tag{12}$$

Note now that if $Z \geq \frac{3}{4} \cdot m$ then $Y \geq X$. It follows then from (10), (12), and union bound that $\mathbb{P}\left[Y \geq k\ell\right] \geq \mathbb{P}\left[\{Z \geq 3m/4\} \cap \{X \geq k\ell\}\right] = 1 - 3n^{-\kappa}$. ◀

▶ **Lemma 35.** *For any constant $\kappa > 0$ and any vertex $\mathbf{x} \in V \setminus V'$, if $\mathbf{x}$ is informed, then after at most $O(\text{poly}(\log \ell))$ rounds, some vertex in V' will become informed, with probability at lest $1 - n^{-\kappa}$.*

Proof. We prove that in $O(\text{poly}(\log \ell))$ rounds some agent will visit $\mathbf{x}$, and after that also visit a vertex in V'. Let $X(t)$ be a random walk starting from the stationary distribution of G and let $\tau_{\mathbf{x}}$ be the first round when it visits $\mathbf{x}$. If N is the number of times the walk $X(t)$ visits $\mathbf{x}$ in the first $\tau_1 = (\eta_1 \log \ell)^2$ rounds, where $\eta_1 > 0$ is a constant, then

$$\mathbb{E}\left[N\right] \geq \frac{\tau_1}{2n},$$

since the walk starts from stationarity and the vertex degrees are between k and $2k$. Also,

$$\begin{aligned}
\mathbb{E}\left[N \mid \tau_{\mathbf{x}}\right] &\leq \sum_{t=0}^{\tau_1 - \tau_{\mathbf{x}}} \mathbb{P}\left[X(\tau_{\mathbf{x}} + t) = \mathbf{x} \mid \tau_{\mathbf{x}}\right] \\
&= \sum_{t'=0}^{\lfloor(\tau_1 - \tau_x)/2\rfloor} \mathbb{P}\left[X(\tau_x + 2t') = \mathbf{x} \mid \tau_{\mathbf{x}}\right], \quad \text{since } G \text{ is bipartite,} \\
&\leq \sum_{t'=0}^{\lfloor(\tau_1 - \tau_x)/2\rfloor} \left(\frac{2k}{k \cdot n} + \frac{20 \cdot 2k}{k \cdot \sqrt{2t'+1}}\right), \quad \text{by Lemma 9 and } k \leq \deg(\mathbf{x}) \leq 2k, \\
&\leq \frac{2(\tau_1 + 1)}{n} + 40 \cdot \sum_{t=0}^{\tau_1} \frac{1}{\sqrt{t+1}} \\
&\leq \frac{2(\tau_1 + 1)}{n} + 80 \cdot \sqrt{\tau_1 + 1} \\
&\leq 100 \cdot \sqrt{\tau_1}, \quad \text{assuming a large enough constant } \eta_1.
\end{aligned}$$

Thus,

$$\mathbb{P}\left[N \geq 1\right] = \frac{\mathbb{E}\left[N\right]}{\mathbb{E}\left[N \mid \tau_x \leq \tau_1\right]} \geq \frac{\sqrt{\tau_1}}{200n}.$$

Let $\mathcal{E}$ be the event that at least one agent visits $\mathbf{x}$ in the first τ_1 rounds. By the independence of the agents' walks, and by taking a sufficiently large η_1, we have

$$\mathbb{P}\left[\mathcal{E}\right] \geq 1 - \left(1 - \frac{\sqrt{\tau_1}}{200n}\right)^{\alpha \cdot n} \geq 1 - e^{-\frac{\alpha}{200} \cdot \sqrt{\tau_1}} \geq 1 - n^{-\kappa}/2. \tag{13}$$

[5] See, e.g., [5, Lemma 3.1], which allows us to apply a Chernoff bound in our setting of dependent random variables.

Next, our goal is to bound the number of rounds until $X(t)$ reaches one of the vertices in V' after round $\tau_{\mathbf{x}}$. Consider a random walk $X'(t)$ on a k-dimensional *torus* with ℓ vertices on each side, for which too we use V' to denote the set of its "central" vertices. Since V' is equidistant from the sides of the grid G, $X(t)$ and $X'(t)$ can be coupled (in a natural way), so that they both hit a vertex in V' at the same round. Let $h = 12 \cdot \Delta_R$ and $\tau_2 = \eta_2 h^{k+1} \log \ell$ for some constant $\eta_2 > 0$. By symmetry of V', we can assume that one of the coordinates of $\mathbf{x}$ is less than $h/2$, w.l.o.g., $0 \le x_1 < h/2$. Consider a walk $Y(t)$ on a smaller k-dimensional torus H with h vertices on each side, where each coordinate of $Y(t)$ is the same as that coordinate of $X'(t)$, modulo h. By [29, Sec. 11.3.2], the expected cover time of Y is at most $O(kh^k \log h) = O(h^{k+1})$, therefore, for a sufficiently large η_2, the walk Y visits all vertices of H in τ_2 rounds, with probability at least $1 - n^{-\kappa}/2$. Let $\mathbf{z}$ be the vertex of H with all its coordinates equal to $h/2 = 6\Delta_R$. If $Y(t) = \mathbf{z}$ and $t < l - h$, then $X'(t) \in V'$, because no vertex in $V \setminus V'$ has its first coordinate equal to $h/2$ (modulo h) and is less than $l - h$ steps away from $\mathbf{x}$. It implies that in τ_2 rounds $X'(t)$ (and thus, $X(t)$) visits some vertex in V', with probability at least $1 - n^{-\kappa}/2$. Combining this with (13), we prove that in $\tau = \tau_1 + \tau_2$ rounds some vertex in V' becomes informed, with probability at least $1 - n^{-\kappa}$. Substituting the values of h and R, we see that $\tau = O(\text{poly}(\log \ell))$, which completes the proof. ◀

Proof of Theorem 6. Fix any vertices $\mathbf{x} \in V'$ and $\mathbf{y} \in V$, and any constant $\kappa > 1$. First suppose that $\mathbf{y}$ is informed initially. If $\mathbf{y} \in V \setminus V'$, then by Lemma 35, some vertex $\mathbf{z} \in V'$ becomes informed in at most $O(\text{poly}(\log \ell))$ rounds with probability at least $1 - n^{-\kappa}$. If $\mathbf{y} \in V'$, we just let $\mathbf{z} = \mathbf{y}$. Then, by Lemma 34 we conclude that $\mathbf{x}$ becomes informed (via $\mathbf{z}$) in at most $O(\ell)$ rounds with probability at least $1 - 3n^{-\kappa}$. Suppose now that $\mathbf{x}$ is informed initially, instead of $\mathbf{y}$. By Lemma 39, then $\mathbf{y}$ becomes informed in $O(\ell)$ rounds, with probability at least $1 - 3n^{-\kappa}$. Using a union bound over all vertices $\mathbf{y}$, we conclude that if $\mathbf{x} \in V'$ is informed then in at most $O(\ell)$ rounds, all vertices in V become informed, with probability at least $1 - 3n^{-\kappa+1}$. Finally, by Corollary 40, we obtain that for any source vertex in V, in at most $O(\ell)$ rounds, all vertices become informed, with probability at least $1 - 6n^{-\kappa+1}$. ◀

References

1 David Aldous and James Allen Fill. Reversible markov chains and random walks on graphs, 2002. Unfinished monograph, recompiled 2014. Available online.

2 Romas Aleliunas, Richard M. Karp, Richard J. Lipton, László Lovász, and Charles Rackoff. Random walks, universal traversal sequences, and the complexity of maze problems. In *Proc. 20th IEEE Symposium on Foundations of Computer Science, FOCS*, pages 218–223, 1979. doi:10.1109/SFCS.1979.34.

3 Noga Alon, Chen Avin, Michal Koucký, Gady Kozma, Zvi Lotker, and Mark R. Tuttle. Many random walks are faster than one. *Combinatorics, Probability & Computing*, 20(4):481–502, 2011. doi:10.1017/S0963548311000125.

4 O. S. M. Alves, F. P. Machado, and S. Yu. Popov. The shape theorem for the frog model. *The Annals of Applied Probability*, 12(2):533–546, May 2002. doi:10.1214/aoap/1026915614.

5 Yossi Azar, Andrei Z. Broder, Anna R. Karlin, and Eli Upfal. Balanced allocations. *SIAM J. Comput.*, 29(1):180–200, 1999. doi:10.1137/S0097539795288490.

6 Andrei Z. Broder, Anna R. Karlin, Prabhakar Raghavan, and Eli Upfal. Trading space for time in undirected s-t connectivity. *SIAM J. Comput.*, 23(2):324–334, 1994. doi:10.1137/S0097539790190144.

7 Flavio Chierichetti, George Giakkoupis, Silvio Lattanzi, and Alessandro Panconesi. Rumor spreading and conductance. *J. ACM*, 65(4):17:1–17:21, 2018. doi:10.1145/3173043.

8 Fan R. K. Chung and Lincoln Lu. Concentration inequalities and martingale inequalities: A survey. *Internet Mathematics*, 3(1):79–127, 2006. doi:10.1080/15427951.2006.10129115.

9 Colin Cooper. Random walks, interacting particles, dynamic networks: Randomness can be helpful. In *Proc. 18th International Colloquium on Structural Information and Communication Complexity, SIROCCO*, pages 1–14, 2011. doi:10.1007/978-3-642-22212-2_1.

10 Colin Cooper, Alan M. Frieze, and Tomasz Radzik. Multiple random walks in random regular graphs. *SIAM J. Discrete Math.*, 23(4):1738–1761, 2009. doi:10.1137/080729542.

11 Alan J. Demers, Daniel H. Greene, Carl Hauser, Wes Irish, John Larson, Scott Shenker, Howard E. Sturgis, Daniel C. Swinehart, and Douglas B. Terry. Epidemic algorithms for replicated database maintenance. *Operating Systems Review*, 22(1):8–32, 1988. doi:10.1145/43921.43922.

12 Tassos Dimitriou, Sotiris E. Nikoletseas, and Paul G. Spirakis. The infection time of graphs. *Discrete Applied Mathematics*, 154(18):2577–2589, 2006. doi:10.1016/j.dam.2006.04.026.

13 Benjamin Doerr, Mahmoud Fouz, and Tobias Friedrich. Social networks spread rumors in sublogarithmic time. In *Proc. 43rd ACM Symposium on Theory of Computing, STOC*, pages 21–30, 2011. doi:10.1145/1993636.1993640.

14 Devdatt P. Dubhashi and Alessandro Panconesi. *Concentration of Measure for the Analysis of Randomized Algorithms*. Cambridge University Press, 2009. doi:10.1017/CBO9780511581274.

15 Klim Efremenko and Omer Reingold. How well do random walks parallelize? In *Proc. Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, APPROX/RANDOM*, pages 476–489, 2009. doi:10.1007/978-3-642-03685-9_36.

16 Robert Elsässer and Thomas Sauerwald. Tight bounds for the cover time of multiple random walks. *Theor. Comput. Sci.*, 412(24):2623–2641, 2011. doi:10.1016/j.tcs.2010.08.010.

17 Uriel Feige, David Peleg, Prabhakar Raghavan, and Eli Upfal. Randomized broadcast in networks. *Random Struct. Algorithms*, 1(4):447–460, 1990. doi:10.1002/rsa.3240010406.

18 William Feller. *An Introduction to Probability Theory and its Applications. Vol. 1.* Wiley series in probability and mathematical statistics. Wiley, 1968.

19 George Giakkoupis, Frederik Mallmann-Trenn, and Hayk Saribekyan. How to spread a rumor: Call your neighbors or take a walk? In *Proc. 38th ACM Symposium on Principles of Distributed Computing, PODC*, pages 24–33, 2019. doi:10.1145/3293611.3331622.

20 George Giakkoupis, Frederik Mallmann-Trenn, and Hayk Saribekyan. How to spread a rumor: Call your neighbors or take a walk? *CoRR*, abs/2006.02368, 2020. URL: https://arxiv.org/abs/2006.02368, arXiv:2006.02368.

21 Peter Gracar and Alexandre Stauffer. Percolation of lipschitz surface and tight bounds on the spread of information among mobile agents. In *Proc. Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, APPROX/RANDOM*, volume 116 of *LIPIcs*, pages 39:1–39:17, 2018. doi:10.4230/LIPIcs.APPROX-RANDOM.2018.39.

22 Jonathan Hermon. Frogs on trees? *Electron. J. Probab.*, 23:40 pp., 2018. doi:10.1214/18-EJP144.

23 Richard M. Karp, Christian Schindelhauer, Scott Shenker, and Berthold Vöcking. Randomized rumor spreading. In *Proc. 41st IEEE Symposium on Foundations of Computer Science, FOCS*, pages 565–574, 2000. doi:10.1109/SFCS.2000.892324.

24 Harry Kesten and Vladas Sidoravicius. Branching random walk with catalysts. *Electronic Journal of Probability*, 8(0), 2003. URL: https://projecteuclid.org/euclid.ejp/1464037578, doi:10.1214/EJP.v8-127.

25 Harry Kesten and Vladas Sidoravicius. The spread of a rumor or infection in a moving population. *The Annals of Probability*, 33(6):2402–2462, Nov 2005. Zbl: 1111.60074. doi:10.1214/009117905000000413.

26 Harry Kesten and Vladas Sidoravicius. A shape theorem for the spread of an infection. *Annals of Mathematics*, 167(3):701–766, May 2008. doi:10.4007/annals.2008.167.701.

27 Axel Kramer, Ingeborg Schwebke, and Günter Kampf. How long do nosocomial pathogens persist on inanimate surfaces? A systematic review. *BMC Infectious Diseases*, 6(1):130, Dec 2006. doi:10.1186/1471-2334-6-130.
28 Henry Lam, Zhenming Liu, Michael Mitzenmacher, Xiaorui Sun, and Yajun Wang. Information dissemination via random walks in d-dimensional space. In *Proc. 23rd ACM-SIAM Symposium on Discrete Algorithms, SODA*, pages 1612–1622, 2012. doi:10.1137/1.9781611973099.128.
29 David A. Levin, Yuval Peres, and Elizabeth L. Wilmer. *Markov Chains and Mixing Times*. American Mathematical Society, 2nd edition, 2017.
30 Russell Lyons and Shayan Oveis Gharan. Sharp bounds on random walk eigenvalues via spectral embedding. *International Mathematics Research Notices*, 2018(24):7555–7605, 05 2017. doi:10.1093/imrn/rnx082.
31 Roberto I. Oliveira and Yuval Peres. Random walks on graphs: New bounds on hitting, meeting, coalescing and returning. In *Proc. 16th Workshop on Analytic Algorithmics and Combinatorics, ANALCO*, pages 119–126, 2019. doi:10.1137/1.9781611975505.13.
32 Alberto Pettarin, Andrea Pietracaprina, Geppino Pucci, and Eli Upfal. Infectious random walks. *CoRR*, abs/1007.1604:21 pp., 2010. URL: http://arxiv.org/abs/1007.1604, arXiv:1007.1604.
33 Alberto Pettarin, Andrea Pietracaprina, Geppino Pucci, and Eli Upfal. Tight bounds on information dissemination in sparse mobile networks. In *Proc. 30th ACM Symposium on Principles of Distributed Computing, PODC*, pages 355–362, 2011. doi:10.1145/1993806.1993882.
34 Serguei Yu. Popov. Frogs and some other interacting random walks models. In *Proc. Discrete Random Walks, DRW*, pages 277–288, 2003. URL: http://dmtcs.episciences.org/3328.

APPENDIX

A Auxiliary Lemmas for Trees

Consider a random walk on a b-ary tree $R_{b,h}$ of height h, with $n = (b^{h+1} - 1)/(b - 1)$ vertices. Denote the number of rounds the walk takes to reach from vertex u to v as $\tau_{u,v}$. The return time τ_u^+ is the number of rounds it takes for the walk starting from u to return to u. Let ρ be the root, and v be any leaf of the tree. Then,

$$\mathbb{E}\left[\tau_\rho^+\right] = \frac{1}{\pi(\rho)} = \frac{2(b^h - 1)}{b - 1} \geq 2b^{h-1}, \tag{14}$$

$$\mathbb{E}\left[\tau_{\rho,v}\right] = 2hb^h + \frac{2hb^h}{b-1} - \frac{2b^{h+1} - 2b}{(b-1)^2} - h \leq 4hb^h, \tag{15}$$

$$\mathbb{E}\left[\tau_{v,\rho}\right] = 2b \cdot \frac{b^h - 1}{(b-1)^2} - h \cdot \frac{b+1}{b-1} \in [b^{h-1}, 8b^{h-1}], \tag{16}$$

where the equalities in (15) and (16) are derived analogously to the binary tree example in [29, Example 10.17].

▶ **Lemma 36.** *For the root ρ of tree $R_{b,h}$ and any leaf vertex v,*

$$\mathbb{P}\left[\tau_{\rho,v} < \min\{\tau_\rho^+, 8hb^{h-1}\}\right] \geq \frac{1}{12hb}.$$

Proof. Let $\mathcal{E}$ denote the event $\{\tau_{\rho,v} < \tau_\rho^+\}$, that a walk starting from ρ hits v before returning to ρ. We have

$$\mathbb{E}\left[\tau_{\rho,v}\right] \geq \mathbb{E}\left[\tau_{\rho,v} \mid \neg\mathcal{E}\right] \cdot \mathbb{P}\left[\neg\mathcal{E}\right] = \left(\mathbb{E}\left[\tau_\rho^+ \mid \neg\mathcal{E}\right] + \mathbb{E}\left[\tau_{\rho,v}\right]\right) \cdot \mathbb{P}\left[\neg\mathcal{E}\right].$$

Also

$$\mathbb{E}\left[\tau_\rho^+ \mid \neg\mathcal{E}\right] \cdot \mathbb{P}\left[\neg\mathcal{E}\right] = \mathbb{E}\left[\tau_\rho^+\right] - \mathbb{E}\left[\tau_\rho^+ \mid \mathcal{E}\right] \cdot \mathbb{P}\left[\mathcal{E}\right],$$

and

$$\mathbb{E}\left[\tau_\rho^+ \mid \mathcal{E}\right] = \mathbb{E}\left[\tau_{\rho,v} \mid \mathcal{E}\right] + \mathbb{E}\left[\tau_{v,\rho}\right] \le \mathbb{E}\left[\tau_{\rho,v}\right] + \mathbb{E}\left[\tau_{v,\rho}\right].$$

Combining these three inequalities we obtain

$$\mathbb{E}\left[\tau_{\rho,v}\right] \ge \mathbb{E}\left[\tau_\rho^+\right] - \left(\mathbb{E}\left[\tau_{\rho,v}\right] + \mathbb{E}\left[\tau_{v,\rho}\right]\right) \cdot \mathbb{P}\left[\mathcal{E}\right] + \mathbb{E}\left[\tau_{\rho,v}\right] \cdot \mathbb{P}\left[\neg\mathcal{E}\right].$$

Substituting $\mathbb{P}\left[\neg\mathcal{E}\right] = 1 - \mathbb{P}\left[\mathcal{E}\right]$, solving for $\mathbb{P}\left[\mathcal{E}\right]$, and using (14) to (16), yields

$$\mathbb{P}\left[\mathcal{E}\right] \ge \frac{\mathbb{E}\left[\tau_\rho^+\right]}{2\mathbb{E}\left[\tau_{\rho,v}\right] + \mathbb{E}\left[\tau_{v,\rho}\right]} \ge \frac{2b^{h-1}}{8hb^h + 8b^{h-1}} = \frac{1}{4(hb+1)} \ge \frac{1}{6hb}.$$

Next, we bound $\mathbb{P}\left[\tau_{\rho,v} < 4hb^{h-1} \mid \mathcal{E}\right]$. Let u be the child of ρ that is also an ancestor of v. Then, given $\mathcal{E}$, until the walk returns to ρ, the walk is a restricted to the subtree of u, which is a b-ary tree of height $h-1$. In particular, in the first step, the walk visits u. Therefore, by (15), we have that $\mathbb{E}\left[\tau_{\rho,v} - 1 \mid \mathcal{E}\right] \le 4(h-1)b^{h-1} < 4hb^{h-1}$. Then, by Markov's inequality,

$$\mathbb{P}\left[\tau_{\rho,v} < 8hb^{h-1} \mid \mathcal{E}\right] \ge 1/2.$$

Finally,

$$\begin{aligned}\mathbb{P}\left[\tau_{\rho,v} < \min\{\tau_\rho^+, 8hb^{h-1}\}\right] &= \mathbb{P}\left[\mathcal{E} \cap \{\tau_{\rho,v} < 8hb^{h-1}\}\right] \\ &= \mathbb{P}\left[\tau_{\rho,v} < 8hb^{h-1} \mid \mathcal{E}\right] \cdot \mathbb{P}\left[\mathcal{E}\right] \ge \frac{1}{12hb}.\end{aligned}$$ ◀

▶ **Lemma 37.** *Let v be a leaf of $R_{b,h}$, and u be its ancestor of height $x \ge 1$. For any $\epsilon > 0$,*

$$\mathbb{P}\left[\tau_{v,u} \ge \epsilon \cdot b^{x-1}\right] \ge 1 - \epsilon.$$

Proof. For brevity denote $k' = \epsilon \cdot b^{x-1}$ and $p = \mathbb{P}\left[\tau_{v,u} \ge k'\right]$. Then, for an integer $i \ge 1$,

$$\begin{aligned}\mathbb{P}\left[\tau_{v,u} \ge i \cdot k'\right] &= \mathbb{P}\left[\tau_{v,u} \ge i \cdot k' \mid \tau_{v,u} \ge (i-1) \cdot k'\right] \cdot \mathbb{P}\left[\tau_{v,u} \ge (i-1) \cdot k'\right] \\ &\le p \cdot \mathbb{P}\left[\tau_{v,u} \ge (i-1) \cdot k'\right], \\ &\qquad \text{by Markov property and because } \mathbb{P}\left[\tau_{v',u} \ge k'\right] \le p \text{ for any } v', \\ &\le p^i, \text{ by iterating the argument.}\end{aligned}$$

This implies that

$$\mathbb{E}\left[\tau_{v,u}\right] = \sum_{t \ge 1} \mathbb{P}\left[\tau_{v,u} \ge t\right] \le \sum_{i \ge 0} k' \cdot \mathbb{P}\left[\tau_{v,u} \ge i \cdot k'\right] \le \frac{k'}{1-p} = \frac{\epsilon \cdot b^{x-1}}{1-p}.$$

By (16), $\mathbb{E}\left[\tau_{v,u}\right] \ge b^{x-1}$. Combining the last two inequalities, we get $p \ge 1 - \epsilon$. ◀

▶ **Lemma 38.** *Let $X(t)$ be the location of a simple random walk that starts at a leaf v of $R_{b,h}$ at round 0. Then, for any even integer $t > 0$,*

$$\mathbb{P}\left[X(t) = v\right] \ge \max\left(\frac{1}{16bt}, \frac{1}{2n}\right).$$

Proof. For even rounds t, $\mathbb{P}[X(t)=v]$ monotonically decreases towards the stationary distribution at v. Thus,

$$\mathbb{P}[X(t)=v] \geq \frac{1}{2(n-1)} \geq \frac{1}{2n}. \tag{17}$$

It implies that we have to show the inequality in the case when the first term under the max is larger, i.e., when $t \leq n/(4b)$. Let $x = 1 + \lceil \log_b(2t) \rceil$. First, we prove that v has an ancestor of height x, i.e., that $x \leq h$.

$$\begin{aligned} x &= 1 + \lceil \log_b(2t) \rceil \\ &\leq 1 + \lceil \log_b(n/(2b)) \rceil \\ &\leq 1 + \lceil \log_b(2 \cdot b^h/(2b) \rceil, \text{ because } n \leq 2 \cdot b^h, \\ &\leq 1 + \lceil \log_b(b^{h-1}) \rceil \\ &= h. \end{aligned}$$

Thus, we can define u as the ancestor of v of height x. Let $\mathcal{R} = \{\tau_{v,u} > t\}$ be the event that the random walk $X(t)$ does not visit u in the first t rounds. Since $t \leq b^{x-1}/2$ by construction, then

$$\mathbb{P}[\mathcal{R}] \geq \mathbb{P}\left[\tau_{v,u} \geq b^{x-1}/2\right] \geq \frac{1}{2},$$

by Lemma 37. Then,

$$\mathbb{P}[X(t)=v] \geq \mathbb{P}[X(t)=v \mid \mathcal{R}] \cdot \mathbb{P}[\mathcal{R}] \geq \frac{1}{2} \cdot \mathbb{P}[X(t)=v \mid \mathcal{R}].$$

Let u' be the child of u that is also an ancestor of v. Let k' be the number of vertices in the subtree of u'. Given the event $\mathcal{R}$, the random walk can only visit vertices in the subtree of u', thus, for an even t, as in (17),

$$\mathbb{P}[X(t)=v \mid \mathcal{R}] \geq \frac{1}{2k'} = \frac{1}{2(1+b+\cdots+b^{x-1})} \geq \frac{1}{4 \cdot b^{x-1}} \geq \frac{1}{16bt},$$

where the last inequality holds because $t > b^{x-2}/2$ by the choice of x. Combining with the previous inequality we finish the proof. ◀

B Other Auxiliary Lemmas

B.1 Symmetry of Visit-Exchange

The following results allow us to assume that in VISIT-EXCHANGE the information starts at a vertex convenient for us.

▶ **Lemma 39.** *For vertices u and v of a connected graph $G=(V,E)$, let $T_{u,v}$ be the number of rounds of* VISIT-EXCHANGE *until v is informed when the information originates at u. Then for any round r, $\mathbb{P}[T_{u,v} \leq r] = \mathbb{P}[T_{v,u} \leq r]$.*

Proof. For round r, let Ω_r be the set of all possible executions in the first r rounds, i.e., every element $\omega \in \Omega_r$ is composed of the paths that are taken by each of the agents. Let $p(\omega)$ be the probability associated with the outcome ω. For an execution ω, let ω^* be the *reversal* of ω: If $X_g(t)$ is the walk taken by agent g then $X_g(t)(\omega^*) = X_g(r-t)(\omega)$ for $t \in \{0,\ldots,r\}$. Clearly $\omega^* \in \Omega_r$.

Since g starts its walk from the stationary distribution π, for any path $u_0, \ldots, u_r$,

$$\mathbb{P}\left[X_g(t) = u_t \text{ for all } t\right] = \pi(u_0) \prod_{t=0}^{r-1} \frac{1}{\deg(u_t)} = \frac{1}{2|E|} \prod_{t=1}^{r-1} \frac{1}{\deg(u_t)} = \mathbb{P}\left[X_g(t) = u_{r-t} \text{ for all } t\right].$$

Applying this equality for every agent g, we get that $p(\omega) = p(\omega^*)$.

Additionally, notice that if in execution ω vertex v gets informed when u is the source, then u gets informed in ω^* when v is the source. Combining the two previous facts gives

$$\mathbb{P}\left[T_{u,v} \le r\right] = \sum_{\omega \in \{T_{u,v} \le r\}} p(\omega) = \sum_{\omega^* \in \{T_{v,u} \le r\}} p(\omega^*) = \mathbb{P}\left[T_{v,u} \le r\right].$$ ◀

▶ **Corollary 40.** *Let u and v be any two vertices of a connected graph $G = (V, E)$, and let T_u denote the spreading time of* VISIT-EXCHANGE *when the information originates at u. If $\mathbb{P}[T_v \le r] = 1 - \delta$ then $\mathbb{P}[T_u \le 2r] \ge 1 - 2\delta$.*

Proof. We have

$$\begin{aligned}
\mathbb{P}\left[T_u \le 2r\right] &\ge \mathbb{P}\left[T_u \le 2r \mid T_{u,v} \le r\right] \cdot \mathbb{P}\left[T_{u,v} \le r\right] \\
&\ge \mathbb{P}\left[T_v \le r\right] \cdot \mathbb{P}\left[T_{u,v} \le r\right] \\
&\ge \mathbb{P}\left[T_v \le r\right] \cdot \mathbb{P}\left[T_{v,u} \le r\right], \quad \text{by Lemma 39,} \\
&\ge \mathbb{P}\left[T_v \le r\right]^2, \quad \text{since } T_{v,u} \le T_v, \\
&\ge 1 - 2\delta.
\end{aligned}$$ ◀

B.2 Return Probability on $\mathbb{Z}$

▶ **Lemma 41.** *Let $X(t)$ be a lazy simple random walk on $\mathbb{Z}$ starting from the origin, with a constant holding probability $\alpha > 0$. Then, there is a constant η, such that*

$$\mathbb{P}\left[X(t) = 0\right] \le \frac{\eta}{\sqrt{t}}.$$

Proof. Suppose N is the number of times the walk moves (rather than staying put) in the first t rounds. N is a sum of independent Bernoulli trials with success probability $1 - \alpha$, thus, $\mathbb{E}[N] = (1 - \alpha)t$. Let $\mathcal{E}$ be the event that $N \ge (1 - \alpha)t/2$. By a Chernoff bound,

$$\mathbb{P}\left[\mathcal{E}\right] \ge 1 - e^{-(1-\alpha)t/8}.$$

Let t' be the smallest even integer that is at least $(1 - \alpha)t/2$ and let X' be a (non-lazy) simple random walk on $\mathbb{Z}$ starting from the origin. Then,

$$\begin{aligned}
\mathbb{P}\left[X(t) = 0\right] &\le \mathbb{P}\left[X(t) = 0 \mid \mathcal{E}\right] + \mathbb{P}\left[\neg\mathcal{E}\right] \\
&\le \mathbb{P}\left[X'(t') = 0\right] + e^{-(1-\alpha)t/8} \\
&= \binom{t'}{t'/2} \cdot 2^{-t'} + e^{-(1-\alpha)t/8} \\
&\le \frac{1}{\sqrt{t'}} + e^{-(1-\alpha)t/8} \\
&\le \sqrt{\frac{2}{(1-\alpha)t}} + e^{-(1-\alpha)t/8} \\
&\le \frac{\eta}{\sqrt{t}},
\end{aligned}$$

for a sufficiently large constant η. ◀

B.3 Chernoff Bounds

▶ **Lemma 42.** *If X is a sum of independent Bernoulli random variables, then for $b > 1$ and $x > b \cdot \mathbb{E}[X]$, $\mathbb{P}[X \geq x] \leq (b/e)^{-x}$.*

Proof. Let $\delta = \frac{x}{\mathbb{E}[X]} - 1$, then, by a standard Chernoff inequality,

$$\begin{aligned}\mathbb{P}[X \geq x] &= \mathbb{P}[X \geq (1+\delta) \cdot \mathbb{E}[X]] \\ &\leq \left(\frac{e^{\delta}}{(1+\delta)^{1+\delta}}\right)^{\mathbb{E}[X]} \leq \left(\frac{e}{1+\delta}\right)^{(1+\delta)\mathbb{E}[X]} \\ &= \left(\frac{e \cdot \mathbb{E}[X]}{x}\right)^{x} \leq \left(\frac{b}{e}\right)^{-x}.\end{aligned}$$

◀

▶ **Lemma 43.** *Let $F_1, \ldots, F_n$ be independent and identical geometrically distributed random variables with parameter p, i.e., for any integer $k \geq 1$,*

$$\mathbb{P}[F_i = k] = (1-p)^{k-1} p.$$

Let $F = \sum_{i=1}^{n} F_i$. Then for any $k \geq 2 \cdot \mathbb{E}[F] = 2n/p$,

$$\mathbb{P}[F \geq k] \leq \exp\left(-\frac{kp}{8}\right).$$

Proof. We define a coupling between random variables $(F_i)_{i=1}^{n}$ and a sequence of Bernoulli trials $(X_j)_{j=1}^{\infty}$ with parameter p. Let $j_0 = 0$ and for $i \geq 1$, let $j_i = \min\{j > j_{i-1} : X_j = 1\}$, i.e., j_i is the index of ith 1 in (X_j). We set $F_i = j_i - j_{i-1}$. With this coupling, $F \geq k$ implies $Y_k = \sum_{j=1}^{k} X_j \leq n$. Therefore, $\mathbb{P}[F \geq k] \leq \mathbb{P}[Y_k \leq n]$, which we can bound using a standard Chernoff bound. We have that $\mathbb{E}[Y_k] = kp \geq 2n$, and thus,

$$\mathbb{P}[F \geq k] \leq \mathbb{P}[Y_k \leq n] \leq \mathbb{P}[Y_k \leq \mathbb{E}[Y_k]/2] \leq \exp\left(-\frac{kp}{2}\right).$$

◀