Rumor Spreading and Vertex Expansion

George Giakkoupis* Thomas Sauerwald†

Abstract

We study the relation between the rate at which rumors spread throughout a graph and the vertex expansion of the graph. We consider the standard rumor spreading protocol where every node chooses a random neighbor in each round and the two nodes exchange the rumors they know. For any n-node graph with vertex expansion α, we show that this protocol spreads a rumor from a single node to all other nodes in $\mathcal{O}(\alpha^{-1}\log^2 n\sqrt{\log n})$ rounds with high probability. Further, we construct graphs for which $\Omega(\alpha^{-1}\log^2 n)$ rounds are needed. Our results complement a long series of works that relate rumor spreading to edge-based notions of expansion, resolving one of the most natural questions on the connection between rumor spreading and expansion.

1 Introduction

Epidemic protocols have become an important primitive for information dissemination in networks. A prominent example is the so-called *randomized rumor spreading protocols*. These protocols disseminate a piece of information, or *rumor*, from a single node of a connected n-node network, to all the other nodes. The paradigm underlying these protocols is that each node chooses a random neighbor to communicate with in every round. This simple and local communication rule has proven to be very effective for several network topologies, and also robust against changes in the network topology, e.g., due to failures [16, 19]. Randomized rumor spreading protocols and variations thereof have been used successfully in various contexts, including maintenance of replicated databases [13], failure detection [36], resource discovery [24], data aggregation [6], and modeling the spread of computer viruses [5].

The most basic and well-studied variant of randomized rumor spreading protocols is the *PUSH protocol*. The protocol proceeds in a sequence of synchronous rounds, and in each round every *informed* node (i.e., every node that learned the rumor in a previous round) chooses a neighbor uniformly at random and sends the rumor to it. The *PULL protocol* is symmetric: In each round, every *uninformed* node chooses a random neighbor, and if that neighbor knows the rumor it sends it to the uniformed node. Finally, the *PUSH-PULL protocol* is the combination of both strategies: In each round, every node chooses a random neighbor to send the rumor to, if the node knows the rumor, or to request the rumor from, otherwise. A primary performance measure of these protocols is their *runtime*, that is, the number of rounds required until a rumor started by a single node spreads to all other nodes.

The above protocols have been shown to be very efficient for several network topologies. In particular, their runtime is exponentially smaller than the network size for topologies ranging from basic networks, such as complete graphs and hypercubes, to more complex structures, such as preferential attachment graphs modeling social networks [22, 34, 13, 19, 26, 8, 15, 16, 20, 21, 14] (for more details, see the Related Work Section).

The success of these protocols on specific networks motivated the search for general network properties that yield fast rumor spreading. One important such property is high expansion. Most of the networks for which rumor spreading is known to be fast share this property. Further, theoretical and empirical studies indicate that social networks also have high conductance [17, 32, 18, 28].

Several recent works have investigated the relation of the runtime of rumor spreading protocols to the *conductance* of the underlying graph [6, 33, 10, 9, 23, 7]. The conductance $\Phi \in (0, 1]$ of a graph is a standard measure of the graph's expansion, and is defined roughly as the minimum ratio of the edges leaving a set of vertices over the volume of that set, i.e., the total number of edges incident to the set (see Section 2 for the precise definition). Those works culminated in a runtime bound of $\mathcal{O}(\Phi^{-1}\log n)$ rounds for the PUSH protocol on any regular graph [33],[1] and the same bound for the PUSH-PULL protocol on general graphs [23]. These upper bounds are tight, as there are graphs with

*University of Calgary, Canada, ggiakkou@ucalgary.ca. Supported in part by a Postdoctoral Fellowship with the Pacific Institute for the Mathematical Sciences (PIMS).

†Max-Planck-Institut für Informatik, Saarbrücken, Germany, sauerwal@mpi-inf.mpg.de.

[1] The authors of [33] consider a different definition for conductance. However, for regular graphs the two definitions are equivalent.

diameter $\Omega(\Phi^{-1}\log n)$ [9].

In this paper we focus on *vertex expansion*, which is another standard notion of expansion. Vertex expansion has proven relevant in many areas, including expander graphs [3, 25], random walks [4, 29], and property testing [12]. The vertex expansion $\alpha \in (0, 1]$ of a graph is, roughly, the minimum ratio of the neighbors that a set of vertices has (and are not in this set), over the cardinality of the set. Although this definition looks similar to that of conductance, vertex expansion and conductance can behave quite differently. For example, a star with n vertices has constant conductance, but vertex expansion $\Theta(1/n)$. On the other hand, an $(n/2)$-regular graph consisting of a matching between two cliques of size $n/2$ has constant vertex expansion, but conductance $\Theta(1/n)$. In fact, for any regular graph, the vertex expansion is at least as large as the conductance.

An interesting property of vertex expansion is that it is monotone under edge addition, whereas conductance is not. For example, consider a 3-regular expander graph—such a graph has constant conductance and vertex expansion. Partition the vertices into two sets S_1, S_2 of size $n/2$, and connect all vertices that are in the same set (and are not already connected). The vertex expansion of the resulting graph is still constant, but its conductance is $\mathcal{O}(1/n)$. Further, rumor spreading takes $\Theta(\log n)$ rounds as in the original graph. Hence in this case, an upper bound on rumor spreading based on vertex expansion would be more appropriate than a bound based on conductance. This monotonicity property of vertex expansion might also be useful for constructing networks that support fast rumor spreading, or for verifying that a network allows fast rumor spreading, e.g., by identifying a (spanning) subgraph with high vertex expansion.

Unlike the relation between conductance and rumor spreading, which is fairly well understood, very little is know about how vertex expansion relates to rumor spreading. The authors of [10] highlighted the question of whether high vertex expansion implies fast rumor spreading as an "outstanding open problem." Some progress in this problem was recently made in [35], where a runtime bound of $\mathcal{O}(\alpha^{-1}\log^5 n)$ rounds was shown for the PUSH protocol on *regular* graphs. This result does not extend to general graphs, since there are simple examples of graphs with constant vertex expansion where the PUSH (or PULL) protocol takes a polynomial number of rounds. For the runtime of the PUSH-PULL protocol on general graphs, only a polynomial upper bound was known so far, of $\mathcal{O}(\alpha^{-1}n^{1-\epsilon})$ rounds [35], where $\epsilon > 0$ is a small constant.

Our Contribution. In this paper we give an almost complete picture of the connection between the runtime of randomized rumor spreading and vertex expansion. Our main result is the following upper bound on the runtime of the PUSH-PULL protocol in general graphs.

THEOREM 1.1. *For any n-vertex graph with vertex expansion at least α, minimum degree δ, and maximum degree Δ, the PUSH-PULL protocol informs all vertices in $\mathcal{O}(\alpha^{-1}\log n\log\Delta\sqrt{\log(2\Delta/\delta)}) = \mathcal{O}(\alpha^{-1}\log^{2.5} n)$ rounds with high probability.*[2]

This result answers in the affirmative the question of [10] whether high vertex expansion implies fast rumor spreading.

We complement the upper bound of Theorem 1.1 by providing an almost matching lower bound: We construct d-regular graphs for which rumor spreading takes $\Omega(\alpha^{-1}\log n\log d)$ rounds, whenever $d/\alpha = \mathcal{O}(n^{1-\epsilon})$ (see Theorem 5.1). Note that when the ratio of maximum over minimum degree is bounded by a constant, the two bounds match. Thus, for d-regular graphs we obtain a tight bound of $\mathcal{O}(\alpha^{-1}\log n\log d)$ rounds on the runtime of the PUSH-PULL protocol. Further we show that the same bound for regular graphs also holds for the PUSH and the PULL protocols by themselves (Theorem 4.1), thus improving the results of [35] for the PUSH protocol.

Following [10], we can view rumor spreading as a graph sparsification procedure. In this context, Theorem 1.1 implies that for any graph with vertex-expansion α, the edges used by the PUSH-PULL protocol induce a subgraph with $\mathcal{O}(\alpha^{-1}n\log^{2.5} n)$ edges, so that every pair of vertices is connected by a path of length $\mathcal{O}(\alpha^{-1}\log^{2.5} n)$. The relation between vertex expansion and rumor spreading could possibly yield further results of this kind in the context of network sampling [27, 28, 30].

Our analysis is substantially different than previous analyses of rumor spreading. Most of the results that bound rumor spreading with conductance essentially measure the growth of the volume of the informed vertices, which, in expectation, can be easily shown to grow at a rate of roughly $1 + \Phi$ per round. While this argument alone is not sufficient to complete the analysis, it provides the right intuition and also serves as a basic building block in the proof (see [23]). In the case of vertex expansion, there seems to be no obvious parameter like the volume or the size of the set of informed vertices that captures the progress of rumor spreading. Instead, in our analysis we measure

[2] An event $\mathcal{E}(n)$ holds *with high probability (w.h.p.)*, if $\mathbf{Pr}(\mathcal{E}(n)) = 1 - \mathcal{O}(n^{-c})$ for an arbitrary constant $c > 0$.

the progress by considering the sum of twice the size of the set of informed vertices plus the size of its outer boundary. We then establish that this quantity increases at a rate of roughly $1+\alpha/\log^{3/2} n$.

Our analysis is also very different from that in [35]. In fact, it seems difficult to extend the techniques used there for regular graphs to general graphs. One reason is that their analysis uses also *upper* bounds on the number of informed nodes (see the upper bound in [35, Lemma 3.2] and how it is used in the key lemma [35, Lemma 3.6]); these bounds do not hold for general graphs.

Finally, our analysis employs some techniques that may be helpful in further studies of randomized rumor spreading. These include a Bootstrap Percolation-like process and its analysis based on the potential method of amortized analysis (in the proof of Lemma 3.7), and the use of the Optional Stopping Theorem to lower bound the progress of the rumor spreading process (in Claim 3.2).

Related Work. The first works on rumor spreading provided a precise analysis of the runtime of the PUSH protocol on complete graphs [22, 34]. A runtime bound of $\mathcal{O}(\log n)$ rounds was later established for hypercubes and random graphs [19]. Other symmetric graphs G similar to the hypercube in which rumor spreading takes $\mathcal{O}(\log n + \operatorname{diam}(G))$ rounds were studied in [15]. A refined analysis for random graphs proving that the runtime is essentially the same as that on complete graphs was provided in [20], and subsequently extended to random regular graphs in [21]. The authors of [8] studied rumor spreading on preferential attachment graphs, and showed that both the PUSH and the PULL protocols require polynomially many rounds, whereas the PUSH-PULL protocol takes only $\mathcal{O}(\log^2 n)$ rounds. The result for the PUSH-PULL protocol was recently improved to $\Theta(\log n)$ [14]. These two results highlight the necessity of using the PUSH-PULL protocol in highly non-regular graphs. (See also [23].)

The first explicit connection between randomized rumor spreading and graph expansion was shown in [33], where a bound of $\mathcal{O}(\Phi^{-1}\log n)$ rounds was proved for the runtime of the PUSH protocol on regular graphs[3]. A comparable bound involving the spectral gap was shown in [6], however, as the result of [33], it does not extend to non-regular graphs. The PUSH and PULL protocols by themselves cannot guarantee fast rumor spreading based solely on the assumption of high conductance (see, e.g, [8] for counter-examples). For the PUSH-PULL protocol, a bound for general graphs slightly weaker than the above bound for regular graphs was shown in [9]. A tight bound of $\mathcal{O}(\Phi^{-1}\log n)$ was subsequently proved in [23]. This result is tight in the sense that there are graphs of diameter $\Omega(\Phi^{-1}\log n)$ [9]. Recently, the authors of [7] introduced a refinement of conductance, called weak conductance, and related this quantity to the time needed to inform a certain fraction of nodes.

[3] The authors of [33] obtained results for non-regular graphs as well, but in this case neighbors are not chosen uniformly at random (see also [10] for an explanation).

2 Definitions and Notation

We consider graphs $G=(V,E)$ that are undirected and connected. We denote by n the number of vertices, $n := |V|$. For any vertex $u \in V$, $N(u)$ denotes the set of neighbors of u, $N(u) := \{v \in V : \{u,v\} \in E\}$, and $d(u)$ is the degree of u, $d(u) := |N(u)|$. The maximum degree of G is $\Delta := \max_{u\in V} d(u)$, and the minimum degree is δ. For any set of vertices $U \subseteq V$, ∂U denotes the *(outer) boundary* of U, $\partial U := \{u \in V \setminus U : \exists\, v \in U, \{u,v\} \in E\}$. The *volume* of U is the sum of the degrees of the vertices in U, $\operatorname{vol}(U) := \sum_{u\in U} d(u)$. Note that $\operatorname{vol}(V) = 2|E|$. For any two sets $U, W \subseteq V$, $E(U,W)$ is the set of edges between U and W, $E(U,W) := \{\{u,w\} \in E : u \in U, w \in W\}$. The *vertex expansion* of G is a real number $0 < \alpha \le 1$ defined by

$$\alpha = \alpha(G) := \min_{U\subseteq V,\ 0<|U|\le n/2} \frac{|\partial U|}{|U|}.$$

The *conductance* of G is also a real $0 < \Phi \le 1$, and is defined by

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

For any graph, $(\delta/\Delta)\Phi \le \alpha \le \Delta\Phi$. (The proof of this simple fact can be found in the Appendix.)

For the analysis of the rumor spreading protocols we assume that the rumor starts from an arbitrary vertex, which learns the rumor in round 0, and begins to spread it in round 1. We denote by I_t the set of informed vertices at the end of round $t \ge 0$ (thus $|I_0| = 1$). We use the following terminology. An informed vertex $u \in V$ *pushes* the rumor to vertex $v \in N(u)$ in some round, if in that round u picks v to transmit the rumor to; we say that this transmission is *initiated* by u. Similarly, an uninformed vertex v *pulls* the rumor from an informed vertex $u \in N(v)$, if v picks u and thus the rumor is transmitted from u to v; this transmission is *initiated* by v.

All logarithms in this paper are to the base 2 unless stated otherwise.

3 Proof of the Upper Bound

In this section we describe the proof of Theorem 1.1.

3.1 Outline of the Proof. Intuitively, we measure the progress of rumor spreading in terms of two quantities: the growth of the set I_t of informed vertices, and the growth of the outer boundary ∂I_t of I_t. We conveniently capture both these quantities by considering the quantity $W_t = 2|I_t| + |\partial I_t|$.[4] Roughly speaking, we show that as long as no more than half of the vertices are informed, W_t increases by a factor of $1 + \Omega(\beta)$ per round "on average," where $\beta = \alpha/(\log \Delta \sqrt{\log(2\Delta/\delta)})$. The key ingredient we use is Lemma 3.1, which states that for any t, there is a number $r \leq 1/\beta$ depending only on the current set of informed vertices I_t, such that within r rounds W_t increases by a factor of at least $1 + r\beta$ with constant probability (w.c.p.). Having established this result we employ the Optional Stopping Theorem (in Claim 3.2) to obtain that W_t doubles in at most $1/\beta$ rounds w.c.p; and since W_t cannot grow larger than $2n$, and thus cannot double more than $\log(2n)$ times, we obtain that $\mathcal{O}(\beta^{-1} \log n)$ rounds suffice to inform half (plus one) of the vertices w.c.p., and also with high probability (w.h.p.). We finish the proof by employing a simple symmetry result for the PUSH-PULL protocol (Lemma 3.3), which yields that $\mathcal{O}(\beta^{-1} \log n)$ additional rounds will inform the remaining vertices w.h.p.

We now give an overview of the proof of our main lemma, Lemma 3.1. For the sake of clarity, we discuss a weaker version of the lemma, where $\beta := \alpha/\log^{3/2} n$. We partition the boundary ∂I_t into sets $A_1, \ldots, A_{\log n}$, where each set A_i consists of the vertices in ∂I_t of degree between $d_i := 2^{i-1}$ and $2d_i$. We then analyze the contribution to the increase in W_t of each A_i individually. Here we consider only those A_i that are sufficiently large and their degree d_i is in a certain range ($|A_i| = \Omega(|\partial I_t|/\log n)$, and $d_i = \mathcal{O}(|A_i|)$ or $\Omega(|\partial I_t|)$); the remaining sets constitute just a small fraction on ∂I_t. We use a case analysis that distinguishes essentially four different cases. The first two cases are fairly straightforward. The first is when for some constant fraction of the vertices $u \in A_i$, a sufficiently large fraction q of the neighbors of each u are in I_t ($q = \Omega(1/\log^{3/2} n)$). Then, a constant fraction of A_i gets informed quickly just by pulling the rumor from I_t (in $r = \mathcal{O}(1/q)$ rounds w.c.p., by Lemma 3.5). We easily show that this yields an increase in W_t equal to the desired increase (i.e., $\Omega((\beta/q) \cdot |W_t|)$) times $|A_i|/|\partial I_t|$.

The second case is when for a constant fraction of the vertices $u \in A_i$, a constant fraction of the neighbors of each u are in $S_t := V \setminus (I_t \cup \partial I_t)$, and also d_i is sufficiently large ($d_i = \Omega(|\partial I_t|)$). Then by informing a single such vertex u, $\Theta(d_i)$ vertices are added to the boundary, and thus W_t is increase by $\Theta(d_i)$. The number r of rounds needed is such that the increase of $\Theta(d_i)$ in W_t is as desired, i.e., $\Omega(r\beta|W_t|)$. (The bound on r is given in Lemma 3.4.)

The third case is more involved. As in the second case, for a constant fraction of the $u \in A_i$, a constant fraction of the neighbors of each u are in S_t. However, now d_i is smaller ($d_i = \mathcal{O}(A_i)$) and thus it is not enough to inform a single vertex. Further, this case assumes that most of the edges from A_i to S_t go to vertices $v \in S_t$ such that each v has at most $\mathcal{O}(d_i)$ neighbors in A_i. We have that in a constant number of rounds the fraction of A_i that pulls the rumor from I_t results in $\mathcal{O}(A_i)$ vertices being added to the boundary w.c.p. (by Lemma 3.6). Again we show that the resulting increase in W_t is the desired one.

For those A_i for which none of the above three cases applies, it holds that a constant fraction of the edges from A_i goes to vertices $v \in V \setminus I_t$ such that each v has a large number of neighbors in A_i (at least $\Omega(d_i/\log n)$); this is the last and most difficult case. We show that there exist two sets $A_i' \subseteq A_i$ and $V' \subseteq V \setminus I_t$ such that each vertex in A_i' has $\Theta(d_i)$ neighbors in V' and each vertex in V' has $\Omega(d_i/\log n)$ neighbors in A_i', and also $|A_i'| = \Theta(|A_i|)$. Then we argue that if we started a new rumor at any vertex $u \in A_i'$ then this rumor would spread quickly (in $\mathcal{O}(\log^{3/2} n)$ rounds w.c.p.) to a constant fraction of A_i', through the vertices in V', and from there to a vertex in I_t. By using the same symmetry argument as before (Lemma 3.3), we obtain that in the same number of rounds the original rumor spreads from I_t to a constant fraction of A_i'. (See Lemma 3.7.) As in the first case, the resulting increase in W_t is equal to the desired increase times $|A_i|/|\partial I_t|$.

If at least one of the A_i meets the conditions of the second or third cases then the lemma follows immediately. Otherwise, we accumulate the progress contributed by each A_i to establish the total increase in W_t. A crucial point is that in the analysis of the first and last cases above, when we count the contribution of A_i we rely only on rumor transmissions initiated by vertices in A_i; thus contribution by different sets are easy to combine.

Roadmap. We give the formal statement of Lemma 3.1 and the derivation of Theorem 1.1 from it in Section 3.2. In Section 3.3 we describe the four lemmata that we use in the proof of Lemma 3.1, i.e., Lemmata 3.4–3.7. The proof of the last lemma, Lemma 3.7, is given in a separate section, Section 3.4. Finally, we derive Lemma 3.1 in Section 3.5.

[4] Our proof works also if we consider instead the maximum between $|I_t|$ and $|I_t| + |\partial I_t|$. However, considering just $|I_t| + |\partial I_t|$ does not work because this sum may grow slowly even if $|I_t|$ grows fast, e.g., when the neighbors of the newly informed vertices are already in the boundary, and thus the increase in $|I_t|$ is "canceled" by a decrease in $|\partial I_t|$.

3.2 Statement of the Main Lemma and Proof of Theorem 1.1. Recall that I_t is the set of informed vertices after round t, and $\partial I_t := \bigcup_{u \in I_t} N(u) \setminus I_t$ is the boundary of I_t. We define

$$Z_t := |I_t| + |\partial I_t|, \quad W_t := |I_t| + Z_t = 2|I_t| + |\partial I_t|.$$

Note that $|I_t|$, Z_t, and W_t are non-decreasing with t, whereas this is not true for $|\partial I_t|$. The next lemma says that W_t increases "on average" per round by a factor of $1 + \Omega(\beta)$, where

$$\beta := \alpha / \big(\log \Delta \sqrt{\log(2\Delta/\delta)}\big).$$

LEMMA 3.1. (MAIN LEMMA) *Fix a round $t \geq 0$ and the set I_t. If $|I_t| \leq n/2$ then there is some integer $1 \leq r \leq c_1/\beta$ depending only on I_t such that with constant probability $W_{t+r} \geq (1 + r\beta/c_2)W_t$, where $c_1, c_2 > 0$ are constants independent of I_t and G.*

Before we prove this result, we use it to derive Theorem 1.1. We start by showing the next claim, which states that more than half of the vertices get informed in the first $\mathcal{O}(\beta^{-1}\log n)$ rounds w.h.p.

CLAIM 3.2. *W.h.p. $|I_{c\beta^{-1}\log n}| > n/2$, for a sufficiently large constant c.*

Proof. Recall that Lemma 3.1 specifies some $r = r(I_t) \leq c_1/\beta$ such that with probability at least $p = \Theta(1)$, $W_{t+r} \geq (1+r\beta/c_2)W_t$. If the *same, fixed* r worked for all values of I_t then the claim would follow easily: A phase of $r = \mathcal{O}(1/\beta)$ rounds would increase W_t by a factor of $(1 + \Omega(r\beta))$ w.c.p. And since W_t cannot be increased beyond $2n$, it follows that $\mathcal{O}(\log n/(r\beta))$ such phases (i.e., $\mathcal{O}(\log n/\beta)$ rounds) would result in $|I_t| > n/2$ w.h.p.

Although Lemma 3.1 itself does not guarantee that there is an r that works for all I_t, we now show that such an r exists. Specifically, we show that in at most $2r_{\max}$ rounds, where $r_{\max} := c_1/\beta$, W_t increases by a constant factor w.c.p. Our proof uses the Optional Stopping Theorem.

Fix I_t. Let $r_1 = r(I_t)$ be the number of rounds that Lemma 3.1 specifies for I_t, and let $r_2 = r(I_{t+r_1})$, $r_3 = r(I_{t+r_1+r_2})$, and so on. I.e, for each $i = 1, 2, \ldots$, we let $r_i := r(I_{t+R_{i-1}})$, where $R_j := \sum_{k=1}^{j} r_k$. Define the random variables Q_i, for $i = 1, 2, \ldots$, by

$$Q_i = \min\left\{W_{t+R_i} - W_{t+R_{i-1}},\ \frac{r_i\beta W_t}{c_2}\right\} - \frac{pr_i\beta W_t}{c_2}.$$

Further, let $\mathcal{Q}_i = \sum_{j=1}^{i} Q_j$. Note that

$$W_{t+R_i} \geq W_t + \mathcal{Q}_i + \frac{pR_i\beta W_t}{c_2}. \tag{3.1}$$

It is straightforward to verify that the sequence of $\mathcal{Q}_i$ is a submartingale: Let $\mathcal{F}_i$ be the σ-algebra corresponding to the history of the rumor spreading process up to round $t + R_i$. Then,

$$\begin{aligned}
&\mathbf{E}[\mathcal{Q}_i - \mathcal{Q}_{i-1} \mid \mathcal{F}_{i-1}] = \mathbf{E}[Q_i \mid \mathcal{F}_{i-1}] \\
&= \mathbf{E}\left[\min\left\{W_{t+R_i} - W_{t+R_{i-1}}, \frac{r_i\beta W_t}{c_2}\right\} \,\middle|\, \mathcal{F}_i\right] - \frac{pr_i\beta W_t}{c_2} \\
&\geq \mathbf{Pr}\left(W_{t+R_i} - W_{t+R_{i-1}} \geq \frac{r_i\beta W_t}{c_2} \,\middle|\, \mathcal{F}_{i-1}\right) \cdot \frac{r_i\beta W_t}{c_2} \\
&\qquad - \frac{pr_i\beta W_t}{c_2} \geq 0,
\end{aligned}$$

since the probability in the second-to-last line is at least p by Lemma 3.1. Define now the stopping time

$$\tau = \min\{i\colon R_i \geq r_{\max}\};$$

so, $r_{\max} \leq R_\tau \leq 2r_{\max}$. By the definition of $\mathcal{Q}_i$,

$$\mathcal{Q}_\tau \leq \frac{R_\tau\beta W_t}{c_2} - \frac{pR_\tau\beta W_t}{c_2} \leq \frac{R_\tau\beta W_t}{c_2} \leq \frac{2r_{\max}\beta W_t}{c_2}.$$

Thus,

$$\begin{aligned}
\mathbf{E}[\mathcal{Q}_\tau] \leq\ & \mathbf{Pr}\left(\mathcal{Q}_\tau \geq -\frac{pr_{\max}\beta W_t}{2c_2}\right)\frac{2r_{\max}\beta W_t}{c_2} \\
&+ \mathbf{Pr}\left(\mathcal{Q}_\tau < -\frac{pr_{\max}\beta W_t}{2c_2}\right)\frac{-pr_{\max}\beta W_t}{2c_2}.
\end{aligned}$$

By the stopping time theorem, $\mathbf{E}[\mathcal{Q}_\tau] \geq \mathcal{Q}_0 = 0$. Combining this and the above inequality gives

$$\begin{aligned}
&\mathbf{Pr}\left(\mathcal{Q}_\tau \geq -\frac{pr_{\max}\beta W_t}{2c_2}\right)\frac{2r_{\max}\beta W_t}{c_2} \\
&\quad + \left(1 - \mathbf{Pr}\left(\mathcal{Q}_\tau \geq -\frac{pr_{\max}\beta W_t}{2c_2}\right)\right)\frac{-pr_{\max}\beta W_t}{2c_2} \geq 0,
\end{aligned}$$

and so,

$$\mathbf{Pr}\left(\mathcal{Q}_\tau \geq -\frac{pr_{\max}\beta W_t}{2c_2}\right) \geq \frac{p}{4+p} = \Theta(1).$$

Therefore, by (3.1), it holds with probability at least $p/(4+p)$ that

$$\begin{aligned}
W_{t+2r_{\max}} &\geq W_{t+R_\tau} \geq W_t - \frac{pr_{\max}\beta W_t}{2c_2} + \frac{pR_\tau\beta W_t}{c_2} \\
&\geq W_t + \frac{pr_{\max}\beta W_t}{2c_2} = W_t + \frac{pc_1W_t}{2c_2},
\end{aligned}$$

since $r_{\max} = c_1/\beta$. We have thus proved that in at most $2r_{\max}$ rounds W_t increases by the factor of at least $1 + pc_1/(2c_2) = 1 + \Theta(1)$ with probability at least $p/(4+p) = \Theta(1)$.

Finally, by a Chernoff bound, $\mathcal{O}(2r_{\max}\log n) = \mathcal{O}(\beta^{-1}\log n)$ rounds suffice w.h.p. to increase W_t to its maximum value, and thus to have $|I_t| > n/2$. This completes the proof of Claim 3.2. ∎

The following simple symmetry result holds for the PUSH-PULL protocol. The proof is essentially the same as that of [9, Lemma 3], and is therefore omitted.

LEMMA 3.3. *For any $V_1, V_2 \subseteq V$, let $\mathcal{E}(V_1, V_2, r)$ denote the event that a rumor known to all vertices in V_1 (and only to them) spreads to at least one vertex in V_2 in at most r rounds. Then $\mathbf{Pr}(\mathcal{E}(V_1, V_2, r)) = \mathbf{Pr}(\mathcal{E}(V_2, V_1, r))$.*

Using Claim 3.2 and Lemma 3.3, we can derive Theorem 1.1 easily: By Claim 3.2 more than half of the vertices are informed in $t = c\beta^{-1}\log n$ rounds w.h.p. Suppose now that $|I_t| > n/2$. Then by the same claim, a rumor started at a given uninformed vertex $v \in V \setminus I_t$ would spread to more than half of the vertices, and thus to at least one vertex in I_t w.h.p. Hence, by Lemma 3.3, the rumor from I_t spreads to v in t additional rounds w.h.p. The theorem follows by applying the union bound to obtain that all uninformed vertices become informed in t additional rounds w.h.p.

3.3 Results Used in the Proof of Lemma 3.1. In this section we describe four lemmata, which bound the number of rounds until some fraction of a set $B \subseteq \partial I_t$ gets informed, or until a sufficiently large number of vertices are added to the boundary as a result of vertices in B getting informed. In all these results we assume that I_t is fixed, and that B is arbitrary but fixed. For each $v \in V$, $h(v)$ denotes the number of neighbors of v in B,

$$h(v) := |N(v) \cap B|.$$

The first two lemmata follow by direct calculations. Recall that $Z_t = |I_t| + |\partial I_t|$.

LEMMA 3.4. *Suppose that for every vertex $u \in B$, $d(u) \leq k$, for some fixed k. Then, with probability at least $1 - 1/e$, at the end of round $t + \lceil \min\{Z_t, k\}/|B| \rceil$ at least one vertex in B is informed.*

Proof. For any vertex $v \in I_t$, the degree of v is $d(v) \leq Z_t$, because $N(v) \subseteq I_t \cup \partial I_t$. The probability that no vertex in I_t pushes the rumor to any vertex in B for $r_1 := \lceil Z_t/|B| \rceil$ rounds is

$$\prod_{v \in I_t} \left(1 - \frac{h(v)}{d(v)}\right)^{r_1} \leq \prod_{v \in I_t} \left(1 - \frac{h(v)}{Z_t}\right)^{r_1} \leq e^{-r_1 \sum_{v \in I_t} h(v)/Z_t} \leq e^{-r_1|B|/Z_t} \leq 1/e.$$

Similarly, the probability that no vertex in B pulls the rumor from I_t for $r_2 := \lceil k/|B| \rceil$ rounds is at most

$$\prod_{u \in B} \left(1 - \frac{1}{d(u)}\right)^{r_2} \leq \prod_{u \in B} \left(1 - \frac{1}{k}\right)^{r_2} \leq e^{-r_2 \sum_{u \in B} 1/k} = e^{-r_2|B|/k} \leq 1/e.$$

Combining the above two results yields the probability that no vertex in B gets informed in $\min\{r_1, r_2\}$ rounds is at most $1/e$. ∎

LEMMA 3.5. *Let $0 < q \leq 1$. Suppose that for every vertex $u \in B$, at least a q-fraction of u's neighbors is in I_t, i.e., $|N(u) \cap I_t| \geq q \cdot d(u)$. Then, with probability at least $1/2$, at the end of round $t + \lceil 2/q \rceil$ at least half of the vertices in B have pulled the rumor from I_t.*

Proof. For each $u \in B$, let X_u be the 0/1 random variable with $X_u = 1$ if and only if u pulls the rumor from some vertex in I_t by the end of round $t + \lceil 2/q \rceil$. Let also $X = \sum_{u \in B} X_u$. Then,

$$\mathbf{Pr}(X_u = 1) \geq 1 - (1-q)^{\lceil 2/q \rceil} \geq 1 - 1/e^2 =: p,$$

and $\mathbf{E}[X] \geq p|B|$. By applying Chernoff bounds we obtain

$$\mathbf{Pr}(X < |B|/2) = \mathbf{Pr}(\mathbf{E}[X] - X > \mathbf{E}[X] - |B|/2) \leq e^{-2(\mathbf{E}[X]-|B|/2)^2/|B|} \leq e^{-2(p-1/2)^2|B|}.$$

If $|B| \geq 3$, the above yields

$$\mathbf{Pr}(X < |B|/2) \leq e^{-2(1-1/e^2-1/2)^2 3} < 1/2.$$

If $|B| = 1, 2$ then

$$\mathbf{Pr}(X < |B|/2) = \mathbf{Pr}(X = 0) \leq \mathbf{Pr}(X_u = 0) \leq 1/e^2 < 1/2.$$

This completes the proof. ∎

The next lemma lower-bounds the number of vertices added to the boundary, when a constant fraction of the vertices from B go to vertices in $S_t := V \setminus (I_t \cup \partial I_t)$ that have not too many neighbors in B. The proof is more involved than of the previous two lemmata, because it must deal with dependencies: For a vertex $v \in S_t$, the event that at least one of its neighbors in B gets informed is not independent of the corresponding events for vertices in S_t that have common neighbors with v in B. We tackle these dependencies by using a version of the Method of Bounded Independent Differences.

LEMMA 3.6. *Suppose that for every vertex $u \in B$, $k \leq d(u) < 2k$, for some fixed k. Suppose also that*

$$\sum_{v \in S_t:\, h(v) \leq k} h(v) \geq \frac{1}{2}\operatorname{vol}(B). \tag{3.2}$$

Then, with probability at least $1 - 1/e$, at the end of round $t + \lceil 14k/|B| \rceil$ at least $|B|/801$ vertices in S_t have a neighbor in B that has pulled the rumor from I_t.

Proof. We distinguish two cases, depending on the relative values of $|B|$ and k. First, we consider the case where k is sufficiently smaller than B; precisely, the case of $k \le |B|/267$.

For each $u \in B$, the probability that u pulls the rumor from I_t in a given round is

$$p_u := |N(u) \cap I_t|/d(u) \ge 1/d(u) \ge 1/(2k) =: p.$$

We make the pessimistic assumption that for every u this probability is *exactly* p. (Formally, we modify the rumor-spreading process such that in each of the rounds $t+1, t+2, \dots,$ each vertex u decides *not* to send a pull request independently with probability $1 - p/p_u$; with the remaining probability p/p_u, u sends the pull request to a random neighbor, as normally. Clearly this modification can only make our bound worse.)

For each $u \in B$, let X_u be the 0/1 random variable with $X_u = 1$ if and only if u pulls the rumor from I_t in round $t+1$ (assuming the modification above). Thus,

$$\mathbf{Pr}(X_u = 1) = p.$$

Define the set

$$D := \{v \in S_t : h(v) \le k\}.$$

For each $v \in D$, let Y_v be the 0/1 random variable with $Y_v = 1$ if and only if $X_u = 1$ for at least one neighbor $u \in B$ of v. Further, let $Y := \sum_{v \in D} Y_v$. We will prove that

$$\mathbf{Pr}(Y \ge |B|/801) \ge 1 - 1/e, \tag{3.3}$$

which implies the lemma.

We have

$$\mathbf{Pr}(Y_v = 1) = 1-(1-p)^{h(v)} \ge 1-e^{-ph(v)} \ge (3/4)ph(v),$$

where the last inequality was obtained using the facts that $ph(v) \le pk = 1/2$, and that $e^{-x} \le 1 - 3x/4$ when $0 \le x \le 1/2$. The expected number of vertices $v \in D$ for which $Y_v = 1$ is then

$$\begin{aligned}\mathbf{E}[Y] &= \sum_{v \in D} \mathbf{Pr}(Y_v = 1) \ge \frac{3p}{4} \sum_{v \in D} h(v) \\ &\ge \frac{3p}{4} \cdot \frac{\mathrm{vol}(B)}{2} \ge \frac{3p}{4} \cdot \frac{k|B|}{2} = \frac{3|B|}{16},\end{aligned}$$

where for the second inequality we used the lemma's assumption that $\sum_{v \in D} h(v) \ge \mathrm{vol}(B)/2$.

Next we prove that Y is concentrated around its expectation. Since Y is a function of the independent (binary) random variables $\{X_u : u \in B\}$, we employ the Method of Bounded Differences. In particular, we use the next result, which follows directly from [31, Theorem 3.9].

Theorem 3.1. (Bounded Differences Inequality) *Let $R_1, \dots, R_n$ be independent 0/1 random variables with $\mathbf{Pr}(R_i = 1) \le p \le 1/2$. Let also f be a bounded real function defined on $\{0,1\}^n$. Define $\mu := \mathbf{E}[f(R_1, \dots, R_n)]$, and $b := \max |f(\mathbf{x}) - f(\mathbf{x}')|$, where the maximum is over all $\mathbf{x}, \mathbf{x}' \in \{0,1\}^n$ that differ only in one position. Then, for any $\lambda > 0$,*

$$\mathbf{Pr}(\mu - f(R_1, \dots, R_n) \ge \lambda) \le \exp\left(-\frac{\lambda^2}{2pnb^2 + 2b\lambda/3}\right).$$

By applying this result to $Y = f(\{X_u : u \in B\})$, in which case we have $\mu = \mathbf{E}[Y]$ and $b \le 2k$, we obtain

$$\begin{aligned}\mathbf{Pr}(\mathbf{E}[Y] - Y \ge \lambda) &\le \exp\left(-\frac{\lambda^2}{2p|B|(2k)^2 + 2(2k)\lambda/3}\right) \\ &= \exp\left(-\frac{\lambda^2}{4k(|B| + \lambda/3)}\right).\end{aligned}$$

And setting

$$\lambda = \mathbf{E}[Y] - \frac{|B|}{16} \ge \frac{|B|}{8},$$

yields

$$\begin{aligned}\mathbf{Pr}(Y \le |B|/16) &\le \exp\left(-\frac{(|B|/8)^2}{4k(|B| + |B|/24)}\right) \\ &\le \exp\left(-\frac{|B|}{267k}\right) \le 1/e,\end{aligned}$$

for $k \le |B|/267$. Thus, Inequality (3.3) holds when $k \le |B|/267$.

It remains to consider the case where $k > |B|/267$. Since now $k = \Omega(|B|)$, we will just compute the probability that (at least) *one* vertex $u \in B$ with a constant fraction of its neighbors in S_t pulls the rumor from I_t in $\lceil k/|B| \rceil$ rounds.

We first show that at least $|B|/7$ of the vertices in B have at least $k/3$ neighbors in S_t each. Let $\hat{B} := \{u \in B \colon |N(u) \cap S_t| \ge k/3\}$ be the set of vertices in B that have at least $k/3$ neighbors in S_t, and let $m := |\hat{B}|$. Let also $\hat{d} := (1/m) \sum_{u \in \hat{B}} d(u)$ be the average degree of the vertices in $\hat{B}$. Then,

$$\begin{aligned}\sum_{v \in S_t} h(v) &= \sum_{u \in B} |N(u) \cap S_t| \\ &\le \sum_{u \in \hat{B}} d(u) + \sum_{u \in B \setminus \hat{B}} |N(u) \cap S_t| \\ &\le m\hat{d} + (|B| - m)k/3.\end{aligned}$$

Also,

$$\sum_{v \in S_t} h(v) \ge \sum_{v \in D} h(v) \ge \mathrm{vol}(B)/2,$$

by the lemma's assumptions, and

$$\operatorname{vol}(B) \geq m\hat{d} + (|B| - m)k.$$

Combining the three inequalities above yields

$$m \geq \frac{k|B|}{3\hat{d} + k} \geq |B|/7,$$

since $\hat{d} \leq 2k$.

Therefore, the probability that at least one of these m vertices in $\hat{B}$ pulls the rumor from I_t in $\lceil 14k/|B| \rceil$ rounds is

$$1 - (1-p)^{m\lceil 14k/|B| \rceil} \geq 1 - e^{-pm\lceil 14k/|B| \rceil} \geq 1 - 1/e.$$

Thus, at least with this probability, the rumor is pulled from I_t to some $u \in B$ with at least $k/3 > (|B|/267)/3 = |B|/801$ neighbors in S_t. This completes the proof of the case $k > |B|/267$, and the proof of Lemma 3.6. ∎

The last lemma we describe gives a lower bound on the fraction of B that gets informed when a constant fraction of the vertices from B goes to vertices in V with not too few neighbors in B.

Lemma 3.7. *Let $k \geq \ell \geq 256$. Suppose that for every vertex $u \in B$, $k \leq d(u) < 2k$. Suppose also that*

$$\sum_{v \in V\,:\, h(v) \geq \ell} h(v) \geq \frac{1}{4}\operatorname{vol}(B).$$

Then, with probability at least $1 - 1/e$, at the end of round $t + 512(\sqrt{k/\ell}\log \ell + k/\ell)$ at least a $(1/40)$-fraction of the vertices in B is informed. This is true even if we only take into account transmissions of the rumor initiated by vertices in B, i.e., only vertices in B push or pull the rumor.

The proof of this result is fairly involved and is described in the next section.

3.4 Proof of Lemma 3.7.

The proof is divided into three steps. In the first step we show that there exist two sets of vertices $B' \subseteq B$ and $V' \subseteq V$, such that each vertex in B' has $\Omega(k)$ neighbors in V' and each vertex in V' has $\Omega(\ell)$ neighbors in B', and also $|B'| = \Omega(|B|)$. In the second step, we show that for any vertex $u \in B'$, a (new) rumor started at u spreads to at least one vertex in I_t within $\mathcal{O}(\sqrt{k/\ell}\log \ell + k/\ell)$ rounds w.c.p. A key claim is that the rumor from u spreads to $\Omega(\ell)$ other vertices in B' within $\mathcal{O}(\sqrt{k/\ell}\log \ell)$ rounds. In the third step, we combine the result of the second step with the symmetry Lemma 3.3 to obtain that the rumor spreads from I_t to a given vertex $u \in B'$ within $\mathcal{O}(\sqrt{k/\ell}\log \ell + k/\ell)$ rounds w.c.p. Thus, by Markov's inequality, a constant fraction of B' learns the rumor from I_t in $\mathcal{O}(\sqrt{k/\ell}\log \ell + k/\ell)$ rounds w.c.p. The lemma then follows by applying the result of the first step that $|B'| = \Omega(|B|)$.

Step 1: The Sets B' and V'. We describe a procedure for constructing two (possibly overlapping) sets of vertices $B' \subseteq B$ and $V' \subseteq V$, such that each vertex in B' has at least $k/8$ neighbors in V', each vertex in V' has at least $\ell/8$ neighbors in B', and the size of B' is $|B'| \geq (1/20)|B|$.

The procedure that constructs B' and V' resembles the Bootstrap Percolation process on G [1]: Two variables $\tilde{B}$ and $\tilde{V}$ are used. Initially, we set $\tilde{B} \leftarrow B$ and $\tilde{V} \leftarrow V$. In each iteration of the process, either we remove from $\tilde{B}$ a vertex u with $|N(u) \cap \tilde{V}| < k/8$, or we remove from $\tilde{V}$ a vertex v with $|N(v) \cap \tilde{B}| < \ell/8$. The procedure stops when no more vertices can be removed from $\tilde{B}$ and $\tilde{V}$, and we let B' and V' be the final sets $\tilde{B}$ and $\tilde{V}$, respectively. Note that a vertex $u \in B$ is initially both in $\tilde{B}$ and $\tilde{V}$, and as the procedure progresses u may be removed from one or both of these sets.

By construction, for every $u \in B'$, $|N(u) \cap V'| \geq k/8$, and for every $v \in V'$, $|N(v) \cap B'| \geq \ell/8$. Next we prove that $|B'| \geq (1/20)|B|$. We will use the following result, which says that $\operatorname{vol}(B', V') := \sum_{v \in B'} |N(v) \cap V'|$ is close to $\operatorname{vol}(B) = \operatorname{vol}(B, V)$.[5]

Claim 3.8. $\operatorname{vol}(B', V') \geq (1/10)\operatorname{vol}(B)$.

Proof. We use the potential method of amortized analysis. Let B_i and V_i denote the current sets $\tilde{B}$ and $\tilde{V}$ in the construction procedure after the i-th iteration, i.e., after i vertices have been removed in total from B and V. ($B_0 = B$ and $V_0 = V$.) We will define the potential Φ_i of the pair B_i and V_i in such a way that the difference $\Phi_{i-1} - \Phi_i$ is an upper bound on the corresponding reduction in the number of edges because of the i-th vertex removal, i.e.,

$$\operatorname{vol}(B_{i-1}, V_{i-1}) - \operatorname{vol}(B_i, V_i) \leq \Phi_{i-1} - \Phi_i. \tag{3.4}$$

Thus, $\operatorname{vol}(B) - \operatorname{vol}(B', V') \leq \sum_{i \geq 0}(\Phi_{i-1} - \Phi_i) \leq \Phi_0$. Further, we will show that

$$\Phi_0 \leq (9/10)\operatorname{vol}(B). \tag{3.5}$$

The last two inequalities then yield the claim.

[5] Note that $\operatorname{vol}(V_1, V_2)$ may be different than $|E(V_1, V_2)|$. Specifically, $\operatorname{vol}(V_1, V_2) = |E(V_1, V_2)| + |E(V_1 \cap V_2, V_1 \cap V_2)|$. Note also that $\operatorname{vol}(V_1, V_2) = \operatorname{vol}(V_2, V_1)$.

The potential Φ_i after the i-th iteration is defined as follows. Recall that $h(v) := |N(v) \cap B|$.

$$\Phi_i := \sum_{u \in B_i} \frac{1}{7}|N(u) \setminus V_i| + \sum_{v \in V_i} \frac{8}{7} \min\{h(v), \ell/8\}.$$

We must prove that inequalities (3.4) and (3.5) hold. We prove (3.4) first. Consider an iteration i in which we remove a vertex u from B_{i-1}. Observe that the second sum in the definition of Φ does not change. So,

$$\begin{aligned}
&\Phi_{i-1} - \Phi_i \\
&\quad = \sum_{u' \in B_{i-1}} \frac{1}{7}|N(u') \setminus V_{i-1}| - \sum_{u' \in B_i} \frac{1}{7}|N(u') \setminus V_i| \\
&\quad = \frac{1}{7}|N(u) \setminus V_{i-1}| \geq \frac{1}{7}(d(u) - k/8) \geq \frac{k}{8},
\end{aligned}$$

where the second-to-last inequality holds because $|N(u) \cap V_{i-1}| \leq k/8$ as u is removed in this iteration, and the last inequality holds because $d(u) \geq k$. And since $\mathrm{vol}(B_{i-1}, V_{i-1}) - \mathrm{vol}(B_i, V_i) = |N(u) \cap V_{i-1}| \leq k/8$, inequality (3.4) holds for this case. Consider now an iteration i in which we remove a vertex v from V_{i-1}. The change in the first sum in the definition of Φ is

$$\begin{aligned}
&\sum_{u' \in B_{i-1}} \frac{1}{7}|N(u') \setminus V_{i-1}| - \sum_{u' \in B_{i-1}} \frac{1}{7}|N(u') \setminus V_i| \\
&\qquad = -\frac{1}{7}|N(v) \cap B_{i-1}|.
\end{aligned}$$

The change in the second sum in Φ's definition is

$$\begin{aligned}
&\sum_{v' \in V_{i-1}} \frac{8}{7} \min\{h(v'), \ell/8\} - \sum_{v' \in V_i} \frac{8}{7} \min\{h(v'), \ell/8\} \\
&\qquad = \frac{8}{7} \min\{h(v), \ell/8\} \geq \frac{8}{7}|N(v) \cap B_{i-1}|,
\end{aligned}$$

where the last inequality holds because $|N(v) \cap B_{i-1}| \leq |N(v) \cap B_0| = h(v)$, and also $|N(v) \cap B_{i-1}| \leq \ell/8$ since v is removed in this iteration. Therefore, $\Phi_{i-1} - \Phi_i \geq |N(v) \cap B_{i-1}|$. And since $\mathrm{vol}(B_{i-1}, V_{i-1}) - \mathrm{vol}(B_i, V_i) = |N(v) \cap B_{i-1}|$, inequality (3.4) follows.

We now prove (3.5). We use the assumption of Lemma 3.7 that $\sum_{v \in V \colon h(v) \geq \ell} h(v) \geq (1/4)\,\mathrm{vol}(B)$.

$$\begin{aligned}
\Phi_0 &= \sum_{v \in V} \frac{8}{7} \min\{h(v), \ell/8\} \\
&\leq \frac{8}{7} \sum_{v \in V \colon h(v) < \ell} h(v) + \frac{1}{7} \sum_{v \in V \colon h(v) \geq \ell} \ell \\
&\leq \frac{8}{7} \sum_{v \in V \colon h(v) < \ell} h(v) + \frac{1}{7} \sum_{v \in V \colon h(v) \geq \ell} h(v) \\
&= \frac{8}{7}\Big(\mathrm{vol}(B) - \sum_{v \in V \colon h(v) \geq \ell} h(v)\Big) + \frac{1}{7} \sum_{v \in V \colon h(v) \geq \ell} h(v) \\
&= \frac{8}{7}\,\mathrm{vol}(B) - \sum_{v \in V \colon h(v) \geq \ell} h(v) \\
&\leq \frac{8}{7}\,\mathrm{vol}(B) - \frac{1}{4}\,\mathrm{vol}(B) \leq \frac{9}{10}\,\mathrm{vol}(B).
\end{aligned}$$

This completes the proof of Claim 3.8. ∎

From this claim, and the assumption that $k \leq d(u) < 2k$ for all $u \in B$, it follows

$$\begin{aligned}
(3.6) \quad |B'| &\geq \frac{\mathrm{vol}(B')}{2k} \geq \frac{\mathrm{vol}(B', V')}{2k} \geq \frac{(1/10)\,\mathrm{vol}(B)}{2k} \\
&\geq \frac{(1/10)k|B|}{2k} = |B|/20.
\end{aligned}$$

Step 2: Spreading a Rumor from B' to I_t. Suppose now that a rumor starts at a given vertex in B'. The next claim states that the number of rounds until this rumor spreads to $\Theta(\ell)$ vertices in B' is $\mathcal{O}(\sqrt{k/\ell}\log \ell)$ w.c.p.

CLAIM 3.9. *For any vertex $s \in B'$, a rumor started at s spreads to at least $\ell/32$ vertices in B' in $512\sqrt{k/\ell}\log \ell$ rounds with probability at least $1 - e^4$. This is true even if we only take into account transmissions of the rumor initiated by vertices in B'.*

Proof. We divide the rumor spreading process into phases of $2r = \Theta(\sqrt{k/\ell})$ rounds each, and we show that in each phase the number of informed vertices in B' doubles w.c.p. Thus, w.c.p. it takes $\mathcal{O}(\log \ell)$ phases until $\Omega(\ell)$ vertices in B' get informed. We assume that only vertices in B' initiate push or pull operations. Formally, a phase consists of

$$r := 256\sqrt{k/\ell}$$

push rounds, during which the informed vertices in B' perform push operations, followed by r *pull rounds*, in which the uninformed vertices in B' perform pull operations. The assumption that there are no pull (push) operations during the push (pull) rounds can only strengthen our result. Further, we assume that $B' \cap V' = \emptyset$, which also makes the result stronger.

We now analyze a single phase. The analysis consist of two steps. Let m be the number of informed vertices in B' at the beginning of the phase. First we show that w.c.p. the total number of vertices in V' that get informed during the r push rounds is at least $\Omega(mr)$. Then we show that if there are at least that many informed vertices in V', then w.c.p. the r pull rounds increase the number of informed vertices in B' by $\Omega(mr^2(\ell/k)) = \Omega(m)$.

Suppose that at the beginning of the phase there are

$$1 \leq m \leq \ell/32$$

informed vertices in B'. W.l.o.g., we assume that the push operations in each round are performed sequentially (in an arbitrary order). In total, rm push operations are performed in a phase. For each $i = 1, \dots, rm$, we denote by V'_i the set of vertices $v \in V'$ that are informed by the first i push operations. Further, we let X_i be the 0/1 random variable with $X_i = 1$ if and only if some of the following two conditions is satisfied:

(i) the i-th push operation informs a vertex in $V' \backslash V'_{i-1}$;

(ii) $|V'_{i-1}| \geq 2k/r$ (i.e., $|V'_{i-1}|$ is already sufficiently large).

We now show that for each $i = 1, \dots, rm$,

$$\mathbf{Pr}(X_i = 1 \mid X_1, \dots, X_{i-1}) \geq 1/32. \tag{3.7}$$

Clearly, if Condition (ii) holds then the above probability equals 1. So, suppose that (ii) does not hold, i.e., $|V_{i-1}| < 2k/r$. Then the vertex $u \in B'$ that performs the i-th push operation has a number of uninformed neighbors in V' that is at least $|N(u) \cap V'| - |V_{i-1}| \geq k/8 - 2k/r \geq k/16$. And since $d(u) < 2k$, the probability that u informs one of those neighbors is at least $(k/16)/(2k) = 1/32$.

Because of (3.7), a simple coupling argument yields that the sum $X := \sum_{i=1}^{rm} X_i$ dominates stochastically the binomial random variable $B(rm, 1/32)$. Thus, by a Chernoff bound,

$$\begin{aligned}\mathbf{Pr}(X \leq rm/64) &\leq \mathbf{Pr}(B(rm, 1/32) \leq rm/64) \\ &\leq e^{-(1/2)^2 (rm/32)/2} = e^{-rm/256} \leq 1/e,\end{aligned}$$

since $r \geq 256$. Therefore, with probability at least $1 - 1/e$, $X \geq rm/64$, and thus,

$$|V'_{rm}| \geq \min\left\{\frac{rm}{64},\ \frac{2k}{r}\right\}.$$

Next we consider the pull rounds. Fix the outcome of the push rounds such that the above inequality holds. In fact, we make the pessimistic assumption that it holds as equality, i.e., $|V'_{rm}| = \min\{rm/64,\ 2k/r\}$. For each $u \in B'$, let Y_u denote the 0/1 random variable with $Y_u = 1$ if and only if u pulls the rumor from V'_{rm} in some of the pull rounds. Then,

$$\begin{aligned}\mathbf{Pr}(Y_u = 1) &= 1 - \left(1 - \frac{|N(u) \cap V'_{rm}|}{d(u)}\right)^r \\ &\geq 1 - \left(1 - \frac{|N(u) \cap V'_{rm}|}{2k}\right)^r \\ &\geq \frac{r|N(u) \cap V'_{rm}|}{4k},\end{aligned}$$

where the last inequality was obtained by using the facts that

$$\frac{r|N(u) \cap V'_{rm}|}{2k} \leq \frac{r|V'_{rm}|}{2k} \leq \frac{r(2k/r)}{2k} = 1,$$

and that $e^{-x} \leq 1 - x/2$ when $0 \leq x \leq 1$. The random variables Y_u, $u \in B'$, are independent, and the expectation of their sum $Y := \sum_{u \in B'} Y_u$ is

$$\begin{aligned}\mathbf{E}[Y] &= \sum_{u \in B'} \mathbf{Pr}(Y_u = 1) \geq \sum_{u \in B'} \frac{r|N(u) \cap V'_{rm}|}{4k} \\ &= \frac{r}{4k} \sum_{v \in V'_{rm}} |N(v) \cap B'| \geq \frac{r}{4k} \sum_{v \in V'_{rm}} \frac{\ell}{8} \\ &= \frac{r\ell|V'_{rm}|}{32k} = \frac{r\ell}{32k} \min\left\{\frac{rm}{64},\ \frac{2k}{r}\right\} \\ &= \min\left\{\frac{r^2 \ell m}{2048k},\ \frac{\ell}{16}\right\} \geq \min\{32m,\ \ell/16\},\end{aligned}$$

by the definition of r. A Chernoff bound yields $\mathbf{Pr}(Y \leq \mathbf{E}[Y]/2) \leq e^{-\mathbf{E}[Y]/8}$. Therefore, with probability at least

$$1 - e^{-\mathbf{E}[Y]/8} \geq 1 - e^{-\min\{4m,\ \ell/128\}} \geq 1 - 1/e^2$$

(since $\ell \geq 256$), the number of informed vertices in B' after the pull rounds is at least

$$\mathbf{E}[Y]/2 \geq \min\{16m,\ \ell/32\};$$

that is, at least 16 times the initial number of informed vertices in B', or at least $\ell/32$. Note that the above probability is conditional on the event that $|V'_{rm}| \geq \min\{rm/64,\ 2k/r)\}$, which has probability at least $1 - 1/e$ as we showed earlier. Therefore, the unconditional probability of the event that the number of informed vertices in B' at the end of the phase is at least $\min\{16m,\ \ell/32\}$ is at least

$$(1 - 1/e)(1 - 1/e^2) \geq 1/2.$$

We say that the phase is *successful* if this event occurs.

Since at most $\lceil \log_{16}(\ell/34) \rceil \leq (\log \ell)/4$ successful phases are needed to inform $\ell/32$ vertices in B', and since each phase is successful with probability at least $1/2$, a Chernoff bound yields that in $\log \ell$ phases (that is, in $2r \log \ell$ rounds) the rumor spreads to at least $\ell/32$ vertices $u \in B'$ with probability at least $1 - 1/e^4$. This completes the proof of Claim 3.9. ∎

Once the rumor started at $s \in B'$ has spread to $\ell/32$ vertices in B', the probability that at least one of those $\ell/32$ vertices pushes the rumor to some vertex in I_t in the next $512k/\ell$ rounds is at least

$$1 - (1 - (1/2k))^{\frac{\ell}{32} \cdot \frac{512k}{\ell}} \geq 1 - e^{\frac{1}{2k} \cdot \frac{\ell}{32} \cdot \frac{512k}{\ell}} \geq 1 - 1/e^4.$$

Therefore, we have the following result.

Corollary 3.1. *A rumor started at some vertex $s \in B'$ spreads to at least one vertex in I_t in $512(\sqrt{k/\ell} \log \ell + k/\ell)$ rounds with probability at least $1 - 2/e^4$. This is true even if we only take into account transmissions of the rumor initiated by vertices in B'.*

Step 3: Finishing the Proof of Lemma 3.7. We observe that the symmetry Lemma 3.3 is still true even if we only take into account transmissions of the rumor initiated by a fixed subset of V. Now, combining this lemma and Corollary 3.1, we obtain that the rumor spreads from I_t to any given vertex $u \in B'$ in at most $r := 512(\sqrt{k/\ell} \log \ell + k/\ell)$ rounds with probability at least $1-2/e^4$. Thus, the expected number of uninformed vertices in B' after r rounds is at most $(2/e^4)|B'|$. Then, by Markov's inequality, the probability that more than half of the vertices in B' are uninformed is at most $(2/e^4)|B'|/(|B'|/2) \leq 1/e^2$. And since by (3.6), $|B'| \geq |B|/20$, it follows that with probability at least $1-1/e^2$, at least $|B|/40$ vertices in B get informed in the next r rounds. This completes the proof of Lemma 3.7.

3.5 Proof of Lemma 3.1.

We partition the boundary ∂I_t into sets A_i, for $i \in \{\lceil \log(\delta + 1) \rceil, \ldots, \lceil \log(\Delta + 1) \rceil\}$, where A_i is the subset of vertices with degree between 2^{i-1} and $2^i - 1$, i.e.,

$$A_i := \{u \in \partial I_t \colon d_i \leq d(u) < 2d_i\}, \quad d_i := 2^{i-1}.$$

The total number of these sets is

$$\begin{aligned} &\lceil \log(\Delta + 1) \rceil - \lceil \log(\delta + 1) \rceil + 1 \\ &\quad < \log(\Delta + 1) + 1 - \log(\delta + 1) + 1 \\ &\quad = \log \frac{4(\Delta + 1)}{\delta + 1} \leq \log \frac{4\Delta}{\delta} =: \rho. \end{aligned}$$

Note that some of these sets may be empty.

We will reason about each set A_i independently. We only consider those A_i that are not too small, and the degree d_i is in a certain range depending on $|A_i|$. More precisely, let $\mathcal{I}$ be the set of all indices i that satisfy the condition

$$|A_i| \geq \frac{|\partial I_t|}{4\rho} \quad \wedge \quad \Big(d_i \leq 16|A_i| \ \vee \ d_i \geq 2|\partial I_t|\Big).$$

We only look at the sets A_i for $i \in \mathcal{I}$. By neglecting the remaining A_i we neglect no more than half of the vertices in ∂I_t: The worst case with respect to the first part of the condition above is that all (but one) of the A_i have size $|A_i| = \lceil |\partial I_t|/(4\rho) \rceil - 1 \leq |\partial I_t|/(4\rho)$; and the worst case with respect to the second part of the condition is that for all $i < \log(2|\partial I_t|)$, $|A_i| = \lceil d_i/16 \rceil - 1 < d_i/16$. Thus,

(3.8)
$$\begin{aligned} &\Big|\bigcup_{i \notin \mathcal{I}} A_i\Big| = \bigcup_{i \notin \mathcal{I}} |A_i| \\ &\leq \frac{|\partial I_t|}{4\rho} \cdot \rho + \frac{1}{16} \left(2|\partial I_t| + |\partial I_t| + \frac{|\partial I_t|}{2} + \frac{|\partial I_t|}{4} + \cdots \right) \\ &\leq |\partial I_t|/2. \end{aligned}$$

Roughly speaking, for each of the sets A_i, $i \in \mathcal{I}$, either we bound the number of rounds until a large fraction of the vertices in A_i gets informed, or we bound the number of rounds until the size of the boundary ∂I_t increases by a certain factor as a result of vertices in A_i becoming informed. In some cases it suffices to just study a single set A_i to prove the lemma, while in others we have to combine the results for all sets. We note that the result for each A_i is obtained by taking into account only transmissions of the rumor initiated by vertices in A_i. Hence, the results for different sets hold independently.

Fix now the set A_i, for some $i \in \mathcal{I}$. We distinguish the following cases.

1. $d_i \leq d_{\min} := 96 \cdot 256 \sqrt{\rho} \log \Delta$.

   In this case a constant fraction of A_i pulls the rumor from I_t in $\mathcal{O}(\sqrt{\rho} \log \Delta)$ rounds w.c.p.:

   For each vertex $u \in A_i \subseteq \partial I_t$, the fraction of u's neighbors in I_t is at least $1/d(u) \geq 1/(2d_{\min})$. Lemma 3.5 then implies that with probability at least $1/2$, half of the vertices in A_i pull the rumor from I_t in at most $4d_{\min}$ rounds. Let $\mathcal{A}_i(q, r)$ denote the event that at least a q-fraction of the vertices in A_i gets informed before the end of round $t + r$, if we only take into account transmissions of the rumor initiated by vertices in A_i. Using this notation we have

$$\mathbf{Pr}\left(\mathcal{A}_i(1/2,\, 4d_{\min})\right) \geq 1/2.$$

2. $d_i \geq 2|\partial I_t|$.

In this case, at least half of the neighbors of each vertex $u \in A_i$ are in $S_t \cup I_t$. We distinguish two subcases: If most vertices have more neighbors in S_t than in I_t, then when some vertex in A_i receives the rumor from I_t, which by Lemma 3.4 happens in $\mathcal{O}(\min\{d_i, Z_t\}/|A_i|)$ rounds w.c.p., the boundary increases by $\Omega(d_i)$. If, on the other hand, most vertices have more neighbors in I_t than in S_t, then by Lemma 3.5 a constant fraction of A_i pulls the rumor from I_t in $\mathcal{O}(1)$ rounds w.c.p. In both cases the lemma follows by considering just this sets A_i. We now provide the details.

Let

$$B := \{u \in A_i \colon |N(u) \cap S_t| \geq |N(u) \cap I_t|\}.$$

We have the following cases.

(a) $|B| \geq |A_i|/2$.

Since each $u \in B$ has at most $|\partial I_t|$ neighbors in ∂I_t, the number of neighbors that u has in S_t is at least $(d(u) - |\partial I_t|)/2 \geq (d_i - |\partial I_t|)/2$. By Lemma 3.4, we have that at least one vertex gets informed in B within

$$r := \min\{Z_t,\, 2d_i\}/|B|$$

rounds with probability at least $1-1/e$ (recall that $Z_t = |I_t| + |\partial I_t|$). Thus, with this probability,

$$Z_{t+r} \geq Z_t + (d_i - |\partial I_t|)/2.$$

To prove Lemma 3.1 we must show that $r = \mathcal{O}(1/\beta) = \mathcal{O}(\alpha^{-1}\sqrt{\rho}\log\Delta)$, and that the above bound on Z_{t+r} yields $W_{t+r} - W_t = \Omega(r\beta W_t) = \Omega(r\alpha W_t/(\sqrt{\rho}\log\Delta))$. We obtain the first result as follows. By the definition of r,

$$\begin{aligned} r &\leq \frac{Z_t}{|B|} \leq \frac{2Z_t}{|A_i|} \leq \frac{2Z_t}{|\partial I_t|/(4\rho)} = \frac{2(|I_t| + |\partial I_t|)}{|\partial I_t|/(4\rho)} \\ &= 8(|I_t|/|\partial I_t| + 1)\rho \leq 8(\alpha^{-1}+1)\rho \\ &\leq 16\alpha^{-1}\rho \leq 48\alpha^{-1}\log\Delta, \end{aligned}$$

since

$$\rho = \log(4\Delta/\delta) \leq \log(4\Delta) \leq \log(\Delta) + 2 \leq 3\log\Delta. \tag{3.9}$$

It remains to prove that $Z_{t+r} \geq Z_t + (d_i - |\partial I_t|)/2$ yields the desired bound on $W_{t+r} - W_t$.

$$\begin{aligned} W_{t+r} - W_t &= (Z_{t+r} - Z_t) + (|I_{t+r}| - |I_t|) \\ &\geq Z_{t+r} - Z_t \geq (d_i - |\partial I_t|)/2 \geq d_i/4, \end{aligned}$$

since $d_i \geq 2|\partial I_t|$. By the definition of r, $r \leq 2d_i/|B|$ and thus $d_i \geq r|B|/2$, so,

$$\begin{aligned} W_{t+r} - W_t &\geq \frac{r|B|}{8} \geq \frac{r|A_i|}{16} \geq \frac{r(|\partial I_t|/4\rho)}{16} \\ &= \frac{r|\partial I_t|}{64\rho} \geq \frac{r|\partial I_t|}{3 \cdot 64\log\Delta}. \end{aligned}$$

Since $W_t = 2|I_t| + |\partial I_t| = |\partial I_t|(2|I_t|/|\partial I_t| + 1) \leq |\partial I_t|(2/\alpha + 1) \leq 3|\partial I_t|/\alpha$ and thus

$$|\partial I_t| \geq \alpha W_t/3, \tag{3.10}$$

we have

$$W_{t+r} - W_t \geq \frac{r\alpha W_t}{9 \cdot 64\log\Delta}.$$

This completes the proof of Lemma 3.1 for this case.

(b) $|B| < |A_i|/2$.

Then $|A_i \setminus B| > |A_i|/2$, and for each vertex $u \in A_i \setminus B$, the number of neighbors that u has in I_t is at least $(d(u) - |\partial I_t|)/2 \geq d(u)/4$. By Lemma 3.5, half of the vertices in $A_i \setminus B$ and thus at least 1/4-th of the vertices in A_i pull the rumor from I_t in at most 8 rounds with probability at least 1/2. Therefore, with probability at least 1/2,

$$\begin{aligned} W_{t+8} - W_t &= (Z_{t+8} - Z_t) + (|I_{t+8}| - |I_t|) \\ &\geq |I_{t+8}| - |I_t| \geq |A_i|/4 \\ &\geq |\partial I_t|/(16\rho) \\ &\overset{(3.9),(3.10)}{\geq} \alpha W_t/(9 \cdot 16\log\Delta). \end{aligned}$$

This proves the lemma.

3. $d_{\min} < d_i \leq 16|A_i|$,

We further distinguish three cases: The first case is when a constant fraction of the edges from A_i goes to vertices $v \in S_t$ that have at most $\mathcal{O}(d_i)$ neighbors in A_i. Then by Lemma 3.6, $\Omega(|A_i|)$ vertices from S_t are added to the boundary in $\mathcal{O}(1)$ rounds w.c.p. The second case is when a constant fraction of the edges from A_i goes to vertices with $\Omega(d_i/\rho)$ neighbors in A_i. Then by applying Lemma 3.7, we obtain that a constant fraction of A_i gets informed in $\mathcal{O}(\sqrt{\rho}\log\Delta)$ rounds w.c.p. If none of the two cases above applies, we observe that a constant fraction of the edges from A_i goes to I_t. Thus by Lemma 3.5 a constant fraction of A_i pulls the rumor from I_t in $\mathcal{O}(1)$ rounds w.c.p. In the first and the last cases the lemma follows by considering just this set A_i. We now present the detailed proof.

For each vertex $v \in V$, let $h_i(v)$ denote the number of neighbors of v in A_i, i.e.,

$$h_i(v) := |N(v) \cap A_i|.$$

We have the following cases.

(a) It holds that

$$\sum_{v \in S_t \colon h_i(v) \le d_i} h_i(v) \ge (1/2)\,\mathrm{vol}(A_i). \tag{3.11}$$

By Lemma 3.6, we have with probability at least $1 - 1/e$ that at least $|A_i|/801$ vertices in S_t have a neighbor in A_i that pulls the rumor from I_t before the end of round $t + \lceil 14d_i/|A_i| \rceil \le t + 224$, as $d_i \le 16|A_i|$. Therefore, with probability at least $1-1/e$, $Z_{t+224} \ge Z_t + |A_i|/801$, and thus, as in Case 2a,

$$\begin{aligned} W_{t+224} - W_t &\ge Z_{t+224} - Z_t \ge |A_i|/801 \\ &\ge |\partial I_t|/(4 \cdot 801\rho) \\ &\overset{(3.9),(3.10)}{\ge} \alpha W_t/(48 \cdot 801 \log \Delta), \end{aligned}$$

which proves the lemma.

(b) For $\ell := d_i/(96\rho)$,

$$\sum_{v \in V \colon h_i(v) \ge \ell} h_i(v) \ge (1/3)\,\mathrm{vol}(A_i). \tag{3.12}$$

Note that $\ell \ge 256$, because $d_i > d_{\min} = 96 \cdot 256\sqrt{\rho}\log\Delta$ and

$$\rho = \log(4\Delta/\delta) \le \log(4\Delta) < (\log\Delta)^2.$$

By Lemma 3.7, with probability at least $1 - 1/e$, at least a $(1/40)$-fraction of A_i is informed before the end of round $t + 512(\sqrt{\rho}\log\Delta + \rho) = t + 512\sqrt{\rho}(\log\Delta + \sqrt{\rho}) \le t + 1024\sqrt{\rho}\log\Delta$, counting only rumor transmissions initiated by vertices in A_i; thus

$$\mathbf{Pr}\left(\mathcal{A}_i(1/40,\ 1024\sqrt{\rho}\log\Delta)\right) \ge 1 - 1/e.$$

(c) Neither of Conditions (3.11) or (3.12) holds. It follows that at least a $(1 - 1/2 - 1/3) = (1/6)$-fraction of $\mathrm{vol}(A_i)$ corresponds to edges going to vertices $v \in V \setminus S_t$ with $h_i(v) < \ell$. Also, the volume that corresponds to edges going to ∂I_t is at most $|\partial I_t| \cdot \ell$. Thus,

$$\begin{aligned} |E(A_i, I_t)| &\ge \frac{\mathrm{vol}(A_i)}{6} - |\partial I_t| \cdot \ell \\ &\ge \frac{d_i|A_i|}{6} - (4|A_i|\rho) \cdot \frac{d_i}{96\rho} \\ &= \frac{d_i|A_i|}{8}. \end{aligned}$$

Let m be the number of vertices in A_i with at least $d_i/16$ neighbors in I_t, then $|E(A_i, I_t)| \le m \cdot 2d_i + (|A_i| - m)d_i/16$. Combining the last two inequalities yields $m \ge |A_i|/32$. We finish this case similarly to Case 2b: By Lemma 3.5, at least a $(1/64)$-fraction of A_i gets informed in $(2d_i)/(d_i/16) = 32$ rounds, with probability at least $1/2$. Thus, with probability at least $1/2$,

$$\begin{aligned} W_{t+32} - W_t &\ge |I_{t+23}| - |I_t| \ge |A_i|/64 \\ &\ge |\partial I_t|/(4 \cdot 64\rho) \\ &\overset{(3.9),(3.10)}{\ge} \alpha W_t/(36 \cdot 64 \log\Delta), \end{aligned}$$

which proves the lemma.

By the analysis above, if at least one A_i meets the conditions of some of the Cases 2a, 2b, 3a, or 3c, then Lemma 3.1 holds.

Otherwise, each A_i meets the conditions of Case 1 or Case 3b, and thus,

$$\begin{aligned} &\mathbf{Pr}\left(\mathcal{A}_i(1/40,\, r)\right) \ge 1/2, \quad \text{where} \\ &r := \max\left\{4d_{\min},\ 1024\sqrt{\rho}\log\Delta\right\} = 96 \cdot 256\sqrt{\rho}\log\Delta. \end{aligned}$$

We now explain how we can combine these results to obtain the lemma. Let R_i be the binary random variable with $R_i = |A_i|/40$ if the event $\mathcal{A}_i(1/40,\, r)$ occurs, and $R_i = 0$, otherwise. Further, let $R = \sum_{i \in \mathcal{I}} R_i$. Then, $\mathbf{E}[R] \ge \sum_{i \in \mathcal{I}} |A_i|/(2 \cdot 40)$. Since by construction each R_i depends only on transmissions of the rumor initiated by vertices in A_i, the R_i are independent random variables. A Hoeffding bound then yields

$$\begin{aligned} \mathbf{Pr}(R \le \mathbf{E}[R]/5) &\le \exp\left(-\frac{2((4/5)\,\mathbf{E}[R])^2}{\sum_{i \in \mathcal{I}}(|A_i|/40)^2}\right) \\ &\le \exp\left(-\frac{(4/5)^2\left(\sum_{i \in \mathcal{I}}|A_i|\right)^2}{2\sum_{i \in \mathcal{I}}|A_i|^2}\right) \\ &\le e^{-(4/5)^2/2} \le 3/4. \end{aligned}$$

And since $|I_{t+r}| - |I_t| \ge R$, it follows that with probability at least $1 - 3/4 = 1/4$,

$$|I_{t+r}| - |I_t| \ge \mathbf{E}[R]/5 \ge \sum_{i \in \mathcal{I}} |A_i|/400 \overset{(3.8)}{\ge} |\partial I_t|/800.$$

Thus, with probability at least 1/4,

$$W_{t+r} - W_t \geq |I_{t+r}| - |I_t| \geq \frac{|\partial I_t|}{800}$$
$$\overset{(3.10)}{\geq} \frac{\alpha W_t}{3 \cdot 800} = \frac{\beta r W_t}{3 \cdot 800 \cdot 96 \cdot 256},$$

since $r = 96 \cdot 256 \sqrt{\rho} \log \Delta = 96 \cdot 256 \alpha/\beta$. This completes the proof of Lemma 3.1.

4 Upper Bound for PUSH and PULL on Regular Graphs

Theorem 1.1 holds also for the PUSH and the PULL protocols for the class of regular graphs.

THEOREM 4.1. *For any d-regular graph with vertex expansion at least α, the PUSH and the PULL protocols inform all vertices in $\mathcal{O}(\alpha^{-1} \log n \log d)$ rounds with high probability.*

This upper bound is matched by the lower bound of Theorem 5.1 presented in the next section, for many choices of d and α.

We now describe the proof of Theorem 4.1. Consider the variant of the PUSH protocol where in each round every informed vertex chooses not the push the rumor independently with probability 1/2; we call this the *lazy PUSH* protocol. The next lemma shows that a round of this protocol is (stochastically) worse than a round of the PULL protocol.

LEMMA 4.1. *Fix a round $t \geq 0$ and the set I_t of informed vertices. Let N_z and N_l be the sets of newly informed vertices after one round of the lazy PUSH protocol and after one round of the PULL protocol, respectively. Then there is a coupling such that $N_z \subseteq N_l$.*

Proof. Let $u_1, u_2, \ldots, u_{|\partial I_t|}$ be the vertices in ∂I_t. We will expose the events $\{u_i \in N_z\}$ sequentially for $i = 1, \ldots, |\partial I_t|$, and prove that regardless of the previous history for $u_1, \ldots, u_{i-1}$, the conditional probability that $u_i \in N_z$ is at most equal to the probability that $u_i \in N_l$. Note that the events $\{u_i \in N_l\}$ are mutually independent, unlike the events $\{u_i \in N_z\}$.

Fix an i and suppose that u_i has k neighbors in I_t. For each $j = 1, \ldots, k$, let v_j denote the i-th of these k neighbors. We define the following events for the round of lazy PUSH. Let $\mathcal{A}_j$ be the event that v_j does *not* push the rumor to any of its *other* neighbors, i.e., the vertices in $N(v_i) \setminus \{u_i\}$. This means that either v_j chooses not to push the rumor, or it pushes the rumor to u. Let also $\mathcal{B}_j$ be the event that v_j pushes the rumor. Finally, let $\mathcal{F}_{i-1}$ be any event describing the states (informed or uninformed) of the vertices $u_1, \ldots, u_{i-1}$. Clearly, the probability that u_i gets informed is maximized when the events $\mathcal{A}_1, \ldots, \mathcal{A}_k$ occur, thus

$$\begin{aligned}
&\mathbf{Pr}(u_i \in N_z \mid \mathcal{F}_{i-1}) \\
&\quad \leq \mathbf{Pr}(u_i \in N_z \mid \mathcal{A}_1, \ldots, \mathcal{A}_k) \\
&\quad = \mathbf{Pr}(\mathcal{B}_1 \vee \cdots \vee \mathcal{B}_k \mid \mathcal{A}_1, \ldots, \mathcal{A}_k) \\
&\quad \leq \mathbf{Pr}(\mathcal{B}_1 \mid \mathcal{A}_1, \ldots, \mathcal{A}_k) + \cdots + \mathbf{Pr}(\mathcal{B}_k \mid \mathcal{A}_1, \ldots, \mathcal{A}_k) \\
&\quad = \mathbf{Pr}(\mathcal{B}_1 \mid \mathcal{A}_1) + \cdots + \mathbf{Pr}(\mathcal{B}_k \mid \mathcal{A}_k).
\end{aligned}$$

By Bayes' law,

$$\begin{aligned}
\mathbf{Pr}(\mathcal{B}_i \mid \mathcal{A}_i) &= \frac{\mathbf{Pr}(\mathcal{A}_i \mid \mathcal{B}_i) \cdot \mathbf{Pr}(\mathcal{B}_i)}{\mathbf{Pr}(\mathcal{A}_i)} \\
&\leq \frac{\mathbf{Pr}(\mathcal{A}_i \mid \mathcal{B}_i) \cdot \mathbf{Pr}(\mathcal{B}_i)}{\mathbf{Pr}(\mathcal{A}_i \wedge \overline{\mathcal{B}_i})} \\
&= \frac{\mathbf{Pr}(\mathcal{A}_i \mid \mathcal{B}_i) \cdot \mathbf{Pr}(\mathcal{B}_i)}{\mathbf{Pr}(\overline{\mathcal{B}_i})} \\
&= \frac{\mathbf{Pr}(\mathcal{A}_i \mid \mathcal{B}_i) \cdot (1/2)}{1 - 1/2} \\
&= \mathbf{Pr}(\mathcal{A}_i \mid \mathcal{B}_i) = 1/d.
\end{aligned}$$

Combining the above yields

$$\mathbf{Pr}(u_i \in N_z \mid \mathcal{F}_{i-1}) \leq k/d.$$

And since the probability that vertex u_i gets informed by the PULL protocol is $\mathbf{Pr}(u_i \in N_l) = k/d$, it follows that there is a coupling such that $N_z \subseteq N_l$. ∎

Consider now the variant of the PUSH-PULL protocol, called *lazy PUSH-PULL*, that is the combination of the lazy PUSH and the (normal) PULL protocols. I.e., in each round every informed vertex chooses not to push the rumor independently with probability 1/2, while every uninformed vertex performs a pull operation in every round. Using Lemma 4.1 we can show that a single round of the lazy PUSH-PULL protocol is worse than two rounds of the PULL protocol (in the same stochastic sense that a round of the lazy PUSH protocol is worse than a round of the PULL protocol). Further, it is straightforward to verify that Theorem 1.1 can be extended to the lazy PUSH-PULL protocol. Therefore, we conclude that for any d-regular graph with vertex expansion α, the PULL protocol informs all nodes in $\mathcal{O}(\alpha^{-1} \log n \log d)$ rounds w.h.p. Finally, since on regular graphs the runtime of the PUSH and the PULL protocols are asymptotically the same, it follows that the same $\mathcal{O}(\alpha^{-1} \log n \log d)$ bound w.h.p. holds for the PUSH protocol as well.

5 Lower Bound

The following theorem generalizes a lower bound of $\Omega(\log^2 n)$ rounds from [35] for graphs with constant vertex expansion.

THEOREM 5.1. *For any d, α (functions of n) with $d \geq 3$ and $d/\alpha \leq n^{1-\epsilon}$, for a fixed $\epsilon > 0$, there exists an infinite sequence of regular graphs $\{G_i\}_{i \geq 1}$ of increasing order n_i such that G_i has degree $d_i = d(n_i)$ and vertex expansion at least $\alpha_i = \Theta(\alpha(n_i))$, and with constant probability, the PUSH-PULL algorithm needs $\Omega(\alpha_i^{-1} \log(n_i) \log(d_i))$ rounds to inform all vertices.*

We now sketch a construction of such a family of graphs. This construction borrows ideas from constructions given in [9] and [35].

We start with a 3-regular expander graph on $n' = \alpha n/(d-2)$ vertices. Then we replace each vertex u by a cycle of length $1/\alpha$, and the three edges pointing to u point now to three vertices in the cycle such that the distance between these vertices is one third of the length of the cycle. Further, we add a few extra edges between the vertices of the cycle in such a way that every vertex has degree exactly 3, and the distance between any two vertices in the cycle decreases by at most a constant factor. The graph obtained is 3-regular with n'/α vertices, it has diameter $\Theta(\alpha^{-1} \log n')$, and most importantly, we show that it has vertex expansion $\Omega(\alpha)$. The last step in the construction is to take the Cartesian product of this graph with a $(d-2)$-clique. We prove that this operation retains the vertex expansion of the graph (although it reduces the conductance by a factor of d). Further, we show that it delays the spread of a rumor by a factor of $\Omega(\log d)$. The reason for this delay is that every vertex within a particular clique chooses a vertex outside this clique only with probability $\Theta(1/d)$. Hence, once a rumor has reached a particular clique, it takes about $\Omega(\log d)$ steps before enough vertices in the clique are informed to spread the rumor to a neighboring clique. Summarizing, the final graph has $(n'/\alpha) \cdot (d-2) = n$ vertices, it has degree d, vertex expansion $\Omega(\alpha)$, diameter $\Theta(\alpha^{-1} \log n')$, and rumor spreading takes a number of rounds that is at least $\Omega(\log d)$ times the diameter, i.e., $\Omega(\alpha^{-1} \log n' \cdot \log d) = \Omega(\alpha^{-1} \log(\alpha n/d) \log d) = \Omega(\alpha^{-1} \log n \log d)$.

5.1 Proof of Theorem 5.1. The proof is divided into two parts. In Section 5.1.1, we give the precise definition of the graph and analyze its expansion and its diameter. In Section 5.1.2, we then lower bound the runtime of the PUSH-PULL protocol on this graph.

5.1.1 Construction In this section we describe the construction of the sequence $G_1, G_2, \ldots$ of graphs. We begin with the following definition that essentially takes a 3-regular graph and replaces each vertex by a cycle of length ρ.

DEFINITION 5.1. *Let $H_n = (V, E)$ be any 3-regular graph with n vertices. For any $\rho \in \mathbb{N}$ that is a multiple of 3 and satisfies $(\rho-3)/3$ is a multiple of 4, define the graph $G_\rho(H_n) = (V_\rho, E_\rho)$ with vertex set $V_\rho := V \times \{1, \ldots, \rho\}$ and edge set E_ρ defined as follows. Take an arbitrary family of bijections $f_u : N(u) \to \{0, \rho/3, 2\rho/3\}$, $u \in V$, and let*

$$\tilde{E}_1 := \{\{(u, f_u(v)), (v, f_v(u))\} \mid \{u,v\} \in E(G)\},$$
$$\tilde{E}_2 := \{\{(u,i),(u,(i+1) \mod \rho)\} \mid i \in \{0, \ldots, \rho-1\}\},$$
$$\tilde{E}_3 := \{\{(u,i),(u,(i+2))\} \mid i \in \{1, \ldots, \rho/3-1\}, i \bmod 4 \in \{1,2\}\},$$
$$\tilde{E}_4 := \{\{(u,i),(u,(i+2))\} \mid i \in \{\rho/3+1, \ldots, 2\rho/3-1\}, i \bmod 4 \in \{2,3\}\},$$
$$\tilde{E}_5 := \{\{(u,i),(u,(i+2))\} \mid i \in \{2\rho/3+1, \ldots, \rho-1\}, i \bmod 4 \in \{3,0\}\}.$$

Then $E_\rho := \tilde{E}_1 \cup \tilde{E}_2 \cup \tilde{E}_3 \cup \tilde{E}_4 \cup \tilde{E}_5$.

We observe that graph $G_\rho(H_n)$ is also 3-regular, and its diameter is $\Theta((1/\rho)\operatorname{diam}(H_n)) = \Theta((1/\rho)\log n)$. The next lemma bounds the conductance of this graph.

LEMMA 5.1. *The graph $G_\rho(H_n)$ has conductance at least $(1/6)\Phi(H_n)/\rho$.*

The lemma implies that the vertex expansion of $G_\rho(H_n)$ is also at least $\Omega(\Phi(H_n)/\rho) = \Omega(\alpha(H_n)/\rho)$, since for any d-regular graph $\Phi \leq \alpha \leq d\Phi$.

Proof. Recall that $G_\rho(H_n)$ has $n \cdot \rho$ vertices. Let $S \subseteq V_\rho$ be any set with $1 \leq |S| \leq n \cdot \rho/2$. For each $u \in H_n$, let us call the set of vertices $\{u\} \times \{0, \ldots, \rho-1\} \subseteq V_\rho$ a cycle. Hence V_ρ can be partitioned into n cycles of length ρ each. Suppose first that at least $(1/4)|S|$ vertices are located in cycles that contain at most $\rho - 1$ vertices in S. As each of these cycles will contribute at least two to $|E_\rho(S, V_\rho \setminus S)|$, we have

$$|E_\rho(S, V_\rho \setminus S)| \geq 2 \cdot \frac{\frac{1}{4} \cdot |S|}{\rho - 1} = \frac{1}{2} \cdot \frac{|S|}{\rho-1}.$$

For the second case, assume that at least $(3/4)|S|$ vertices are located in a cycle which contains ρ vertices. Let $X \subseteq V$ be the corresponding indices of the cycles that only contain vertices in S. Note that $|X| \geq \frac{(3/4)|S|}{\rho}$, but also $|X| \leq |S|/\rho \leq n/2$. Let $\tilde{S}$ be the subset of S that only contains vertices of S that are located in a cycle which contains ρ vertices in S. Then by construction,

$$|E_\rho(\tilde{S}, V_\rho \setminus \tilde{S})| \geq |E(X, V \setminus X)| \geq 3\Phi(H_n) \cdot |X| \geq 3\Phi(H_n) \cdot \frac{1}{\rho} \cdot \frac{3}{4}|S| = \frac{9}{4} \cdot \frac{1}{\rho}\Phi(H_n)|S|.$$

Let us now argue that a similar lower bound also holds for $|E_\rho(S, V_\rho \setminus S)|$. Note that each edge $\{u_i, v_j\} \in E_\rho(\tilde{S}, V_\rho \setminus \tilde{S})$, $u_i \in \tilde{S}, v_j \in V_\rho \setminus \tilde{S}$ also appears in $E_\rho(S, V_\rho \setminus S)$, unless $v_j \in S \setminus \tilde{S}$. Now for each edge $\{u_i, v_j\} \in E_\rho$, $u_i \in \tilde{S}, v_j \in S$ which is in $E_\rho(\tilde{S}, V_\rho \setminus \tilde{S}) \cap E_\rho(\tilde{S}, S \setminus \tilde{S}) = E_\rho(\tilde{S}, S \setminus \tilde{S})$, there is at least one edge $\{v_j, v_l\} \in E_\rho(S, V_\rho \setminus S)$ with $1 \le l \le \rho$. Since H_n is 3-regular, we regard each such edge $\{v_j, v_l\} \in E_\rho(S, V_\rho \setminus S)$ at most three times. Hence,

$$|E_\rho(S, V_\rho \setminus S)| \ge \frac{1}{3} \cdot \frac{9}{4}\Phi(H_n) \cdot \frac{1}{\rho} \cdot |S| = \frac{3}{4} \cdot \frac{1}{\rho}\Phi(H_n)|S|.$$

■

Recall that the graph $G_\rho(H_n)$ is 3-regular. In order to increase the degree to d in a way that the vertex expansion remains the same but the rumor spreading time increases, we will consider the Cartesian product of $G_\rho(H_n)$ with a complete graph on $d-2$ vertices. We now recall the definition of the Cartesian product of two graphs.

Definition 5.2. *Let $G_1 = (V_1, E_1)$ and $G_2 = (V_2, E_2)$ be two graphs. We define the graph $G := G_1 \times G_2$ such that it has vertex set $V := V_1 \times V_2$ and contains edges between each pair of vertices (u_1, u_2) and (v_1, v_2) for which either $\{u_1, v_2\} \in E_1$ and $u_2 = v_2$, or $\{u_2, v_2\} \in E_2$ and $u_1 = v_1$.*

The next lemma both improves and generalizes result [35, Lemma 4.2].

Lemma 5.2. *Let G_1 be a graph with n_1 vertices and vertex expansion α_1, and let G_2 be a graph with n_2 vertices and vertex expansion $\alpha_2 \le \alpha_1$. Then, the vertex expansion of $G := G_1 \times G_2$ is at least $\alpha_1 \cdot \alpha_2/96$.*

Before we prove the lemma, we point out that similar results were shown by Chung and Tetali [11] for the edge expansion (isoperimetric number) of Cartesian products of graphs. However, such a result does not hold in general for the conductance. For example, consider the graph $G = G_1 \times G_2$, where G_1 and G_2 are cliques of size $n/2$ and 2, respectively. Although both G_1 and G_2 have constant conductance, the graph G has conductance $\Theta(1/n)$ (see also Remark 5.1.)

Proof of Lemma 5.2. We first prove the following simple claim, which says that the vertex expansion of a large set can essentially be reduced to the vertex expansion of the complement set.

Claim 5.3. *Let G be any graph with vertex expansion at least α. Let S be any subset of vertices with $n/2 < |S| \le n$. Then,*

$$|\partial S| \ge \frac{\alpha}{2}(n - |S|).$$

Proof. Suppose for contradiction that $|\partial S| < \frac{\alpha}{2}(n-|S|)$. As $V \setminus S$ is the disjoint union of ∂S and $V \setminus (S \cup \partial S)$, this would imply that $|V \setminus (S \cup \partial S)| \ge n - |S| - \frac{\alpha}{2}(n-|S|) \ge \frac{1}{2}(n-|S|)$. But since $\partial(V \setminus (S \cup \partial S)) \subseteq \partial S$, we conclude by using the vertex expansion for the set $V \setminus (S \cup \partial S)$ that

$$\begin{aligned}|\partial S| &\ge |\partial(V \setminus (S \cup \partial S))| \ge \alpha \cdot |V \setminus (S \cup \partial S)| \\ &\ge \alpha \cdot \frac{1}{2}(n - |S|),\end{aligned}$$

which yields the desired contradiction. ■

We now turn to the proof of Lemma 5.2. Fix any set $S \subseteq V = V(G)$ with $1 \le |S| \le n_1 n_2/2$. We proceed by a case distinction.

1. At least $\alpha_2/16 \cdot |S|$ of the nodes are in a graph $\{i\} \times V_2$, $1 \le i \le n_1$ which contains at most $(3/4)n_2$ vertices in S, i.e.,

$$\begin{aligned}&|\{u \in S, u = (i,j) \in V_1 \times V_2 : \\ &\qquad |\{\{i\} \times V_2 \cap S\}| \le (3/4)n_2\}| \ge \frac{\alpha_2}{16}|S|.\end{aligned}$$

   Let $X \subseteq \{1, \dots, n_1\}$ be the indices i with the property above. Applying Claim 5.3,

$$\begin{aligned}|\partial S| &\ge \sum_{i \in X} |\partial S \cap (\{i\} \times V_2)| \\ &\ge \sum_{i \in X} \frac{\alpha_1}{2} \cdot \min\{|S \cap (\{i\} \times V_2)|, \\ &\qquad\qquad n_2 - |S \cap (\{i\} \times V_2)|\} \\ &\ge \sum_{i \in X} \frac{\alpha_1}{2} \cdot \frac{1}{3} \cdot |S \cap (\{i\} \times V_2)| \\ &\ge \frac{\alpha_1}{6} \cdot \sum_{i \in X} |S \cap (\{i\} \times V_2)| \\ &\ge \frac{\alpha_1 \alpha_2}{96} \cdot |S|.\end{aligned}$$

2. More than $(1 - \alpha_2/16) \cdot |S|$ of the nodes are in a graph $\{i\} \times G_2$, $1 \le i \le n_1$, with more than $\frac{3}{4}n_2$ vertices in S. Let $X \subseteq \{1, \dots, n_1\}$ be the indices i with the property above. Moreover, let $\tilde{S}$ be the set of vertices with the property above. First observe that $|X| \le (2/3)n_1$ and $|X| \ge (1 - \alpha_2/4) \cdot |S|/n_1$. Our aim is to prove a lower bound on $\partial\tilde{S}$. To this end, consider first the set of vertices $\widehat{S}$ which contains *all* vertices in X, i.e., $\widehat{S} := X \times V_2$. Using Claim 5.3 and the definition of the Cartesian product, we find that

$$\begin{aligned}|\partial\widehat{S}| &\ge n_2 \cdot \partial_{G_1} X \\ &\ge n_2 \cdot \frac{\alpha_1}{2} \min\{|X|, n_1 - |X|\} \\ &\ge n_2 \cdot \frac{\alpha_1}{4}|X|,\end{aligned}$$

where the last inequality used $|X| \le (2/3)n_1$. Note that $\partial\widehat{S}$ can be written as

$$\partial\widehat{S} = \cup_{j\in Y}(j \times V_2).$$

where $Y = \partial_{G_1}X \subseteq \{1,\dots,n_1\}$. If we transform $\widehat{S}$ into $\tilde{S}$, then we remove from each subgraph $j \times V_2$, $j \in Y$, at most $(1/4)n_2$ vertices from the boundary. Hence,

$$|\partial\tilde{S}| \ge \frac{3}{4}n_2 \cdot \frac{\alpha_1}{4}|X| \ge \frac{3}{4}n_2 \cdot \frac{\alpha_1}{4}\cdot\frac{3}{4}\cdot\frac{|S|}{n_2} = \frac{9}{64}\cdot\alpha_1|S|.$$

On the other hand, $|S \setminus \tilde{S}| \le \alpha_2/16 \cdot |S|$ by our case assumption. Using this along with the above inequalities and finally recalling that $\alpha_1 \ge \alpha_2$, we arrive at

$$\begin{aligned}|\partial S| &\ge |\partial\tilde{S}| - |S\setminus\tilde{S}|\\ &\ge \frac{9}{64}\cdot\alpha_1|S| - \alpha_2/16|S|\\ &\ge \frac{9}{64}\cdot\alpha_1|S| - \alpha_1/16|S| = \frac{5}{64}\alpha_1\cdot|S|.\end{aligned}$$

This completes the proof of Lemma 5.2. ∎

We now recall a result from [2] about the construction of 3-regular expander graphs.

THEOREM 5.2. ([2, THEOREM 1]) *Fix a small enough constant $\epsilon > 0$. Then for every sufficiently large integer n, there is a graph H_n with girth at least $(1/10)\ln n$ and vertex expansion at least ϵ.*

Finally, we are now able to complete the construction. We start with the graph H_n from Theorem 5.2, plug this graph into Definition 5.1 and take the Cartesian product of the resulting graph with a complete graph. Combining our previous lemmata, we can lower bound the vertex expansion of that graph.

LEMMA 5.4. *For any α such that $\rho := 1/\alpha$ satisfies the constraints of Definition 5.1, any $d \ge 3$, and any n that is a multiple of $(d-2)/\alpha$, define*

$$G := G_{(1/\alpha)}(H_{\alpha n/(d-2)}) \times K_{d-2},$$

where H_n is the graph from Theorem 5.2, and K_{d-2} is a complete graph with $d-2$ vertices.[6] *Then, graph G has degree $3 + (d-3) = d$, number of vertices $\big((1/\alpha)\cdot\alpha n/(d-2)\big)\cdot(d-2) = n$, and vertex-expansion at least $\Omega(\alpha)$.*

[6]For the special case $d = 3$, we simply consider the graph $G_{(1/\alpha)}(H_{\alpha n/(d-2)})$.

Proof. Recall that by Theorem 5.2, $H_{\alpha n/(d-2)}$ has constant vertex expansion. Hence by Lemma 5.1, $G_{(1/\alpha),n}(H_{\alpha n/(d-2)})$ has vertex expansion $\Omega(\alpha)$. Further, by Lemma 5.2, G has vertex expansion $\Omega(\alpha)$, too, since the graph K_{d-2} has vertex expansion $\Omega(1)$. ∎

For the sake of completeness, let us also mention that the conductance of G is considerably smaller than the vertex expansion if $d = \omega(1)$.

REMARK 5.1. *The graph G defined in Lemma 5.4 has conductance $\mathcal{O}(\alpha/d)$.*

5.1.2 Runtime Analysis.

We now analyze the runtime of the PUSH-PULL protocol on the graph G defined in Lemma 5.4. In the following, we identify the set of nodes $\{(u,i,j)\colon 1 \le j \le d-2\}$ by a so-called *supernode* $(u,i) \in G_{(1/\alpha),n}(H_{\alpha n/(d-2)})$.

Note that for constant d, we can use the trivial lower bound $\Omega(\mathrm{diam}(G)) = \Omega((1/\alpha)\log n)$. Therefore, we may assume in the following that $d = \omega(1)$.

LEMMA 5.5. *Consider the graph G from Lemma 5.4. For an arbitrary vertex $u \in H_{\alpha n/(d-2)}$, let $(u,0),(u,1),\dots,(u,1/\alpha-1)$ be the corresponding supernodes in G. Let T_1 be the first round when one of the supernodes $(u,0),(u,1/(3\alpha))$ or $(u,2/(3\alpha))$ receives the information for the first time. Moreover, let $T_2 \ge T_1$ be the first round when another one of these three supernodes becomes informed, where we only allow transmissions initiated by the nodes in $u \times \{0,1,\dots,(1/\alpha)-1\}\times\{1,\dots,d\}$. Then, for every $\epsilon > 0$ there is a constant $c = c(\epsilon)$ such that*

$$\mathbf{Pr}\left(T_2 - T_1 \ge c\cdot\frac{1}{\alpha}\cdot\log d\right) \ge 1-\epsilon.$$

The proof of Lemma 5.5 is given in the full version.

LEMMA 5.6. *Let G be a tree with n nodes where all nodes have degree 3 except for the leaves at depth $D = \log n/(\log 6)+1$. Assume further that each node u gets a random weight w_u which satisfies $\mathbf{Pr}(w_u \ge \gamma) \ge 1-1/64$ for some value $\gamma \ge 0$, independently of all other weights. Then with probability at least $1-3\cdot2^{-D}$, it holds for each path from the root to a leaf $P = (u_1,u_2,\dots,u_D)$ that $\sum_{i=1}^{D} w_{u_i} \ge (D/2)\gamma$.*

Proof. Fix any path $P = (u_0,\dots,u_D)$. The probability that more than half of the nodes have a weight less than γ can be bounded from above by

$$\binom{D}{\frac{1}{2}D}\cdot\left(\frac{1}{64}\right)^{\frac{1}{2}D} \le 2^D\cdot 8^{-D} = 4^{-D}.$$

Taking the union bound, we obtain that *all* paths from the root to a leaf, at least half of the nodes have a weight more than γ with probability $1-3\cdot 2^{D-1}\cdot 4^{-D} \geq 1-3\cdot 2^{-D}$. ■

We complete the runtime analysis by combining Lemma 5.5 and Lemma 5.6 as follows. Suppose that the rumor is initially placed at some node $(s,0,0) \in V(G)$. Then by Theorem 5.2, there is a constant $\epsilon > 0$ so that the neighborhood of the node s in the graph $H_{\alpha n/d}$ up to distance $D := \epsilon \log_2(\alpha n/d)$ is a tree. Let L be the set of vertices in $H_{\alpha n/d}$ with $\mathrm{dist}(s,L) = D$. To inform the first node (u,i,r) with $u \in L, 0 \leq i \leq 1/\alpha - 1, 1 \leq r \leq d-2$, the rumor has to be spread corresponding to a path $P = (s, v_1, v_2, \ldots, u = v_D)$ in H. However, by Lemma 5.5 we know that for each $1 \leq k \leq D$ the number of rounds for the rumor to reach a vertex (v_k, x, y) with $0 \leq x \leq 1/\alpha - 1$, $1 \leq y \leq d-2$ from a vertex (v_{k-1}, x', y') with $0 \leq x' \leq 1/\alpha - 1$, $1 \leq y' \leq d-2$ is at least $\Omega(\frac{1}{\alpha} \cdot \log(d))$ with probability $1 - 1/64$. Hence we can apply Lemma 5.6 to conclude that with probability $1 - o(1)$, it takes at least $\Omega(\frac{1}{\alpha} \cdot \log(d) \cdot \log(\alpha n/d))$ rounds to inform the first node (u_i, r) with $u \in L, 0 \leq i \leq 1/\alpha - 1, 1 \leq r \leq d-2$.

References

[1] J. Adler. Bootstrap percolation. *Physica A: Statistical and Theoretical Physics*, 171(3):453–470, 1991.

[2] M. Ajtai. Recursive construction for 3-regular expanders. *Combinatorica*, 14(4):379–416, 1994.

[3] N. Alon. Eigenvalues and expanders. *Combinatorica*, 6(2):83–96, 1986.

[4] L. Babai. Local expansion of vertex-transitive graphs and random generation in finite groups. In *Proceedings of the 23rd ACM Symposium on Theory of Computing (STOC'91)*, pages 164–174, 1991.

[5] N. Berger, C. Borgs, J. T. Chayes, and A. Saberi. On the spread of viruses on the Internet. In *Proceedings of the 16th ACM-SIAM Symposium on Discrete Algorithms (SODA'05)*, pages 301–310, 2005.

[6] S. Boyd, A. Ghosh, B. Prabhakar, and D. Shah. Randomized gossip algorithms. *IEEE Transactions on Information Theory*, 52(6):2508–2530, 2006.

[7] K. Censor-Hillel and H. Shachnai. Fast information spreading in graphs with large weak conductance. In *Proceedings of the 22nd ACM-SIAM Symposium on Discrete Algorithms (SODA'11)*, pages 440–448, 2011.

[8] F. Chierichetti, S. Lattanzi, and A. Panconesi. Rumor spreading in social networks. In *Proceedings of the 36th International Colloquium on Automata, Languages, and Programming (ICALP'09)*, pages 375–386, 2009.

[9] F. Chierichetti, S. Lattanzi, and A. Panconesi. Almost tight bounds for rumour spreading with conductance. In *Proceedings of the 42nd ACM Symposium on Theory of Computing (STOC'10)*, pages 399–408, 2010.

[10] F. Chierichetti, S. Lattanzi, and A. Panconesi. Rumour spreading and graph conductance. In *Proceedings of the 21th ACM-SIAM Symposium on Discrete Algorithms (SODA'10)*, pages 1657–1663, 2010.

[11] F. R. K. Chung and P. Tetali. Isoperimetric inequalities for cartesian products of graphs. *Combinatorics, Probability & Computing*, 7(2):141–148, 1998.

[12] A. Czumaj and C. Sohler. Testing expansion in bounded-degree graphs. In *Proceedings of the 48th IEEE Symposium on Foundations of Computer Science (FOCS'07)*, pages 570–578, 2007.

[13] A. Demers, D. Greene, C. Hauser, W. Irish, J. Larson, S. Shenker, H. Sturgis, D. Swinehart, and D. Terry. Epidemic algorithms for replicated database maintenance. In *Proceedings of the 6th ACM-SIGOPT Symposium on Principles of Distributed Computing (PODC'87)*, pages 1–12, 1987.

[14] B. Doerr, M. Fouz, and T. Friedrich. Social networks spread rumors in sublogarithmic time. In *Proceedings of the 43rd ACM Symposium on Theory of Computing (STOC'11)*, pages 21–30, 2011.

[15] R. Elsässer and T. Sauerwald. Broadcasting vs. mixing and information dissemination on cayley graphs. In *Proceedings of the 24th International Symposium on Theoretical Aspects of Computer Science (STACS'07)*, pages 163–174, 2007.

[16] R. Elsässer and T. Sauerwald. Cover time and broadcast time. In *Proceedings of the 26th International Symposium on Theoretical Aspects of Computer Science (STACS'09)*, pages 373–384, 2009.

[17] S. Eubank, H. Guclu, V. A. Kumar, M. V. Marathe, A. Srinivasan, Z. Toroczkai, and N. Wang. Modelling disease outbreaks in realistic urban social networks. *Nature*, 429:180–184, 2004.

[18] S. Eubank, V. A. Kumar, M. V. Marathe, A. Srinivasan, and N. Wang. Structure of social contact networks and their impact on epidemics. *AMS-DIMACS Special Issue on Epidemiology*, 2006.

[19] U. Feige, D. Peleg, P. Raghavan, and E. Upfal. Randomized broadcast in networks. *Random Structures and Algorithms*, 1(4):447–460, 1990.

[20] N. Fountoulakis, A. Huber, and K. Panagiotou. Reliable broadcasting in random networks and the effect of density. In *Proceedings of the 29th IEEE Conference on Computer Communications (INFOCOM'10)*, pages 2552–2560, 2010.

[21] N. Fountoulakis and K. Panagiotou. Rumor spreading on random regular graphs and expanders. In *Proceedings of the 14th International Workshop on Randomization and Approximation Techniques in Computer Science (RANDOM'10)*, pages 560–573, 2010.

[22] A. Frieze and G. Grimmett. The shortest-path problem for graphs with random arc-lengths. *Discrete Applied Mathematics*, 10:57–77, 1985.

[23] G. Giakkoupis. Tight bounds for rumor spreading in graphs of a given conductance. In *Proceedings of the 28th International Symposium on Theoretical Aspects of Computer Science (STACS'11)*, pages 57–68, 2011.

[24] M. Harchol-Balter, F. T. Leighton, and D. Lewin. Resource discovery in distributed networks. In *Proceedings of the 18th ACM-SIGOPT Symposium on Principles of Distributed Computing (PODC'99)*, pages 229–237, 1999.

[25] S. Hoory, N. Linial, and A. Wigderson. Expander graphs and their applications. *Bulletin of the American Mathematical Society*, 43(4):439–561, 2006.

[26] R. Karp, C. Schindelhauer, S. Shenker, and B. Vöcking. Randomized rumor spreading. In *Proceedings of the 41st IEEE Symposium on Foundations of Computer Science (FOCS'00)*, pages 565–574, 2000.

[27] J. Leskovec and C. Faloutsos. Sampling from large graphs. In *Proceedings of the 12th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD'06)*, pages 631–636, 2006.

[28] J. Leskovec, K. J. Lang, A. Dasgupta, and M. W. Mahoney. Statistical properties of community structure in large social and information networks. In *Proceedings of the 17th International Conference on World Wide Web (WWW'08)*, pages 695–704, 2008.

[29] L. Lovász. Random walks on graphs: A survey. *Combinatorics, Paul Erdös is Eighty*, 2:1–46, 1993.

[30] A. S. Maiya and T. Y. Berger-Wolf. Sampling community structure. In *Proceedings of the 19th International Conference on World Wide Web (WWW'10)*, pages 701–710, 2010.

[31] C. McDiarmid. Concentration. In *Probabilistic Methods for Algorithmic Discrete Mathematics*, pages 1–46. Springer, 1998.

[32] M. Mihail, C. H. Papadimitriou, and A. Saberi. On certain connectivity properties of the internet topology. *J. Comput. Syst. Sci.*, 72(2):239–251, 2006.

[33] D. Mosk-Aoyama and D. Shah. Fast distributed algorithms for computing separable functions. *IEEE Transactions on Information Theory*, 54(7):2997–3007, 2008.

[34] B. Pittel. On spreading a rumor. *SIAM Journal on Applied Mathematics*, 47(1):213–223, 1987.

[35] T. Sauerwald and A. Stauffer. Rumor spreading and vertex expansion on regular graphs. In *Proceedings of the 22nd ACM-SIAM Symposium on Discrete Algorithms (SODA'11)*, pages 462–475, 2011.

[36] R. van Renesse, Y. Minsky, and M. Hayden. A gossip-style failure detection service. In *Proceedings of the 15th IFIP International Conference on Distributed Systems Platforms (Middleware'98)*, pages 55–70, 1998.

A Relation Between Vertex Expansion and Conductance

FACT A.1. *For any graph,* $(\delta/\Delta)\Phi \le \alpha \le \Delta\Phi$.

Proof. The first inequality can be obtained as follows.

$$\begin{aligned}
\alpha &= \min_{U\subseteq V,\ 0<|U|\le n/2} \frac{|\partial U|}{|U|} \\
&\ge \min_{U\subseteq V,\ 0<|U|\le n/2} \frac{|E(U,V\setminus U)|/\Delta}{|U|} \\
&= \min_{U\subseteq V,\ 0<|U|<n} \frac{|E(U,V\setminus U)|/\Delta}{\min\{|U|,|V\setminus U|\}} \\
&\ge \min_{U\subseteq V,\ 0<|U|<n} \frac{|E(U,V\setminus U)|/\Delta}{\min\{\mathrm{vol}(U)/\delta,\mathrm{vol}(V\setminus U)/\delta\}} \\
&= \min_{U\subseteq V,\ 0<\mathrm{vol}(U)<\mathrm{vol}(V)} \frac{|E(U,V\setminus U)|/\Delta}{\min\{\mathrm{vol}(U)/\delta,\mathrm{vol}(V\setminus U)/\delta\}} \\
&= \min_{U\subseteq V,\ 0<\mathrm{vol}(U)\le\mathrm{vol}(V)/2} \frac{|E(U,V\setminus U)|/\Delta}{\mathrm{vol}(U)/\delta} \\
&= (\delta/\Delta)\Phi.
\end{aligned}$$

Similarly, for the second inequality we have

$$\begin{aligned}
\alpha &\le \min_{U\subseteq V,\ 0<|U|\le n/2} \frac{|E(U,V\setminus U)|}{|U|} \\
&= \min_{U\subseteq V,\ 0<|U|<n} \frac{|E(U,V\setminus U)|}{\min\{|U|,|V\setminus U|\}} \\
&\le \min_{U\subseteq V,\ 0<|U|<n} \frac{|E(U,V\setminus U)|}{\min\{\mathrm{vol}(U)/\Delta,\mathrm{vol}(V\setminus U)/\Delta\}} \\
&= \Delta\Phi.
\end{aligned}$$

■