Tight Bounds for Rumor Spreading with Vertex Expansion

George Giakkoupis*

Abstract

We establish a bound for the classic PUSH-PULL rumor spreading protocol on general graphs, in terms of the vertex expansion of the graph. We show that $O(\log^2(n)/\alpha)$ rounds suffice with high probability to spread a rumor from any single node to all n nodes, in any graph with vertex expansion at least α. This bound matches a known lower bound, and settles the natural question on the relationship between rumor spreading and vertex expansion asked by Chierichetti, Lattanzi, and Panconesi (SODA 2010). Further, some of the arguments used in the proof may be of independent interest, as they give new insights, for example, on how to choose a small set of nodes in which to plant the rumor initially, to guarantee fast rumor spreading.

1 Introduction

We study a classic randomized protocol for information dissemination in networks, known as *(randomized) rumor spreading*. The protocol proceeds in a sequence of synchronous rounds.[1] Initially, in round 0, an arbitrary node learns a piece of information, the *rumor*. This rumor is then spread iteratively to other nodes: In each round, every *informed* node (i.e., every node that learned the rumor in a previous round) chooses a random neighbor and sends the rumor to that neighbor. This is the PUSH version of the protocol. The PULL version is symmetric: In each round, every *uninformed* node contacts a random neighbor, and if this neighbor knows the rumor it sends it to the uninformed node. Finally, the PUSH-PULL algorithm is the combination of both: 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.

The above protocols were proposed almost thirty years ago, and have been the subject of extensive study, especially in the past decade. The most studied question concerns the number of rounds these protocols need to spread a rumor in various network topologies. It has been shown that $O(\log n)$ rounds suffice with high probability (w.h.p.) for several families for networks, from basic communication networks, such as complete graphs and hypercubes, to more complex structures, such as preferential attachment graphs modeling social networks (see the Related Work Section).

A main motivation for the study of rumor spreading is its application to algorithms for broadcasting in communication networks [8, 14, 22]. Rumor spreading provides a scalable alternative to the flooding protocol (where each node sends the information to *all* its neighbors in a round), and a simpler and more robust alternative to deterministic solutions. The advantages of simplicity (each node makes a simple local decision in each round; no knowledge of the global topology is needed; no state is maintained), scalability (each node initiates just one connection per round), and robustness (the protocol tolerates random node/link failures without the use of error recovery mechanisms) make rumor spreading protocols particularly suited for today's distributed networks of massive-scale. Such networks, e.g., peer-to-peer, mobile ad-hoc, or sensor networks, are highly dynamic, suffer from frequent link and node failures, or nodes have limited computational, communication, and energy resources.

Another motivation for the study of rumor spreading protocols is that they provide intuition as to how information spreads in social networks [5]. Understanding these simple rumor spreading protocols may lead to a better understanding of more realistic epidemic processes on social and other complex networks.

Our main focus in this paper is the connection between rumor spreading and graph expansion properties of networks. Many of the topologies for which rumor spreading is known to be fast have high expansion. Further, empirical studies indicate that social networks also have good expansion properties [13, 24].

Several works have studied the relationship of rumor spreading with the *conductance* of the network graph [25, 7, 6, 19]. The conductance $\phi \in (0, 1]$ is a standard expansion measure defined roughly as the minimum ratio of the edges leaving a set of nodes over the total number of edges incident on these nodes (see Section 2). The main result of the above works is an upper bound of $O(\log(n)/\phi)$ rounds for PUSH-PULL to inform all nodes w.h.p., for any n-node graph with conductance

*INRIA Rennes, France, george.giakkoupis@inria.fr. Funded, in part, by INRIA Associate Team RADCON and ERC Starting Grant GOSSPLE 204742.

[1] There are also asynchronous versions of rumor spreading (see, e.g., [1]), but in this paper we focus on synchronous protocols.

at least ϕ. This bound is tight, as there are graphs with diameter $\Omega(\log(n)/\phi)$ [6].

The above result has found applications to the design of various information dissemination protocols [4, 2, 21]. These protocols rely on the fact that PUSH-PULL spreads information fast in subgraphs of high conductance, and they combine PUSH-PULL with more sophisticated rules on how each node chooses the neighbor to contact in each round. Those protocols achieve fast information spreading in a broader class of networks, and some even achieve for all graphs time bounds that are close (within poly-logarithmic factors) to the network diameter, which is the natural lower bound for information dissemination in networks.

More recently, another standard measure of expansion, the *vertex expansion*, has been studied in connection with rumor spreading. The vertex expansion $\alpha \in (0, 1]$ of a graph is, roughly, the minimum ratio of the neighbors that a set of nodes has (and are not in the set) over the size of the set. In general, vertex expansion is incomparable to conductance, as there are graphs with high vertex expansion but low conductance, and vice versa. (For an account of the differences between the two measures see [20].) The question of whether high vertex expansion implies fast rumor spreading (similar to high conductance) was highlighted as an interesting and challenging open problem in [7]. This problem was subsequently studied in [27, 20], and the main result obtained was an upper bound for PUSH-PULL of $O(\log^{2.5}(n)/\alpha)$ rounds w.h.p., for any graph with vertex expansion at least α. The precise bound is $O(\log n \cdot \log \Delta \cdot \sqrt{\log(2\Delta/\delta)}/\alpha)$, where Δ and δ are the maximum and minimum node degrees, respectively. Further, a lower bound of $\Omega(\log n \cdot \log(\Delta)/\alpha)$ was shown, assuming $\Delta/\alpha \leq n^{1-\epsilon}$ for a constant $\epsilon > 0$.

Our Contribution. Our main result is the following upper bound for PUSH-PULL in terms of vertex expansion, which matches the lower bound from [20].

THEOREM 1.1. *Let $G = (V, E)$ be a graph with $|V| = n$ nodes, maximum degree at most Δ, and vertex expansion at least α. For any such graph G and constant $\beta > 0$, with probability $1 - O(n^{-\beta})$* PUSH-PULL *informs all nodes of G in $O\big(\log n \cdot \log(\Delta)/\alpha\big)$ rounds.*

This result, together with the $O(\log(n)/\phi)$ bound with conductance, resolve completely the natural question asked by Chierichetti, Lattanzi, and Panconesi [5, 7], on the relationship between rumor spreading and the two most standard measures of graph expansion.

A rough outline of the proof is as follows. Let S be the set of informed nodes after a given round, and let ∂S be the boundary of S, i.e., the set of nodes from $V - S$ that have a neighbor in S. The proof defines a new simple measure of the expansion of S, called *boundary expansion*. If S has low boundary expansion, then we prove that a constant fraction of the boundary ∂S gets informed in an expected number of $O(\log \Delta)$ rounds; this is the core argument of the proof. If, instead, S has high boundary expansion, then we have that an expected number of $\Omega(|\partial S|)$ nodes from $V - (S \cup \partial S)$ are added to the boundary in a single round. It follows that a simple potential function $\Psi(S)$ that counts 1 for each informed node and 1/2 for each node in the boundary, increases "on average" per round by at least $\Omega(|\partial S|/ \log \Delta) = \Omega(\alpha \Psi(S)/ \log \Delta)$; this is the right increase rate we need to show the $O\big(\log n \cdot \log(\Delta)/\alpha\big)$ time bound.

Some of the arguments used in the proof may be of independent interest as tools for the analysis of rumor spreading. We demonstrated this by applying those arguments to show the following two smaller results.

1. Our first result serves as a "warm-up" for the main proof, as it uses a simpler version of the core argument of our analysis. We show that if a rumor is initially known to a subset of nodes that is a *dominating set*, then $O(\log n)$ rounds of PUSH-PULL suffice w.h.p. to inform the other nodes. (In fact we can use just PULL instead of PUSH-PULL.)[2] This result is somewhat relevant to problems in viral marketing [10, 23]: It says that if we want to plant a rumor, or ad, in a (small) initial set of nodes in an arbitrary network, so that the remaining nodes get informed quickly by rumor spreading, then it suffices that the set we choose be a dominating set.

2. Our second result reuses the main argument from the proof of Result 1 above, and establishes that PUSH-PULL spreads a rumor in $O(\log n)$ rounds w.h.p. in any graph of diameter (at most) 2. This can be viewed as an extension to the classic result that rumor spreading takes $O(\log n)$ rounds in graphs of diameter 1, i.e., in complete graphs. We point out that unlike complete graphs, graphs of diameter 2 may have low expansion. Also the result does not hold for graphs of diameter 3. (The proof of the above result can be found in the appendix.)

Related Work. The first works on rumor spreading provided a precise analysis of PUSH on complete graphs [18, 26]. Time bounds of $O(\log n)$ rounds were later proved for hypercubes and random graphs [14].

[2] Note that the analysis of the more sophisticated information dissemination algorithm proposed in [2, 21], yields time bounds of $O(\log^3 n)$ and $O(\log^2 n)$, respectively, for this setting.

Other symmetric graphs similar to the hypercube in which rumor spreading takes $O(\log n + \text{diam})$ rounds were studied in [12]. A refined analysis for random graphs proving essentially the same time bound as for complete graphs was provided in [15], and extended to random regular graphs in [16]. The authors of [5] studied rumor spreading on preferential attachment graphs, which are used as models for social networks, and showed that PUSH or PULL need polynomially many rounds, whereas PUSH-PULL needs only $O(\log^2 n)$ rounds. The last bound was subsequently improved to $\Theta(\log n)$ [9]. Another class of random graphs used to model social networks was considered in [17], and it was shown that PUSH-PULL needs just $\Theta(\log \log n)$ rounds to inform all but an ϵ-fraction of nodes.

For general graphs, a bound of $O(\log(n)/\Phi)$ was shown in [25] for a version of PUSH-PULL with non-uniform probabilities for neighbor selection, where Φ is the conductance of the matrix of selection probabilities (which is different than the conductance ϕ of the graph). A comparable bound in terms of the mixing time of an appropriate random walk was shown in [1]. For PUSH-PULL, a polynomial bound in $\log(n)/\phi$ was shown in [5] via a connection to a spectral sparsification process. An improved, almost tight bound was shown in [6], and the tight $O(\log(n)/\phi)$ bound was shown in [19]. For PUSH or PULL this bound holds for regular graphs, but not for general graphs.

The bound for PUSH-PULL with conductance has been used in subsequent works [3, 4, 2, 21], mainly to argue that rumors spread fast in *subgraphs* of high conductance. A refinement of conductance, called *weak conductance*, which is greater or equal to ϕ was introduced in [3] and was related to the time for PUSH-PULL to inform a certain fraction of nodes. A gossip protocol for the problem in which every node has a rumor initially, and must receive the rumors of all other nodes was proposed in [4]. The protocol alternates rounds of PUSH-PULL with rounds of deterministic communication, and guarantees fast information spreading in all graphs with high weak conductance. In [2, 21] protocols for the same gossip problem were proposed that need only $O(\text{diam} \cdot \text{polylog}(n))$ or $O(\text{diam} + \text{polylog}(n))$ rounds. One of these protocols is even deterministic [21]. Similar to the conductance-based bounds, we believe that the results with vertex expansion we present in this paper could potentially help in the design of new protocols for information dissemination.

2 Definitions and Notation

Throughout the paper we assume that graph $G = (V, E)$ is connected, and $|V| = n$. For a node $v \in V$, we denote by $N(v)$ the set of v's neighbors in G, and $\deg(v) = |N(u)|$ is the degree of v. By Δ and δ we denote the maximum and minimum node degrees of G. For a set $S \subseteq V$, we denote by ∂S the *boundary* of S, that is, the set of neighbors of S that are not in S; formally, $\partial S = \{v \in V - S : N(v) \cap S \neq \emptyset\}$. We will write S^+ to denote set $S \cup \partial S$. A set S is a *dominating set* iff each node $v \in V$ either belongs to S or has a neighbor in S, i.e., $S^+ = V$. The vertex expansion of a non-empty set $S \subseteq V$, and the vertex expansion of graph G are defined respectively as

$$\alpha(S) = \frac{|\partial S|}{|S|}, \quad \alpha(G) = \min_{0<|S|\leq \frac{n}{2}} \alpha(S).$$

The volume of a set S is $\text{vol}(S) = \sum_{v \in S} \deg(v)$. By $E(S, V-S)$ we denote the set of edges with one endpoint from S and the other from $V - S$. The conductance of a non-empty set $S \subseteq V$, and the conductance of graph G are defined respectively as

$$\phi(S) = \frac{|E(S, V - S)|}{\text{vol}(S)}, \quad \phi(G) = \min_{0<\text{vol}(S)\leq \frac{\text{vol}(V)}{2}} \phi(S).$$

For any graph G we have that $0 < \phi(G), \alpha(G) \leq 1$, and $(\delta/\Delta) \cdot \phi(G) \leq \alpha(G) \leq \Delta \cdot \phi(G)$ [20].

3 Warm-up: Rumors Spread Fast from a Dominating Set

In this section we show the following result.

THEOREM 3.1. *Let $S \subseteq V$ be a dominating set of $G = (V, E)$, and suppose that all nodes $u \in S$ know the rumor initially. Then* PULL *informs all remaining nodes in $O(\log n)$ rounds, with probability $1 - O(n^{-\beta})$ for any constant $\beta > 0$.*

We start with an overview of the proof. By a standard lemma on the symmetry between PUSH and PULL (Lemma 3.2), we have that in order to bound the time for PULL to spread a rumor from S to a given node $u \notin S$, it suffices to bound instead the time for PUSH to spread a rumor from u to *some* node in S. We bound the latter time as follows (Lemma 3.3). Let I_t^u be the set of informed nodes after t rounds of PUSH, when the rumor starts from u. We consider the earliest round τ for which the expected growth of I_t^u in the next round $\tau + 1$ is smaller than by a constant factor $\epsilon > 0$, i.e., $\mathbf{E}[|I_{\tau+1}^u - I_\tau^u| \mid I_\tau^u] < \epsilon \cdot |I_\tau^u|$. We argue that $\tau = O(\log n)$ w.h.p. (Claim 3.4), which is intuitively clear, as up to round τ the number of informed nodes increases by a constant factor in expectation per round. Next we bound the harmonic mean of the degrees of nodes in I_τ^u; precisely, we show that $\sum_{v \in I_t^u} \frac{1}{\deg(v)} = \Omega(1)$ (Claim 3.5). If $I_\tau^u \cap S = \emptyset$ then each node from I_τ^u

has some neighbor in S, and from the above bound on the degrees it follows that in a round, the probability that some node from I_τ^u pushes the rumor to a neighbor in S is $\Omega(1)$. Thus, within $O(\log n)$ rounds after round τ, the rumor has reached some node in S w.h.p.

Next we give the detailed proof.

Lemma 3.2. *Let $T_{push}(V_1, V_2)$, for $V_1, V_2 \subseteq V$, be the number of* PUSH *rounds until a rumor initially known to all $u \in V_1$ (and only them) spreads to at least one $v \in V_2$; define $T_{pull}(V_1, V_2)$ similarly. For any $V_1, V_2 \subseteq V$, the random variables $T_{push}(V_1, V_2)$ and $T_{pull}(V_2, V_1)$ have the same distribution.*

The proof of Lemma 3.2 is essentially the same as that of [6, Lemma 3], and is therefore omitted. Next we state our main lemma; its proof is given in Section 3.1.

Lemma 3.3. *Let S be a dominating set and $u \in V - S$. Using* PUSH*, a rumor originated at u spreads to at least one node $v \in S$ in $O(\log n)$ rounds, with probability $1 - O(n^{-\beta})$ for any constant $\beta > 0$.*

Theorem 3.1 follows now easily: From Lemma 3.3 we obtain that a rumor originated at a given $u \in V - S$ spreads to at least one $v \in S$ after $O(\log n)$ rounds of PUSH, with probability $1 - O(n^{-\beta-1})$. Lemma 3.2 then implies that a rumor known to all $v \in S$ reaches u after $O(\log n)$ rounds of PULL, with the same probability, $1 - O(n^{-\beta-1})$. Applying now the union bound over all $u \in V - S$ yields the theorem.

3.1 Proof of Lemma 3.3.

Let I_t^u be the set of informed nodes after t rounds of PUSH, when the rumor starts from node $u \in V - S$. Let τ be the earliest round such that the expected increase of I_t^u in the next round is smaller than $\epsilon \cdot |I_t^u|$, i.e.,

$$\tau = \min\{t\colon \mathbf{E}[|I_{t+1}^u| \mid I_t^u] < (1+\epsilon) \cdot |I_t^u|\},$$

for some positive constant $\epsilon < 1$.

Claim 3.4. *With probability $1 - n^{-\beta}$, $\tau = O(\log n)$.*

Proof. Let X_t, for $t \geq 1$, be a 0/1 random variable that is 1 iff $|I_t^u| \geq (1+\epsilon/2) \cdot |I_{t-1}^u|$ or $t \geq \tau$. Then for any k,

$$\Pr(\tau \leq k) \geq \Pr\left(\sum_{t=1}^{k} X_t \geq \log_{1+\epsilon/2} n\right), \tag{3.1}$$

because: if $\tau > k$, then $|I_k^u| \geq (1+\epsilon/2)^{\sum_{t=1}^k X_t}$ and thus $\sum_{t=1}^k X_t \leq \log_{1+\epsilon/2}(|I_k^u|) < \log_{1+\epsilon/2} n$; and taking the contrapositive gives that $\sum_{t=1}^k X_t \geq \log_{1+\epsilon/2} n$ implies $\tau \leq k$.

Next we show that for any t,

$$\Pr(X_t = 1 \mid X_1 \ldots X_{t-1}) \geq \frac{\epsilon/2}{1-\epsilon/2}. \tag{3.2}$$

Fix the outcome of the first $t-1$ rounds, and suppose that $\tau > t-1$. (Otherwise, we have $X_t = 1$ and the inequality above holds trivially.) We bound $\Pr(X_t = 0)$ using Markov's Inequality:

$$\begin{aligned}\Pr(X_t = 0) &\leq \Pr(|I_t^u| < (1+\epsilon/2) \cdot |I_{t-1}^u|) \\ &= \Pr(2\,|I_{t-1}^u| - |I_t^u| > (1-\epsilon/2) \cdot |I_{t-1}^u|) \\ &\leq \frac{2\,|I_{t-1}^u| - \mathbf{E}[|I_t^u|]}{(1-\epsilon/2) \cdot |I_{t-1}^u|}, \quad \text{by Markov's Ineq.}\end{aligned}$$

Further, since $\tau > t-1$ we have from τ's definition that $\mathbf{E}[|I_t^u|] \geq (1+\epsilon) \cdot |I_{t-1}^u|$. Hence,

$$\Pr(X_t = 0) \leq \frac{2\,|I_{t-1}^u| - (1+\epsilon) \cdot |I_{t-1}^u|}{(1-\epsilon/2) \cdot |I_{t-1}^u|} = \frac{1-\epsilon}{1-\epsilon/2},$$

and thus $\Pr(X_t = 1) \geq \frac{\epsilon/2}{1-\epsilon/2}$.

From Eq. (3.2), it follows that $\sum_{t=1}^k X_t$ stochastically dominates binomial random variable $B\big(k, \frac{\epsilon/2}{1-\epsilon/2}\big)$. Thus, using Chernoff bounds we obtain $\Pr\big(\sum_{t=1}^k X_t \geq \log_{1+\epsilon/2} n\big) \geq \Pr\big(B\big(k, \frac{\epsilon/2}{1-\epsilon/2}\big) \geq \log_{1+\epsilon/2} n\big) \geq 1 - n^{-\beta}$ for $k \geq c \log n$, for a sufficiently large constant c. From this and Eq. (3.1), the claim follows. ∎

The next claim bounds the harmonic mean of the degrees of nodes $v \in I_\tau^u$.

Claim 3.5. *If $\mathbf{E}[|I_{t+1}^u| \mid I_t^u] < (1+\epsilon) \cdot |I_t^u|$ then $\sum_{v \in I_t^u} \frac{1}{\deg(v)} \geq 1 - \sqrt{\epsilon}$.*[3]

Proof. The proof is by contrapositive. Fix I_t^u and let $k = |I_t^u|$. For $1 \leq i \leq k$, let u_i be the i-th node in I_t^u, and let $d_i = \deg(u_i)$. We will assume that $\sum_{i=1}^k 1/d_i < 1 - \sqrt{\epsilon}$ and prove that $\mathbf{E}[|I_{t+1}^u|] \geq (1+\epsilon)k$.

We count the number of uninformed nodes that in round $t+1$ receive exactly one copy of the rumor. This is clearly a lower bound on the number of nodes that get informed in round $t+1$. Let $d_i' = |N(u_i) \cap I_t^u|$ be the number of neighbors that u_i has in I_t^u. The probability that in round $t+1$ node u_i pushes the rumor to some uninformed node is then $1 - d_i'/d_i$. And if this happens, the probability that the recipient node does not receive the rumor from any other node in the round is at least

$$\prod_{j \neq i}(1 - 1/d_j) \geq 1 - \sum_{j \neq i}(1/d_j) \geq 1 - (1 - \sqrt{\epsilon}) = \sqrt{\epsilon}.$$

[3] Inequality $\sum_{v \in I_t^u} \frac{1}{\deg(v)} \geq 1 - \sqrt{\epsilon}$ yields that the harmonic mean of $\deg(v)$, over all $v \in I_t^u$, is at most $|I_t^u|/(1-\sqrt{\epsilon})$.

Thus the probability that u_i sends the rumor to an uninformed node that does not receive another copy of the rumor is at least $(1-d'_i/d_i)\sqrt{\epsilon}$. Hence, the expected number of uninformed nodes that receive exactly one rumor copy in round $t+1$ is at least $\sum_{i=1}^{k}(1-d'_i/d_i)\sqrt{\epsilon}$. And since

$$\sum_{i=1}^{k}(1-d'_i/d_i) = k - \sum_{i=1}^{k}(d'_i/d_i) \geq k - \sum_{i=1}^{k}(k/d_i)$$
$$\geq k - k(1-\sqrt{\epsilon}) = k\sqrt{\epsilon},$$

we obtain a lower bound of $k\sqrt{\epsilon}\cdot\sqrt{\epsilon} = k\epsilon$ on the expected number of uninformed nodes that receive exactly one copy of the rumor in round $t+1$. Hence, the same lower bound holds for the total number on nodes informed in the round, and we conclude that $\mathbf{E}[|I^u_{t+1}|] \geq (1+\epsilon)k$. ∎

Using Claim 3.5 it is easy to show an $O(\log n)$ bound w.h.p. on the number of additional rounds after round τ, until the rumor spreads to at least one node from S: Suppose that $I^u_\tau = U$ for some set $U \subseteq V - S$. (If $U \not\subseteq V - S$ then some node from S is already informed.) Each $v \in U$ has at least one neighbor in the dominating set S, and thus the probability that none of these nodes pushes the rumor to a neighbor in S in a given round $t > \tau$, is upper bounded by

$$\prod_{v\in U}\left(1 - \frac{1}{\deg(v)}\right) \leq e^{-\sum_{v\in U}\frac{1}{\deg(v)}} \leq e^{-1+\sqrt{\epsilon}},$$

by Claim 3.5. This probability bound holds for each round $t > \tau$ independently of the outcome of previous rounds. It follows that $\ell := (\beta \ln n)/(1-\sqrt{\epsilon}) = O(\log n)$ additional rounds after round τ suffice to spread the rumor to a node in S with probability $1 - (e^{-1+\sqrt{\epsilon}})^\ell = 1 - n^{-\beta}$. Combining this with Claim 3.4, which bounds τ by $O(\log n)$ with probability $1 - n^{-\beta}$, and applying the union bound gives that the rumor spreads from u to at least one $v \in S$ in $O(\log n)$ rounds with probability $1 - 2n^{-\beta}$. This completes the proof of Lemma 3.3.

Another application of the same proof idea is described in Section A of the appendix.

4 Proof of the Main Result

In this section we show our main result, Theorem 1.1. We start with an overview of the proof. Let I_t denote the set of informed nodes after the first t rounds. We study the growth of the quantity $\Psi_t := |I_t| + |\partial I_t|/2 = (|I_t| + |I_t^+|)/2$. You can think of Ψ_t as a potential function: each informed node has as a potential of 1, each uninformed node with an informed neighbor has potential 1/2, and the remaining uninformed nodes have potential zero; Ψ_t is then the total potential after round t. We have $1 < \Psi_t \leq n$. To prove the theorem, we show that the expected number of rounds needed to double Ψ_t is bounded by $O(\log(\Delta)/\alpha)$, as long as $|I_t| \leq n/2$. It follows that $O(\log n \cdot \log(\Delta)/\alpha)$ rounds suffice w.h.p. to inform more than half of the nodes, and by a symmetry argument, $O(\log n \cdot \log(\Delta)/\alpha)$ additional rounds suffice to inform the remaining nodes.

For Ψ_t to double in $O(\log(\Delta)/\alpha)$ rounds, it suffices that it increases by $\Omega(|\partial I_t|/\log\Delta)$ "on average" per round, as $|\partial I_t| \geq \alpha \cdot |I_t|$ and thus $|\partial I_t| = \Omega(\alpha\Psi_t)$. Such an increase can be achieved either by informing $\Omega(|\partial I_t|/\log\Delta)$ nodes from ∂I_t, or by informing fewer nodes which however have a total number of $\Omega(|\partial I_t|/\log\Delta)$ neighbors in $\partial(I_t^+)$. Along this intuition, we distinguish the following two cases, in terms of a simple expansion measure we define for I_t, called boundary expansion (Definition 4.1).

The first case is when the boundary expansion of I_t is low (upper-bounded by a constant $\epsilon_h < 1$). This is the more challenging case, and is the core of our analysis. Our main lemma in this case is Lemma 4.3, which establishes that a constant fraction of the boundary ∂I_t gets informed in an expected number of $O(\log\Delta)$ rounds. The proof builds upon and extends the ideas used in the proof of Theorem 3.1. We note that in the setting of Theorem 3.1, the boundary expansion of the set S of informed nodes is zero.

The second case is when the boundary expansion of I_t is high (lower-bounded by a constant $\epsilon_h > 0$). Then from our definition of boundary expansion it follows that the expected number of nodes from $\partial(I_t^+)$ that have an informed neighbor after the next round is $\Omega(|\partial I_t|)$, i.e., $\mathbf{E}[|I^+_{t+1} - I^+_t| \mid I_t] = \Omega(|\partial I_t|)$. Our main lemma in this case is Lemma 4.10, which turns the above lower bound on the expected per round growth of I_t^+ into an upper bound on the expected number of rounds until I_t^+ grows by some quantity $b \cdot |\partial I_t|$, which depends on the degrees of nodes in ∂I_t.

Finally we bound the expected time needed to double Ψ_t (Lemma 4.12), by combining the results of the two cases above and using an inductive argument.

The rest of this section is structured as follows. We define the measure of boundary expansion in Section 4.1. In Section 4.2 we lower-bound the growth of I_t when boundary expansion is low. In Section 4.3 we lower-bound the growth of I_t^+ when boundary expansion is high. And in Section 4.4 we put the pieces together to complete the proof of Theorem 1.1.

4.1 Boundary Expansion.

Definition 4.1. *Let $S \subset V$ be a non-empty set. Let U be a random subset of ∂S such that each node $u \in \partial S$*

belongs to U *with probability* $1/\deg(u)$ *independently of the other nodes. The* boundary expansion $h(S)$ of S *is the ratio of the expected number of nodes* $v \in \partial(S^+)$ *that have some neighbor in* U*, over the size of* ∂S*, i.e.,* $h(S) = \mathbf{E}\left[|\{v \in \partial(S^+) \colon N(v) \cap U \neq \emptyset\}|\right] / |\partial S|$.

It is $0 \leq h(S) < 1$. The lower bound of 0 is matched iff S is a dominating set; and the upper bound holds because $\mathbf{E}\left[|\{v \in \partial(S^+) \colon N(v) \cap U \neq \emptyset\}|\right]$ is upper-bounded by the expected number of the edges between U and $\partial(S^+)$, which is $\sum_{u \in \partial S} |N(u) - S^+|/\deg(u) \leq \sum_{u \in \partial S} (\deg(u) - 1)/\deg(u) < |\partial S|$.

Intuitively, $h(S)$ is small if for many nodes $u \in \partial S$, we have that either just a small fraction of u's neighbors belong to $\partial(S^+)$, or, if u has a lot of neighbors in $\partial(S^+)$, then a large fraction of them are common neighbors to many nodes from ∂S.

If $I_t = S$ then $h(S) \cdot |\partial S|$ is a lower bound on the expected number of new nodes that have some informed neighbor after round $t + 1$, i.e.,

$$\mathbf{E}[|I^+_{t+1} - I^+_t| \mid I_t = S] \geq h(S) \cdot |\partial S|, \tag{4.3}$$

as each node $u \in \partial I_t$ pulls the rumor from I_t in round $t + 1$ with probability at least $1/\deg(u)$.

In Section 4.3 we will need the following refined definition, which describes the boundary expansion of S contributed by a given subset T of ∂S.

Definition 4.2. *Let* $S \subset V$ *and* $T \subseteq \partial S$*. Let* U_T *be a random subset of* T *such that each node* $u \in T$ *belongs to* U_T *with probability* $1/\deg(u)$ *independently of the other nodes. The* boundary expansion of S due to T *is* $h_T(S) = \mathbf{E}\left[|\{v \in \partial(S^+) \colon N(v) \cap U_T \neq \emptyset\}|\right] / |\partial S|$.

For $T = \partial S$ the above definition is identical to Definition 4.1, i.e., $h_{\partial S}(S) = h(S)$.

4.2 The Case of Low Boundary Expansion. In this section we prove the following result, which is the core lemma of our analysis.

Lemma 4.3. *Suppose that* $I_t = S$ *for some set* $S \subset V$ *with boundary expansion* $h(S) \leq \epsilon_h$*, where* $0 \leq \epsilon_h < 1$ *is an arbitrary constant. There is a constant* $\epsilon = \epsilon(\epsilon_h) > 0$ *such that the expected number of rounds until* $\epsilon \cdot |\partial S|$ *nodes from* ∂S *get informed is* $O(\log \Delta)$.

We start with an overview of the proof. Similar to the proof of Theorem 3.1, to bound the time for a given node $u \in \partial S$ to get informed, we bound instead the time for a rumor originated at u to spread to some node from S. Establishing this bound, however, is more difficult now than in the setting of Theorem 3.1. Recall that in the proof of Theorem 3.1, to bound the time until a rumor from $u \in \partial S$ reaches S, we first bound the time when the set I^u_t of informed nodes stops growing by a constant expected factor in each round, and then bound the harmonic mean of the degrees of nodes in I^u_t at that time; the bound on the degrees implies that if $I^u_t \cap S = \emptyset$, then with large probability some node from I^u_t will send the rumor to a neighbor in S. This last statement depends critically on the assumption that S is a dominating set, and thus every node from I^u_t has a neighbor in S. This in not the case, however, in the current setting.

To tackle this problem we consider a "restricted" rumor spreading process, on an induced subgraph of G. We identify a set of nodes *participating* in rumor spreading (Definition 4.4), such that, intuitively, each participating node has at least a constant probability to contact or be contacted by another participating node in a round. Only nodes from S^+ or $\partial(S^+)$ can be participating. Participating nodes from S^+ are *active*, i.e., they initiate a connection to a random neighbor in each round, while participating nodes from $\partial(S^+)$ are *passive*, i.e., they *accept* connections from active neighbors but do not *initiate* connections to random neighbors. Using the assumption that S has low boundary expansion we show that at least a constant fraction of ∂S is participating (Lemma 4.5). For each participating node $u \in \partial S$ then we show that an expected number of $O(\log \Delta)$ rounds suffice for a rumor originated at u to spread to some node from S (Lemma 4.6). The proof of this last result follows a similar outline to that of Lemma 3.3, although it is a bit more involved.

Below, in Section 4.2.1 we formally define the set of participating nodes and show that a constant fraction of ∂S is participating. Then in Section 4.2.2 we bound the time for a rumor originated at a participating node from ∂S to reach S. Finally, we complete the proof of Lemma 4.3 in Section 4.2.3.

4.2.1 Participating Nodes.

Definition 4.4. *The set* P *of* participating *nodes is a maximal subset of* V *such that for any* $u \in P$ *and for* $A := P \cap S^+$*,*

$$\frac{|N(u) \cap P|}{\deg(u)} + \sum_{v \in N(u) \cap A} \frac{1}{\deg(v)} \geq \epsilon_p, \text{ if } u \in A; \tag{4.4}$$

$$\sum_{v \in N(u) \cap A} \frac{1}{\deg(v)} \geq \epsilon_p, \text{ if } u \in P - A, \tag{4.5}$$

where $0 < \epsilon_p < (1 - \epsilon_h)/3$ *is a constant;*[4] *set* A *is the set of* active *nodes, and* $P - A$ *is the set of* passive *nodes.*

[4] We will see that P is unique, although this is not essential for the analysis.

It follows from this definition that all participating nodes belong to $S^+ \cup \partial(S^+)$; the ones in S^+ are active, and those in $\partial(S^+)$ are passive. Intuitively, P is defined such that each participating node has at least a constant probability to contact or be contacted by another participating node in a round. The term $|N(u) \cap P|/\deg(u)$ in Eq. (4.4) is the probability that active node u chooses a participating neighbor in a round. In Eq. (4.5) we do not have this term because, as mentioned earlier, passive nodes do not initiate connections. The sum that is common in both equations adds the probabilities of the events that u is chosen by v, for all active neighbors v of u. If this sum is small (bounded by a constant), then it is of the same order as the probability that u is chosen by at least one of its active neighbors.

The set P can be obtained by a simple procedure that recursively removes from V all nodes that do not satisfy (4.4) or (4.5). (This proves also that P is unique.) Formally, we start with set $P_0 = V$. In the i-th step of the procedure we obtain set P_i by removing from P_{i-1} all nodes $u \in A_{i-1} := P_{i-1} \cap S^+$ for which $|N(u) \cap P_{i-1}|/\deg(u) + \sum_{v \in N(u) \cap A_{i-1}} 1/\deg(v) < \epsilon_p$, and all nodes $u \in P_{i-1} - A_{i-1}$ for which $\sum_{v \in N(u) \cap A_{i-1}} 1/\deg(v) < \epsilon_p$. Clearly, the procedure finishes after at most n steps. Let $P^* := P_i$ for the last step i. We argue now that $P^* = P$. Because of the maximality of P, it suffices to show that $P \subseteq P^*$: Suppose, for contradiction, that $P \not\subseteq P^*$, and consider the first step i for which $P \not\subseteq P_i$. Then, we have $P \subseteq P_{i-1}$ and $A \subseteq A_{i-1}$, and some node $u \in P$ is removed from P_{i-1} in step i. If $u \in A$ then it follows from (4.4) that $|N(u) \cap P_{i-1}|/\deg(u) + \sum_{v \in N(u) \cap A_{i-1}} 1/\deg(v) \geq \epsilon_p$, which contradicts the assumption that u is removed in step i. Similarly, if $u \in P - A$, then it follows from (4.5) that $\sum_{v \in N(u) \cap A_{i-1}} 1/\deg(v) \geq \epsilon_p$, which again contradicts that u is removed.

We observe that if a subset of V is used as the starting set P_0 (instead of having $P_0 = V$), then the set obtained by the procedure is a subset of P.

We will now show that at least a constant fraction of nodes $u \in \partial S$ is participating.

Lemma 4.5. $|\partial S \cap P| \geq \left(1 - \frac{\epsilon_h}{(1-\epsilon_p)(1-2\epsilon_p)}\right) \cdot |\partial S|$.[5]

Proof. We use a potential function argument. We consider the recursive procedure given above to obtain P, and define a potential Φ_i after each step i. This potential is a non-negative quantity. Further, for an appropriate choice of the initial set P_0 for the procedure, we achieve that the potential function is non-increasing. For this P_0 we show that the potential Φ_0 before the first step satisfies $\Phi_0 \leq \frac{\epsilon_h}{1-\epsilon_p} \cdot |\partial S|$. We also show that for each node $u \in A_{i-1}$ that is removed from P_{i-1} in step i, the potential drops by at least $1 - 2\epsilon_p$, and the removal of a node $u \notin A_{i-1}$ does not increase the potential. It follows that the total number of nodes $u \in S^+$ removed in all steps is bounded by $\Phi_0/(1 - 2\epsilon_p)$. Thus, the same bound holds for the nodes $u \in \partial S$ removed, i.e., $|\partial S| - |\partial S \cap P| \leq \Phi_0/(1 - 2\epsilon_p)$. Rearranging and using that $\Phi_0 \leq \frac{\epsilon_h}{1-\epsilon_p} \cdot |\partial S|$ yields the claim.

Next we fill in the details omitted in the outline above. We start with the definition of the potential function. Intuitively, the potential Φ_i after step i measures the probability "wasted" in connections between participating and non-participating nodes. It has two components. The first component, $\Phi_{i,1}$, is the sum over all $u \in A_i$ of the probability that u chooses a neighbor $v \notin P_i$; the second component, $\Phi_{i,2}$, is the sum over all $u \in S^+ - A_i$ of the probability that u would choose a neighbor $v \in P_i$ if u were active. We give two equivalent formulas for each of $\Phi_{i,1}, \Phi_{i,2}$, to be used later on.

$$\Phi_{i,1} = \sum_{u \in A_i} \sum_{v \in N(u) - P_i} \frac{1}{\deg(u)} = \sum_{u \notin P_i} \sum_{v \in N(u) \cap A_i} \frac{1}{\deg(v)};$$

$$\Phi_{i,2} = \sum_{u \in S^+ - A_i} \sum_{v \in N(u) \cap P_i} \tfrac{1}{\deg(u)} = \sum_{u \in P_i} \sum_{v \in N(u) \cap (S^+ - A_i)} \tfrac{1}{\deg(v)}.$$

Then, $\Phi_i = \Phi_{i,1} + \Phi_{i,2}$.

The starting set P_0 we use consists of all nodes $u \in S^+$ plus those $u \in \partial(S^+)$ for which

$$\sum_{v \in N(u) \cap \partial S} \frac{1}{\deg(v)} \geq 2\epsilon_p. \tag{4.6}$$

The above condition is similar to (4.5), but the threshold is twice that in (4.5). We show now that

$$\Phi_0 \leq \frac{\epsilon_h}{1 - \epsilon_p} \cdot |\partial S|. \tag{4.7}$$

Let B be the set of nodes from $\partial(S^+)$ that do not satisfy (4.6), i.e.,

$$B = \left\{ u \in \partial(S^+) \colon \sum_{v \in N(u) \cap \partial S} \frac{1}{\deg(v)} < 2\epsilon_p \right\}.$$

We have $P_0 = S^+ \cup (\partial(S^+) - B)$ and $A_0 = S^+$. Substituting these values to the second formula for $\Phi_{i,1}$ and the first one for $\Phi_{i,2}$ yields

$$\Phi_0 = \Phi_{0,1} + \Phi_{0,2} = \sum_{u \in B} \sum_{v \in N(u) \cap \partial S} \frac{1}{\deg(v)} + 0. \tag{4.8}$$

[5]From the assumption in Definition 4.4 that $\epsilon_p < (1 - \epsilon_h)/3$, it follows that $1 - \frac{\epsilon_h}{(1-\epsilon_p)(1-2\epsilon_p)} > 0$.

Next we relate the double sum above with ϵ_h. From the definition of $h(S)$ (Definition 4.1), it follows

$$h(S) = \frac{1}{|\partial S|} \sum_{v \in \partial(S^+)} \left(1 - \prod_{u \in N(v) \cap \partial S} \left(1 - \frac{1}{\deg(u)}\right)\right).$$

Since the left side is at most ϵ_h, and replacing $\partial(S^+)$ by $B \subseteq \partial(S^+)$ on the right side can only decrease the sum's value, it follows

$$\epsilon_h \geq \frac{1}{|\partial S|} \sum_{u \in B} \left(1 - \prod_{v \in N(u) \cap \partial S} \left(1 - \frac{1}{\deg(v)}\right)\right).$$

Further, for any $u \in B$, the product on the right side is

$$\begin{aligned}
&\prod_{v \in N(u) \cap \partial S} \left(1 - \frac{1}{\deg(v)}\right) \leq e^{-\sum_{v \in N(u) \cap \partial S} \frac{1}{\deg(v)}} \\
&\quad \leq 1 - \sum_{v \in N(u) \cap \partial S} \frac{1}{\deg(v)} + \frac{1}{2} \left(\sum_{v \in N(u) \cap \partial S} \frac{1}{\deg(v)}\right)^2 \\
&\quad \leq 1 - \left(\sum_{v \in N(u) \cap \partial S} \frac{1}{\deg(v)}\right)(1 - \epsilon_p),
\end{aligned}$$

as $\sum_{v \in N(u) \cap \partial S} 1/\deg(v) < 2\epsilon_p$ for $u \in B$. It follows

$$\epsilon_h \geq \frac{1 - \epsilon_p}{|\partial S|} \sum_{u \in B} \left(\sum_{v \in N(u) \cap \partial S} \frac{1}{\deg(v)}\right) \overset{(4.8)}{=} \frac{(1 - \epsilon_p)\Phi_0}{|\partial S|},$$

and rearranging yields Eq. (4.7).

It remains to show that for each node $u \in A_{i-1}$ that is removed in step i, the potential decreases by at least $1 - 2\epsilon_p$, and for each node $u \in P_{i-1} - A_{i-1}$ removed the potential does not increase. W.l.o.g. we assume that only one node $u \in P_{i-1}$ is removed in step i. (If $k > 1$ nodes should be removed we just break step i into k sub-steps.)

First, suppose that a node $u \in P_{i-1} - A_{i-1}$ is removed in step i. Then $P_i = P_{i-1} - \{u\}$ and $A_i = A_{i-1}$. From the second formulas for $\Phi_{i,1}$ and $\Phi_{i,2}$ we obtain

$$\Phi_{i,1} - \Phi_{i-1,1} = \sum_{v \in N(u) \cap A_i} \frac{1}{\deg(v)} =: \phi_{inc},$$

$$\Phi_{i,2} - \Phi_{i-1,2} = - \sum_{v \in N(u) \cap (S^+ - A_{i-1})} \frac{1}{\deg(v)} =: -\phi_{dcr}.$$

Since $A_i = A_{i-1}$ we have

$$\phi_{inc} + \phi_{dcr} = \sum_{v \in N(u) \cap S^+} \frac{1}{\deg(v)} \geq 2\epsilon_p,$$

where the inequality holds because of Eq. (4.6), since $u \in \partial(S^+)$. Further, since u is removed in step i, it satisfies the condition $\sum_{v \in N(u) \cap A_{i-1}} 1/\deg(v) < \epsilon_p$, which yields $\phi_{inc} < \epsilon_p$. From the above two inequalities on ϕ_{inc} and ϕ_{dcr} we get $\phi_{inc} < \epsilon_p < \phi_{dcr}$, and thus, $\Phi_i - \Phi_{i-1} = \phi_{inc} - \phi_{dcr} < 0$.

Suppose now that a node $u \in A_{i-1}$ is removed. Then $P_i = P_{i-1} - \{u\}$ and $A_i = A_{i-1} - \{u\}$, and

$$\Phi_{i,1} - \Phi_{i-1,1} = - \underbrace{\sum_{v \in N(u) - P_{i-1}} \tfrac{1}{\deg(u)}}_{\phi_1} + \underbrace{\sum_{v \in N(u) \cap A_i} \tfrac{1}{\deg(v)}}_{\phi_2},$$

$$\Phi_{i,2} - \Phi_{i-1,2} = \underbrace{\sum_{v \in N(u) \cap P_i} \tfrac{1}{\deg(u)}}_{\phi_3} - \sum_{v \in N(u) \cap (S^+ - A_{i-1})} \tfrac{1}{\deg(v)}.$$

In the expression for ϕ_2 we can replace A_i by A_{i-1}, as $u \notin N(u)$ and thus $N(u) \cap A_i = N(u) \cap A_{i-1}$. Similarly, in the expression for ϕ_3 we can replace P_i by P_{i-1}. It follows that $\phi_1 + \phi_3 = \sum_{v \in N(u)} 1/\deg(u) = 1$, and also $\phi_2 + \phi_3 < \epsilon_p$ as otherwise u would not be removed in step i. Thus,

$$\begin{aligned}
\Phi_i - \Phi_{i-1} &= (\Phi_{i,1} - \Phi_{i-1,1}) + (\Phi_{i,2} - \Phi_{i-1,2}) \\
&\leq -\phi_1 + \phi_2 + \phi_3 \leq -\phi_1 + 2\phi_2 + \phi_3 \\
&= -(\phi_1 + \phi_3) + 2(\phi_2 + \phi_3) < -1 + 2\epsilon_p.
\end{aligned}$$

This completes the proof of Lemma 4.5. ∎

4.2.2 Spreading a Rumor from an Active Node.

We prove now the following lemma, which bounds the expected time until a rumor originated at some active node $u \in \partial S$ reaches S.

LEMMA 4.6. *Let $u \in \partial S \cap P$. Using* `PUSH-PULL`*, a rumor originated at u spreads to at least one node $v \in S$ in an expected number of $O(\log \Delta)$ rounds.*

This result is similar to Lemma 3.3, but holds only for active nodes rather than all $u \in \partial S$, and assumes `PUSH-PULL` rather than just `PUSH`. The proof analyzes the spread of u's rumor on the subgraph induced by the set P of participating nodes. Each active node (i.e., from set $A = P \cap S^+$) chooses a random neighbor in each round, and contacts that neighbor if it is participating. Passive nodes (i.e, from set $P - A = P \cap \partial(S^+)$) do not choose neighbors; they communicate only with the active nodes that contact them in each round.

Let I_t^u be the set of informed nodes after t rounds. Similar to the proof of Lemma 3.3, we define

$$\tau = \min\{t \colon \mathbf{E}[|I_{t+1}^u| \mid I_t^u] < (1 + \varepsilon) \cdot |I_t^u| \ \vee \ |I_t^u| \geq \Delta^2\},$$

where $0 < \varepsilon < \epsilon_p/2$ is a constant. The new condition $|I_t^u| \geq \Delta^2$ is added because we want to show a bound of

$O(\log \Delta)$, instead of $O(\log n)$ as in Lemma 3.3. The square in Δ is to ensure that if at least Δ^2 nodes are informed then at least Δ of them are active, as we elaborate later. The next result is an analogue of Claim 3.4, and bounds the expectation of τ.

CLAIM 4.7. $\mathbf{E}[\tau] = O(\log \Delta)$.

Proof. Let X_t, for $t \geq 1$, be a 0/1 random variable that is 1 iff $|I_t^u| \geq (1+\varepsilon/3) \cdot |I_{t-1}^u|$ or $t \geq \tau$. Further, let

$$\tau' = \min\left\{ i\colon \sum_{t=1}^{i} X_t \geq 2\log_{1+\varepsilon/3}\Delta \right\}.$$

We have

$$\mathbf{E}[\tau] \leq \mathbf{E}[\tau'], \tag{4.9}$$

because: for any k, if $\tau > k$ then $|I_k^u| \geq (1+\varepsilon/3)^{\sum_{t=1}^{k} X_t}$ and $|I_t^u| < \Delta^2$, and thus,

$$\sum_{t=1}^{k} X_t \leq \log_{1+\varepsilon/3}(|I_k^u|) < \log_{1+\varepsilon/3}(\Delta^2) = 2\log_{1+\varepsilon/3}\Delta,$$

which implies $\tau' > k$. Hence, for any k we have that $\tau > k$ implies $\tau' > k$, and thus $\mathbf{E}[\tau] \leq \mathbf{E}[\tau']$.

Next we show that for any t,

$$\Pr(X_t = 1 \mid X_1 \dots X_{t-1}) \geq 1 - e^{-\varepsilon^2/18}. \tag{4.10}$$

Fix the outcome of the first $t-1$ rounds, and suppose that $\tau > t-1$. (Otherwise, $X_t = 1$ and the inequality above holds trivially.) We have

$$\Pr(X_t = 0) \leq \Pr(|I_t^u| < (1+\varepsilon/3) \cdot |I_{t-1}^u|).$$

We will bound now the probability on the right side. Observe that $|I_t^u| = \sum_{v \in V} Y_v$, where Y_v is the indicator variable of the event that $v \in I_t^u$. The random variables Y_v are not independent: if two participating nodes $v, v' \notin I_{t-1}^u$ have a common neighbor $w \in I_{t-1}^u$, then the events that w pushes the rumor to v or to v' in round t are correlated. However, the variables Y_v are *negatively associated* [11]; this follows from [11, Example 4.5].[6] Negative association allows us to use standard Chernoff bounds to bound $\Pr(|I_t^u| < (1+\varepsilon/3) \cdot |I_{t-1}^u|)$: Since we have $\tau > t-1$, it follows from τ's definition that $\mathbf{E}[|I_t^u|] \geq (1+\varepsilon) \cdot |I_{t-1}^u|$, and

$$\begin{aligned} &\Pr(|I_t^u| < (1+\varepsilon/3) \cdot |I_{t-1}^u|) \\ &\qquad \leq \Pr\left(|I_t^u| < (1+\varepsilon/3) \cdot \frac{\mathbf{E}[|I_t^u|]}{1+\varepsilon}\right) \\ &\qquad \leq \Pr(|I_t^u| < (1-\varepsilon/3)\,\mathbf{E}[|I_t^u|]) \\ &\qquad \leq e^{-(\varepsilon/3)^2\,\mathbf{E}[|I_t^u|]/2} \leq e^{-\varepsilon^2/18}. \end{aligned}$$

Hence, $\Pr(X_t = 1) = 1 - \Pr(X_t = 0) \geq 1 - e^{-\varepsilon^2/18}$.

From Eq. (4.10) it follows that τ' is stochastically dominated by a negative-binomial random variable that counts the number of independent bernoulli trials until we have $2\log_{1+\varepsilon/3}\Delta$ successes, and each trial has a success probability of $1 - e^{-\varepsilon^2/18}$. Therefore, $\mathbf{E}[\tau'] \leq (2\log_{1+\varepsilon/3}\Delta)/(1-e^{-\varepsilon^2/18}) = O(\log \Delta)$. From this and Eq. (4.9), the claim follows. ∎

The next result is an analogue of Claim 3.5, and bounds the harmonic mean of the degrees of *active* nodes in I_τ^u

CLAIM 4.8. *If* $\mathbf{E}[|I_{t+1}^u| \mid I_t^u] < (1+\varepsilon) \cdot |I_t^u|$ *or* $|I_t^u| \geq \Delta^2$, *then* $\sum_{v \in I_t^u \cap S^+} \frac{1}{\deg(v)} \geq \zeta$, *where* $\zeta = (\epsilon_p - 2\varepsilon)/3$.

Proof. First we show that if $|I_t^u| \geq \Delta^2$ holds then $\sum_{v \in I_t^u \cap S^+} 1/\deg(v) \geq \zeta$—this is the easier case. Observe that each informed inactive node $w \in I_t^u \cap \partial(S^+)$ has at least one informed neighbor v in ∂S, namely, the one that pushed the rumor to w. Further, each $v \in \partial S$ has at most $\deg(v) - 1 \leq \Delta - 1$ neighbors in $\partial(S^+)$ (and at least one in S). It follows that if $|I_t^u| \geq \Delta^2$, then at least Δ of the nodes in I_t^u belong to S^+, i.e., $|I_t^u \cap S^+| \geq \Delta$, because otherwise, $|I_t^u| = |I_t^u \cap S^+| + |I_t^u \cap \partial(S^+)| \leq |I_t^u \cap S^+| + |I_t^u \cap S^+| \cdot (\Delta - 1) = |I_t^u \cap S^+| \cdot \Delta < \Delta^2$. And if we have $|I_t^u \cap S^+| \geq \Delta$, then

$$\sum_{v \in I_t^u \cap S^+} 1/\deg(v) \geq |I_t^u \cap S^+|/\Delta \geq \Delta/\Delta = 1 > \zeta.$$

It remains to show that if $\mathbf{E}[|I_{t+1}^u| \mid I_t^u] < (1+\varepsilon) \cdot |I_t^u|$ then $\sum_{v \in I_t^u \cap S^+} 1/\deg(v) \geq \zeta$. In the rest of the proof we assume that $\sum_{v \in I_t^u \cap S^+} 1/\deg(v) < \zeta$, and will prove that $\mathbf{E}[|I_{t+1}^u| \mid I_t^u] \geq (1+\varepsilon) \cdot |I_t^u|$.

The proof builds upon the ideas used for Claim 3.5. Similar to Claim 3.5, we count the number of uninformed nodes to which exactly one copy of the rumor is *pushed* in round $t+1$ (these nodes may also receive a second copy via pull). Further, we count the number of uninformed nodes that *pull* the rumor in that round. The sum of the above two numbers is then a lower bound on twice the total number of nodes informed in round

[6] The example cited refers to a general balls and bins model. In our case: bins are the nodes, and for each informed node v we place a ball to a random neighbor of v, while for each uninformed node v we place a ball to v iff a randomly chosen neighbor of v is informed. The example then shows that the number of balls in the bins are negatively associated.

$t+1$. We lower-bound the expectation of this sum using the definition of active and passive nodes.

Fix I_t^u. For each informed active node $v \in I_t^u \cap S^+$, let $\beta(v) = |N(v) \cap (P - I_t^u)|/\deg(v)$ be the fraction of v's neighbors that are participating and uninformed. Then the probability that v pushes the rumor to such a neighbor and no other node pushes the rumor to the same neighbor is lower-bounded by

$$\beta(v) \cdot \prod_{v' \in I_t^u \cap S^+} \left(1 - \tfrac{1}{\deg(v')}\right) \geq \beta(v) \cdot \left(1 - \sum_{v' \in I_t^u \cap S^+} \tfrac{1}{\deg(v')}\right) \geq \beta(v) \cdot (1 - \zeta).$$

Further, for each informed node $v \in I_t^u$, let $\gamma(v) = \sum_{v' \in N(v) \cap (A - I_t^u)} 1/\deg(v')$ be the expected number of uninformed active nodes that pull the rumor from v. It follows that the expected total number of nodes that get informed in round $t+1$ is

(4.11)
$$\mathbf{E}\left[|I_{t+1}^u| - |I_t^u|\right] \geq \frac{1}{2} \sum_{v \in I_t^u \cap S^+} \beta(v) \cdot (1 - \zeta) + \frac{1}{2} \sum_{v \in I_t^u} \gamma(v).$$

Next we bound the above sums of $\beta(v)$ and $\gamma(v)$ using Definition 4.4. For each informed active node $v \in I_t^u \cap S^+$, we have

$$\frac{|N(v) \cap P|}{\deg(v)} = \frac{|N(v) \cap I_t^u|}{\deg(v)} + \beta(v) \leq \frac{|I_t^u|}{\deg(v)} + \beta(v),$$

and

$$\sum_{v' \in N(v) \cap A} \frac{1}{\deg(v')} = \sum_{v' \in N(v) \cap (A \cap I_t^u)} \frac{1}{\deg(v')} + \gamma(v) \leq \sum_{v' \in A \cap I_t^u} \frac{1}{\deg(v')} + \gamma(v) < \zeta + \gamma(v).$$

Then (4.4) yields $|I_t^u|/\deg(v) + \beta(v) + \zeta + \gamma(v) \geq \epsilon_p$. Summing over all $v \in I_t^u \cap S^+$ and rearranging gives

$$\sum_{v \in I_t^u \cap S^+} \big(\beta(v) + \gamma(v)\big) \geq |I_t^u \cap S^+| \cdot (\epsilon_p - \zeta) - \sum_{v \in I_t^u \cap S^+} \frac{|I_t^u|}{\deg(v)} \geq |I_t^u \cap S^+| \cdot (\epsilon_p - \zeta) - |I_t^u| \cdot \zeta.$$

Next, for each informed passive node $v \in I_t^u - S^+$, we obtain similarly using (4.5) that $\zeta + \gamma(v) \geq \epsilon_p$, and summing over all such v gives

$$\sum_{v \in I_t^u - S^+} \gamma(v) \geq |I_t^u - S^+| \cdot (\epsilon_p - \zeta).$$

Adding the last two inequalities above yields

$$\sum_{v \in I_t^u \cap S^+} \beta(v) + \sum_{v \in I_t^u} \gamma(v) \geq |I_t^u| \cdot (\epsilon_p - \zeta) - |I_t^u| \cdot \zeta = |I_t^u| \cdot (\epsilon_p - 2\zeta).$$

From this and (4.11) it follows

$$\mathbf{E}[|I_{t+1}^u| - |I_t^u|] \geq \frac{1}{2}(1 - \zeta) \cdot |I_t^u| \cdot (\epsilon_p - 2\zeta) \geq \varepsilon \cdot |I_t^u|,$$

where the last inequality is obtained using that $\zeta = (\epsilon_p - 2\varepsilon)/3$. This completes the proof of Claim 4.8. ■

Using Claim 4.8 it is easy to show an $O(1)$ bound on the expected number of additional rounds after round τ, until some node in S gets informed: Fix I_τ^u and suppose that $I_\tau^u \cap S = \emptyset$ (otherwise some node from S is already informed). Then $I_\tau^u \cap S^+ = I_\tau^u \cap \partial S$, and Claim 4.8 gives $\sum_{v \in I_t^u \cap \partial S} 1/\deg(v) \geq \zeta$. Since each node $v \in I_t^u \cap \partial S$ has at least one neighbor from S, the probability that none of these nodes pushes the rumor to a neighbor from S in a given round $t > \tau$, is at most $\prod_{v \in I_t^u \cap \partial S} (1 - 1/\deg(v)) \leq e^{-\sum_{v \in I_t^u \cap \partial S} 1/\deg(v)} \leq e^{-\zeta}$. It follows that the expected number of rounds until some of the nodes $v \in I_\tau^u \cap \partial S$ pushes the rumor to a node in S is at most $1/(1 - e^{-\zeta}) = O(1)$. Combining this with Claim 4.7 gives an upper bound of $O(\log \Delta)$ on the expected number of rounds until some node from S gets informed, concluding the proof of Lemma 4.6.

4.2.3 Completing the Proof of Lemma 4.3.

The next standard lemma is the analogue of Lemma 3.2 for `PUSH-PULL`. Versions of this result can be found, e.g., in [6, 20, 2].

LEMMA 4.9. *Let $T(V_1, V_2)$ be the number of* `PUSH-PULL` *rounds until a rumor initially known to all $u \in V_1$ (and only them) spreads to at least one $v \in V_2$. Then $T(V_1, V_2)$ and $T(V_2, V_1)$ have the same distribution.*

From the lemma above and Lemma 4.6, it follows that the rumor spreads from S to a given $u \in \partial S \cap P$ in an expected number of at most $\ell = O(\log \Delta)$ rounds. Markov's Inequality then gives that in 2ℓ rounds node u has been informed with probability at least $1/2$, and from the linearity of expectation, in 2ℓ rounds at least $1/2$ of the nodes from $\partial S \cap P$ have been informed in expectation. Using Markov's again we obtain that in 2ℓ rounds more than $3/4$ of the nodes from $\partial S \cap P$ are still uninformed with probability at most $(1/2)/(3/4) = 2/3$. Thus, $1/4$ of the nodes from $\partial S \cap P$ get informed within 2ℓ rounds with probability at least $1/3$, and this implies that $1/4$ of the nodes from $\partial S \cap P$ get informed in an expected number of most $2\ell/(1/3) = 6\ell$ steps. From this and Lemma 4.5, which states that $|\partial S \cap P| = \Theta(|\partial S|)$, we obtain Lemma 4.3.

4.3 The Case of High Boundary Expansion. In this section we show the following result.

Lemma 4.10. *Suppose that $I_t = S$ for some set $S \subset V$ with boundary expansion $h(S) \geq \epsilon_h$, where $\epsilon_h > 0$ is an arbitrary constant. There is a positive real number $b = b(S) \leq 1/\alpha$ such that the expected number of rounds until $b \cdot |\partial S|$ nodes from $\partial(S^+)$ have some informed neighbor is $O(b \log \Delta)$.*

We start with an overview of the proof. Eq. (4.3) gives that the expected number of nodes from $\partial(S^+)$ that have some informed neighbor after one round is at least $h(S) \cdot |\partial S| \geq \epsilon_h \cdot |\partial S| = \Omega(|\partial S|)$. To prove the lemma we need to bound also the variance of those nodes. Intuitively, the variance will be higher when the degrees of nodes in ∂S are larger, and in this case larger values for b are needed to satisfy the lemma.

The proof distinguishes three cases. The first case is when the larger contribution to the boundary expansion of S is from nodes $u \in \partial S$ of degree $\deg(u) \leq c \cdot |\partial S|$, for some constant c. Formally, we have $h_T(S) \geq \epsilon_h/3$, for $T = \{u \in \partial S\colon \deg(u) \leq c \cdot |\partial S|\}$ (see Definition 4.2). Using the second moment method, we show that after one round of PULL we have with probability $\Omega(1)$ that $\Omega(|\partial S|)$ nodes from $\partial(S^+)$ have some informed neighbor. It follows that the expected number of rounds until $\Theta(|\partial S|)$ nodes from $\partial(S^+)$ have some informed neighbor is $O(1)$, thus the lemma holds for $b = \Theta(1)$.

The next case is when the larger contribution to $h(S)$ comes from nodes $u \in \partial S$ of degree between $c \cdot |\partial S|$ and $|S|$. We argue that for some degree k from that range, the number of nodes $u \in \partial S$ with degree $k \leq \deg(u) \leq 2k$ and $\Theta(k)$ neighbors in $\partial(S^+)$ is at least $\Omega(|\partial S|/\log \Delta)$. Then the probability of informing at least one such u in a round of PULL is $p = \Omega(|\partial S|/(k \log \Delta))$. It follows that the expected number of rounds until $\Theta(k)$ nodes from $\partial(S^+)$ have some informed neighbor is $1/p = O((k/|\partial S|) \log \Delta)$, thus the lemma holds for $b = \Theta(k/|\partial S|)$.

The last case is when the largest contribution to $h(S)$ is from nodes $u \in \partial S$ of degree $\deg(u) \geq d^*$, where $d^* = \max\{|S|, c|\partial S|\}$. We argue that $\Omega(|\partial S|)$ nodes $u \in \partial S$ have $\Omega(d^*)$ neighbors in $\partial(S^+)$. As the degree of those nodes u may be very large, we rely on push transmissions to inform them. Using that the degree of any node from S is bounded by $|S^+|$, we argue that in one round of PUSH, at least one of the above nodes u is informed with probability $p = \Omega(|\partial S|/|S^+|)$. It follows that the expected number of rounds until $\Theta(d^*) = \Theta(|S^+|)$ nodes from $\partial(S^+)$ have some informed neighbor is $1/p = O(|S^+|/|\partial S|)$, thus the lemma holds for $b = \Theta(|S^+|/|\partial S|)$.

We give the detailed proof next.

4.3.1 Proof of Lemma 4.10. We partition ∂S into three sets,

$$T_1 = \{u \in \partial S\colon \deg(u) \leq c \cdot |\partial S|\},$$
$$T_2 = \{u \in \partial S\colon c \cdot |\partial S| < \deg(u) \leq |S|\}\},$$
$$T_3 = \{u \in \partial S\colon \deg(u) > \max\{|S|, c \cdot |\partial S|\}\},$$

where $c = (\epsilon_h/3)^2/8$. From Definitions 4.1 and 4.2, it follows $h_{T_1}(S) + h_{T_2}(S) + h_{T_2}(S) \geq h(S) \geq \epsilon_h$, and thus, $h_{T_i}(S) \geq \epsilon_h/3$ for at least one $i \in \{1, 2, 3\}$.

The next result lower-bounds $|I^+_{t+1} - I^+_t|$, i.e., the number of nodes from $\partial(S^+)$ that have an informed neighbor after one round.

Claim 4.11. *Let $\varepsilon = \epsilon_h/3$.*

(a) If $h_{T_1}(S) \geq \varepsilon$ then $\Pr\left(|I^+_{t+1} - I^+_t| \geq \varepsilon|\partial S|/2\right) \geq \frac{1}{2}$.

(b) If $h_{T_2}(S) \geq \varepsilon$ then there is a $k \in \{c|\partial S|, \dots, |S|\}$ such that for $l = \log \frac{2\min\{|S|, \Delta\}}{\max\{c|\partial S|, \delta\}} \leq \log(2\Delta)$,

$$\Pr\left(|I^+_{t+1} - I^+_t| \geq \varepsilon k/2\right) \geq 1 - e^{-\frac{\varepsilon|\partial S|}{4kl}} = \Omega\left(\frac{|\partial S|}{kl}\right).$$

(c) If $h_{T_3}(S) \geq \varepsilon$ then for $k = \max\{|S|, c|\partial S|\}$,

$$\Pr\left(|I^+_{t+1} - I^+_t| \geq \varepsilon k/2\right) \geq 1 - e^{-\frac{\varepsilon c|\partial S|}{4k}} = \Omega\left(\frac{|\partial S|}{k}\right).$$

Proof. (a) Before round $t+1$, we fix for each node $u \in T_1$ a neighbor $v_u \in S$. Let U_{T_1} be the set of nodes $u \in T_1$ that get informed in round $t+1$ by pulling the rumor from their neighbor u_v. Further, let f_{T_1} be the number of nodes from $\partial(S^+)$ that have a neighbor in U_{T_1}. Clearly, $|I^+_{t+1} - I^+_t| \geq f_{T_1}$, thus to prove (a) it suffices to show that $\Pr(f_{T_1} \geq \varepsilon \cdot |\partial S|/2) \geq 1/2$. We show this next using Chebyshev's Inequality.

Each node $u \in T_1$ belongs to U_{T_1} with probability $1/\deg(u)$ independently of the other nodes, similar to U_T in Definition 4.2. It follows that $h_{T_1}(S) = \mathbf{E}[f_{T_1}]/|\partial S|$, thus

$$\mathbf{E}[f_{T_1}] = h_{T_1}(S) \cdot |\partial S| \geq \varepsilon \cdot |\partial S|.$$

For the variance of f_{T_1}, we will prove below that

$$\mathbf{Var}[f_{T_1}] \leq \sum_{u \in T_1} \deg(u). \qquad (4.12)$$

From this it follows

$$\mathbf{Var}[f_{T_1}] \leq |T_1| \cdot (c \cdot |\partial S|) \leq c \cdot |\partial S|^2 = \varepsilon^2 \cdot |\partial S|^2/8,$$

as $c = \varepsilon^2/8$. Chebyshev's Inequality then yields the desired result:

$$\Pr(f_{T_1} \leq \varepsilon \cdot |\partial S|/2) \leq \Pr\left(|f_{T_1} - \mathbf{E}[f_{T_1}]| \geq \varepsilon \cdot |\partial S|/2\right) \leq \frac{\mathbf{Var}[f_{T_1}]}{(\varepsilon \cdot |\partial S|/2)^2} \leq 1/2.$$

It remains to prove Eq. (4.12). For each $v \in \partial(S^+)$, let X_v be a 0/1 random variable that is 1 iff v has a neighbor in U_{T_1}. Then, $f_{T_1} = \sum_{v\in\partial(S^+)} X_v$, and

$$\mathbf{Var}[f_{T_1}] = \sum_{(v_1,v_2)\in(\partial(S^+))^2} \mathbf{Cov}[X_{v_1}, X_{v_2}].$$

The covariance term $\mathbf{Cov}[X_{v_1}, X_{v_2}]$ equals

$$\Pr(X_{v_1} = X_{v_2} = 1) - \Pr(X_{v_1} = 1)\cdot\Pr(X_{v_2} = 1).$$

We can express $\Pr(X_{v_1} = X_{v_2} = 1)$ as the sum of the following two terms: 1) the probability that v_1 and v_2 have a common neighbor in U_{T_1}; by the union bound, this is at most equal to $\sum_{u\in N(v_1)\cap N(v_2)\cap T_1} 1/\deg(u)$; and 2) the probability that each of v_1 and v_2 has a neighbor in U_{T_1} but they have no common neighbors in U_{T_1}; this is at most equal to $\Pr(X_{v_1} = 1)\cdot\Pr(X_{v_2} = 1)$. It follows $\mathbf{Cov}[X_{v_1}, X_{v_2}] \le \sum_{u\in N(v_1)\cap N(v_2)\cap T_1} 1/\deg(u)$, thus

$$\mathbf{Var}[f_{T_1}] \le \sum_{(v_1,v_2)\in(\partial(S^+))^2}\ \sum_{u\in N(v_1)\cap N(v_2)\cap T_1} 1/\deg(u).$$

For each node $u \in T_1$, the term $1/\deg(u)$ appears in the double sum above exactly $|N(u)\cap\partial(S^+)|^2$ times: once for each pair $(v_1, v_2) \in (N(u)\cap\partial(S^+))^2$. Thus,

$$\begin{aligned}\mathbf{Var}[f_{T_1}] &\le \sum_{u\in T_1} (|N(u)\cap\partial(S^+)|^2/\deg(u))\\ &\le \sum_{u\in T_1} (\deg(u)^2/\deg(u)) = \sum_{u\in T_1}\deg(u).\end{aligned}$$

(b) Let $T_2' = \{u \in T_2 \colon |N(u)\cap\partial(S^+)|/\deg(u) \ge \varepsilon/2\}$ be the set of nodes $u \in T_2$ with the property that an $(\varepsilon/2)$-fraction of u's neighbors belong to $\partial(S^+)$. We have $|T_2'| \ge \varepsilon\cdot|\partial S|/2$, because otherwise the assumption $h_{T_2}(S) \ge \varepsilon$ is contradicted:

$$\begin{aligned}h_{T_2}(S)\cdot|\partial S| &\le \sum_{u\in T_2}\frac{|N(u)\cap\partial(S^+)|}{\deg(u)}\\ &\le |T_2'|\cdot 1 + (|T_2| - |T_2'|)\cdot(\varepsilon/2)\\ &< \varepsilon\cdot|\partial S|/2 + |\partial S|\cdot(\varepsilon/2) = \varepsilon\cdot|\partial S|.\end{aligned}$$

Since nodes in T_2' have degrees in the range between $\max\{c|\partial S|, \delta\}$ and $\min\{|S|, \Delta\}$, it follows that for some k in this range, at least $|T_2'|/l$ nodes $u \in T_2'$ have degree $k \le \deg(u) \le 2k$. If at least one of these nodes gets informed in round $t+1$ then $|I_{t+1}^+ - I_t^+| \ge \varepsilon k/2$, and the probability that this happens is at least

$$1-(1-1/(2k))^{|T_2'|/l} \ge 1-e^{-|T_2'|/(2lk)} \ge 1-e^{-\varepsilon|\partial S|/(4lk)}.$$

(c) Similar to (b), we let $T_3' = \{u \in T_3 \colon |N(u)\cap\partial(S^+)|/\deg(u) \ge \varepsilon/2\}$, and it is $|T_3'| \ge \varepsilon\cdot|\partial S|/2$. If a node $u \in T_3'$ gets informed in round $t+1$, then we have $|I_{t+1}^+ - I_t^+| \ge \varepsilon k/2$, where $k = \max\{|S|, c\cdot|\partial S|\}$. Thus, to prove the claim it suffices to lower-bound the probability of informing at least one node $u \in T_3'$ in round $t+1$. Unlike the proofs for (a) and (b) which rely on pull transmissions of the rumor, we use push transmissions here. For each $u \in T_3'$, we fix a neighbor $v_u \in S$ of u before round $t+1$. The probability that the rumor is pushed from v_u to u is $1/\deg(v_u) \ge 1/|S^+|$. Further, if $i > 1$ nodes $u \in T_3'$ have the same v_u, the probability that none of them receives the rumor via a push from v_u is $1-i/\deg(v_u) \le (1-1/\deg(v_u))^i$, i.e., it is smaller than if the nodes u had distinct neighbors v_u. It follows that the probability at least one node $u \in T_3'$ receives the rumor via a push from its neighbor v_u is lower-bounded by

$$1-(1-1/|S^+|)^{|T_3'|} \ge 1-e^{-|T_3'|/|S^+|} \ge 1-e^{-\frac{\varepsilon|\partial S|/2}{2k/c}},$$

where for the last inequality we used that $|S^+| = |S| + |\partial S| \le k + k/c \le 2k/c$, as $c \le 1$. This completes the proof of Claim 4.11. ∎

From Claim 4.11, Lemma 4.10 follows easily: We can assume that before each round $i > t$ all informed nodes $u \notin S$ become uninformed, as this can only decrease the number of nodes $v \in \partial(S^+)$ with some informed neighbor after round i. With this assumption, the number of $v \in \partial(S^+)$ with an informed neighbor after round i becomes independent of the outcome of the previous rounds $t+1, \ldots, i-1$. From Claim 4.11(a) then we obtain that if $h_{T_1}(S) \ge \epsilon_h/3$, the probability that at least $\varepsilon\cdot|\partial S|/2$ nodes $v \in \partial(S^+)$ have an informed neighbor after a given round $i > t$ is at least 1/2. It follows that for $b = \varepsilon/2$ the expected number of rounds until $b\cdot|\partial S|$ nodes $v \in \partial(S^+)$ have an informed neighbor is at most $1/(1/2) = O(b)$. Similarly, Claim 4.11(b) yields that if $h_{T_2}(S) \ge \epsilon_h/3$, then for some $k \in \{c\cdot|\partial S|, \ldots, |S|\}$ and for $b = \varepsilon k/(2|\partial S|) < 1/\alpha$, the expected number of rounds until $b\cdot|\partial S|$ nodes $v \in \partial(S^+)$ have an informed neighbor is $1/\Omega(|\partial S|/kl) = O(bl)$. Finally, Claim 4.11(c) gives that if $h_{T_3}(S) \ge \epsilon_h/3$, then for $k = \max\{|S|, c\cdot|\partial S|\}$ and $b = \varepsilon k/(2|\partial S|) < 1/\alpha$, the expected number of rounds until $b\cdot|\partial S|$ nodes $v \in \partial(S^+)$ have an informed neighbor is $1/\Omega(|\partial S|/k) = O(b)$. This completes the proof of Lemma 4.10.

4.4 Completing the Proof of Theorem 1.1. We will now use Lemmas 4.3 and 4.10 to establish a lower bound on the growth of $\Psi_t = (|I_t| + |I_t^+|)/2$. We show that the expected number of rounds needed to double Ψ_t or increase $|I_t|$ above $n/2$ (whichever occurs first) is bounded by $O(\log(\Delta)/\alpha)$.

Lemma 4.12. *Let*

$$T_t = \min\{i\colon \Psi_{t+i} \geq 2\Psi_t \vee |I_{t+i}| > n/2\}.$$

Then $\mathbf{E}[T_t \mid I_t] \leq s\log(\Delta)/\alpha$, *for a fixed constant* $s > 0$.

Proof. The proof uses an inductive argument, and is based on the following corollary of Lemmas 4.3 and 4.10. Let I_i^S and Ψ_i^S denote respectively I_{t+i} and Ψ_{t+i} given that $I_t = S$.

Claim 4.13. *For any non-empty set* $S \subset V$, *there is a positive real number* $b_S \leq 1/\alpha$ *such that for*

$$\tau_S = \min\{i\colon \Psi_i^S \geq \Psi_0^S + b_S \cdot |\partial S|\},$$

we have $\mathbf{E}[\tau_S] \leq cb_S \log\Delta$, *for some constant* $c > 0$.

Proof. If $h(S) \leq \epsilon_h < 1$, then it follows from Lemma 4.3 that set I_i^S increases by $\Theta(|\partial S|)$ nodes—and thus Ψ_i^S increases by at least that amount—in an expected number of $O(\log\Delta)$ rounds; so the claim holds for $b_s = \Theta(1)$ in this case. If $h(S) \geq \epsilon_h > 0$, then it follows from Lemma 4.10 that $(I_i^S)^+$ increases by $b\cdot|\partial S|$ nodes—thus Ψ_i^S increases by at least $b\cdot|\partial S|/2$—in an expected number of $O(b\log\Delta)$ rounds, for some positive $b \leq 1/\alpha$; so the claim holds for $b_s = b/2$ in this case. ∎

Define

$$T(S,k) = \min\left\{i\colon \Psi_i^S \geq \Psi_0^S + k \vee |I_i^S| > n/2\right\}.$$

The quantity $\mathbf{E}[T_t \mid I_t = S]$ we must bound is then the same as $\mathbf{E}[T(S,\Psi_0^S)]$. We will prove that if $|S| \leq n/2$ and $0 < k \leq \Psi_0^S$, then

$$\mathbf{E}[T(S,k)] \leq \left(4 - \frac{3\Psi_0^S}{k+\Psi_0^S}\right)\cdot\frac{c\log\Delta}{\alpha}, \tag{4.13}$$

where c is the constant of Claim 4.13. Lemma 4.12 then follows as a special case of the above inequality: Setting $k = \Psi_0^S$ yields $\mathbf{E}[T(S,\Psi_0^S)] \leq 5c\log(\Delta)/(2\alpha)$, and since $\mathbf{E}[T(S,\Psi_0^S)] = \mathbf{E}[T_t \mid I_t = S]$ it follows that $\mathbf{E}[T_t \mid I_t = S] \leq s\log(\Delta)/\alpha$ for $s = 5c/2$.

It remains to show Eq. (4.13). The proof is by induction on k. We distinguish two cases.

Case 1: $k \leq b_S \cdot |\partial S|$. This is the base case of the inductive proof. Since $k \leq b_S \cdot |\partial S|$, we have that $T(S,k) \leq T(S, b_S\cdot|\partial S|) \leq \tau_S$, and from Claim 4.13, $\mathbf{E}[\tau_s] \leq cb_S\log\Delta \leq c\log(\Delta)/\alpha$. Thus

$$\mathbf{E}[T(S,k)] \leq c\log(\Delta)/\alpha.$$

Eq. (4.13) now follows, because $4 - \frac{3\Psi_0^S}{k+\Psi_0^S} \geq 1$ as $k \geq 0$.

Case 2: $k > b_S\cdot|\partial S|$. We divide $T(S,k)$ into two terms, the rounds until the potential has increase by $b_S\cdot|\partial S|$, and the remaining rounds until it has increased by k in total. Let $S' = I_{\tau_S}^S$. Then

$$T(S,k) = T(S, b_S\cdot|\partial S|) + T(S', k + \Psi_0^S - \Psi_0^{S'}).$$

The first term on the right side is at most τ_S, thus from Claim 4.13 it follows

$$\mathbf{E}[T(S, b_S\cdot|\partial S|)] \leq cb_S\log\Delta.$$

We can bound the expectation of the second term, $T(S', k+\Psi_0^S - \Psi_0^{S'})$, by applying the induction hypothesis, as $S \subsetneq S'$ and thus $k + \Psi_0^S - \Psi_0^{S'} < k$. Thus by applying (4.13) with S' and $k+\Psi_0^S-\Psi_0^{S'}$ in place of S and k, respectively, we obtain that if $k+\Psi_0^S-\Psi_0^{S'} > 0$,

$$\begin{aligned}\mathbf{E}[T(S', k+\Psi_0^S-\Psi_0^{S'}) \mid S'] &\leq \left(4 - \frac{3\Psi_0^{S'}}{k+\Psi_0^S}\right)\cdot\frac{c\log\Delta}{\alpha}\\ &\leq \left(4 - \frac{3(\Psi_0^S + b_S|\partial S|)}{k+\Psi_0^S}\right)\cdot\frac{c\log\Delta}{\alpha}.\end{aligned}$$

If $k+\Psi_0^S-\Psi_0^{S'} \leq 0$, the above expectation is 0, thus it is still bounded by the quantity in the second line. From the three equations above it follows

$$\mathbf{E}[T(S,k)] \leq \left(\alpha b_S + 4 - \frac{3\Psi_0^S}{k+\Psi_0^S} - \frac{3b_S|\partial S|}{k+\Psi_0^S}\right)\cdot\frac{c\log\Delta}{\alpha}.$$

To prove (4.13) it suffices to show that $\alpha b_S - \frac{3b_S|\partial S|}{k+\Psi_0^S} \leq 0$, or equivalently, that $\alpha(k+\Psi_0^S) \leq 3|\partial S|$: Using the assumption of (4.13) that $k \leq \Psi_0^S$, and the fact that $|\partial S|/|S| \geq \alpha$ as $|S| \leq n/2$, we get $\alpha(k+\Psi_0^S) \leq 2\alpha\Psi_0^S = 2\alpha(|S| + |\partial S|/2) \leq 2(|\partial S| + \alpha|\partial S|/2) \leq 3|\partial S|$.

This completes the proof of Eq. (4.13) and the proof of Lemma 4.12. ∎

We will now use Lemma 4.12 to derive Theorem 1.1. We partition rumor spreading in phases of $y = 2s\log(\Delta)/\alpha$ rounds each, and count the phases in which Ψ_t doubles or $|I_t|$ has exceeded $n/2$. Let X_i, for $i \geq 1$, be a 0/1 random variable that is 1 iff $T_{(i-1)y} \leq y$. From Lemma 4.12, $\mathbf{E}[T_{(i-1)y} \mid I_{(i-1)y}] \leq y/2$, and Markov's Ineq. gives $\Pr\left(T_{(i-1)y} \leq y \mid I_{(i-1)y}\right) \geq 1/2$. It follows $\Pr(X_i = 1 \mid X_1 \ldots X_{i-1}) \geq 1/2$, and from this, $\sum_{j\leq i} X_j$ stochastically dominates binomial random variable $B(i, 1/2)$. Then Chernoff bounds give for $i^* = 2(\beta+3)\log n$, $\Pr\left(\sum_{j\leq i^*} X_j < \log n\right) \leq e^{-2(i^*/2-\log n)^2/i^*} < n^{-\beta-1}$. This implies that i^*y rounds suffice to inform more than $n/2$ nodes with probability at least $1 - n^{-\beta-1}$: If $|I_{i^*y}| \leq n/2$ then fewer than $\log n$ among the $X_1 \ldots X_{i^*}$ are 1, as in this case

$\Psi_{i^*y} \geq 2^{\sum_{j \leq i^*} X_j}$, and $\Psi_{i^*y} < n$. Hence, we have $\Pr(|I_{i^*y}| \leq n/2) \leq \Pr(\sum_{j \leq i^*} X_j < \log n) < n^{-\beta-1}$.

Once a set $V_1 \subseteq V$ with $|V_1| > n/2$ has been informed, any given node $u \notin V_1$ gets informed within i^*y additional rounds with probability at least $1 - n^{-\beta-1}$: From Lemma 4.9, the probability that u learns the rumor from V_1 within i^*y rounds equals the probability that some node from V_1 learns a rumor originated at u within i^*y rounds. And the latter probability is at least equal to the probability that more than $n/2$ nodes learn u's rumor within i^*y rounds, because then at least one of these nodes will belong to V_1 as $|V_1| > n/2$.

From the above and the union bound, it follows that all nodes get informed in at most $2i^*y = O(\log n \cdot \log(\Delta)/\alpha)$ rounds with probability at least $1-n^{-\beta}$. This completes the proof of Theorem 1.1.

Acknowledgements. I would like to thank Keren Censor-Hillel and Thomas Sauerwald for helpful discussions on the motivation for this work.

References

[1] Stephen P. Boyd, Arpita Ghosh, Balaji Prabhakar, and Devavrat Shah. Randomized gossip algorithms. *IEEE Trans. Inf. Theory*, 52(6):2508–2530, 2006.

[2] Keren Censor-Hillel, Bernhard Haeupler, Jonathan A. Kelner, and Petar Maymounkov. Global computation in a poorly connected world: Fast rumor spreading with no dependence on conductance. In *Proc. 44th STOC*, pages 961–970, 2012.

[3] Keren Censor-Hillel and Hadas Shachnai. Partial information spreading with application to distributed maximum coverage. In *Proc. 27th PODC*, pages 161–170, 2010.

[4] Keren Censor-Hillel and Hadas Shachnai. Fast information spreading in graphs with large weak conductance. In *Proc. 22nd SODA*, pages 440–448, 2011.

[5] Flavio Chierichetti, Silvio Lattanzi, and Alessandro Panconesi. Rumor spreading in social networks. In *Proc. 36th ICALP*, pages 375–386, 2009.

[6] Flavio Chierichetti, Silvio Lattanzi, and Alessandro Panconesi. Almost tight bounds for rumour spreading with conductance. In *Proc. 42nd STOC*, pages 399–408, 2010.

[7] Flavio Chierichetti, Silvio Lattanzi, and Alessandro Panconesi. Rumour spreading and graph conductance. In *Proc. 21st SODA*, pages 1657–1663, 2010.

[8] Alan Demers, Dan Greene, Carl Hauser, Wes Irish, John Larson, Scott Shenker, Howard Sturgis, Dan Swinehart, and Doug Terry. Epidemic algorithms for replicated database maintenance. In *Proc. 6th PODC*, pages 1–12, 1987.

[9] Benjamin Doerr, Mahmoud Fouz, and Tobias Friedrich. Social networks spread rumors in sublogarithmic time. In *Proc. 43rd STOC*, pages 21–30, 2011.

[10] Pedro Domingos and Matthew Richardson. Mining the network value of customers. In *Proc. 7th KDD*, pages 57–66, 2001.

[11] Devdatt P. Dubhashi and Alessandro Panconesi. *Concentration of Measure for the Analysis of Randomized Algorithms*. Cambridge University Press, 2009.

[12] Robert Elsässer and Thomas Sauerwald. Broadcasting vs. mixing and information dissemination on Cayley graphs. In *Proc. 24th STACS*, pages 163–174, 2007.

[13] Stephen Eubank, Hasan Guclu, V.S. Anil Kumar, Madhav V. Marathe, Aravind Srinivasan, Zoltn Toroczkai, and Nan Wang. Modelling disease outbreaks in realistic urban social networks. *Nature*, 429:180–184, 2004.

[14] Uriel Feige, David Peleg, Prabhakar Raghavan, and Eli Upfal. Randomized broadcast in networks. *Random Struct. Algorithms*, 1(4):447–460, 1990.

[15] Nikolaos Fountoulakis, Anna Huber, and Konstantinos Panagiotou. Reliable broadcasting in random networks and the effect of density. In *Proc. 29th INFOCOM*, pages 2552–2560, 2010.

[16] Nikolaos Fountoulakis and Konstantinos Panagiotou. Rumor spreading on random regular graphs and expanders. In *Proc. 14th RANDOM*, pages 560–573, 2010.

[17] Nikolaos Fountoulakis, Konstantinos Panagiotou, and Thomas Sauerwald. Ultra-fast rumor spreading in social networks. In *Proc. 23rd SODA*, pages 1642–1660, 2012.

[18] Alan Frieze and Geoffrey Grimmett. The shortest-path problem for graphs with random arc-lengths. *Discrete Applied Math.*, 10:57–77, 1985.

[19] George Giakkoupis. Tight bounds for rumor spreading in graphs of a given conductance. In *Proc. 28th STACS*, pages 57–68, 2011.

[20] George Giakkoupis and Thomas Sauerwald. Rumor spreading and vertex expansion. In *Proc. 23rd SODA*, pages 1623–1641, 2012.

[21] Bernhard Haeupler. Simple, fast and deterministic gossip and rumor spreading. In *Proc. 24th SODA*, pages 705–716, 2013.

[22] Richard Karp, Christian Schindelhauer, Scott Shenker, and Berthold Vöcking. Randomized rumor spreading. In *Proc. 41st FOCS*, pages 565–574, 2000.

[23] David Kempe, Jon Kleinberg, and Éva Tardos. Maximizing the spread of influence through a social network. In *Proc. 9th KDD*, pages 137–146, 2003.

[24] Jure Leskovec, Kevin J. Lang, Anirban Dasgupta, and Michael W. Mahoney. Statistical properties of community structure in large social and information networks. In *Proc. 17th WWW*, pages 695–704, 2008.

[25] Damon Mosk-Aoyama and Devavrat Shah. Fast distributed algorithms for computing separable functions. *IEEE Trans. Inf. Theory*, 54(7):2997–3007, 2008.

[26] Boris Pittel. On spreading a rumor. *SIAM J. Applied Math.*, 47(1):213–223, 1987.

[27] Thomas Sauerwald and Alexandre Stauffer. Rumor spreading and vertex expansion on regular graphs. In

Proc. 22nd SODA, pages 462–475, 2011.

APPENDIX

A Rumors Spread Fast in Graphs of Diameter 2

In this section we give an example of another result that can be proved using the machinery developed for the proof of our main result. We show that PUSH-PULL completes in a logarithmic number of rounds in any graph of diameter (at most) 2. This result can be viewed as an extension to the classic result that rumor spreading takes logarithmic time in graphs of diameter 1, i.e., in complete graphs. We point out that unlike complete graphs, some graphs of diameter 2 have very low expansion, e.g., two cliques of the same size with one common vertex. The result does not extend to graphs of diameter 3, as for example, rumor spreading takes linear time in the dumbbell graph, which consist of two cliques of the same size and a single edge between one node from each clique.

THEOREM A.1. *For any graph $G = (V, E)$ of diameter 2,* PUSH-PULL *informs all nodes of G in $O(\log n)$ rounds with probability $1 - O(n^{-\beta})$, for any constant $\beta > 0$.*

Proof Sketch. We fix an arbitrary pair of nodes $u, v \in V$, and show that a rumor originated at u reaches v in $O(\log n)$ rounds w.h.p. We divide rumor spreading into three phases, each of length $c \log n$ for a sufficiently large constant c.

In the first phase we consider only push operations. We define τ as in the proof of Lemma 3.3, and from Claims 3.4 and 3.5 we obtain that w.h.p. in $O(\log n)$ rounds the rumor has spread to all nodes of some set $V_u \subseteq V$ for which $\sum_{u' \in V_u} 1/\deg(u') = \Omega(1)$.

In the last (third) phase we consider just pull operations. From the symmetry between push and pull, and the argument used for the first phase, it follows that w.h.p. there is a set $V_v \subseteq V$ with $\sum_{v' \in V_v} 1/\deg(v') = \Omega(1)$, such that if some node from V_v knows the rumor at the beginning of the last phase, then v learns the rumor as well within the next $O(\log n)$ rounds of the phase. (The set V_v depends only on the random choices of nodes in the rounds of the phase.)

Suppose now that sets V_u and V_v as above exist (this is true w.h.p.), and fix all random choices in the first and third phases (and thus sets V_u and V_v). We will argue that in the second phase the rumor spreads from V_u to some node in V_v w.h.p. This is trivially true if $V_u \cap V_v \neq \emptyset$, hence assume $V_u \cap V_v = \emptyset$. From the assumption that the graph has diameter 2, it follows that for every node $v' \in V_v$, at least one of the next two conditions holds:

1. Node v' has a neighbor in V_u; or
2. For every $u' \in V_u$ there is a node $w_{u'v'} \notin V_u \cup V_v$ that is a common neighbor of u' and v'.

If Condition 1 holds for all $v' \in V_v$, and thus each $v' \in V_v$ has some informed neighbor at the beginning of the second phase, then from inequality $\sum_{v' \in V_v} 1/\deg(v') = \Omega(1)$ it follows that some $v' \in V_v$ will pull the rumor from an informed neighbor w.h.p. within $O(\log n)$ rounds. Suppose now that Condition 1 does not hold for some $v' \in V_v$, thus Condition 2 must hold. Then from inequality $\sum_{u' \in V_u} 1/\deg(u') = \Omega(1)$ it follows that w.h.p. at least one of the nodes $u' \in V_u$ will push the rumor to its neighbor $w_{u'v'}$ within $O(\log n)$ rounds. And thus v' will have some informed neighbor after that. Hence, after $O(\log n)$ rounds in the second phase, w.h.p. all $v' \in V_v$ have some informed neighbor. It follows then that some $v' \in V_v$ will pull the rumor in an additional $O(\log n)$ rounds w.h.p., as we argued earlier. ∎