

Rumor Spreading and Conductance

Flavio Chierichetti, George Giakkoupis, Silvio Lattanzi, Alessandro Panconesi

► **To cite this version:**

Flavio Chierichetti, George Giakkoupis, Silvio Lattanzi, Alessandro Panconesi. Rumor Spreading and Conductance. Journal of the ACM (JACM), Association for Computing Machinery, 2018, 65 (4), pp.17:1-17:21. 10.1145/3173043 . hal-01942162

HAL Id: hal-01942162

<https://hal.inria.fr/hal-01942162>

Submitted on 3 Dec 2018

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

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

Rumor Spreading and Conductance

Flavio Chierichetti*
Sapienza, University of Rome

George Giakkoupis
INRIA, Rennes[†]

Silvio Lattanzi
Google Research, New York

Alessandro Panconesi[‡]
Sapienza, University of Rome

October 6, 2017

Abstract

In this paper we study the completion time of the PUSH-PULL variant of rumor spreading, also known as randomized broadcast. We show that if a network has n nodes and conductance ϕ then, with high probability, PUSH-PULL will deliver the message to all nodes in the graph within $\mathcal{O}(\log n/\phi)$ many communication rounds. This bound is best possible. We also give an alternative proof that the completion time of PUSH-PULL is bounded by a polynomial in $\log n/\phi$, based on graph sparsification. Although the resulting asymptotic bound is not optimal, this proof shows an interesting and, at the outset, unexpected connection between rumor spreading and graph sparsification. Finally, we show that if the degrees of the two endpoints of each edge in the network differ by at most a constant factor, then both PUSH and PULL alone attain the optimal completion time of $\mathcal{O}(\log n/\phi)$, with high probability.

1 Introduction

Rumor spreading, also known as randomized broadcast or randomized gossip, is a well-known distributed algorithm. Starting with one source node with a message, the algorithm proceeds in a sequence of synchronous rounds with the goal of broadcasting the message, i.e., delivering it to all nodes in the network. In each round, every node that knows the message selects a neighbor uniformly at random to which the message is forwarded. This is the so-called PUSH strategy. The PULL variant is symmetric. In each round, every node that does not yet have the message selects a neighbor uniformly at random and asks for the information, which is then transferred to the node provided that the queried neighbor possesses it. Finally, the PUSH-PULL strategy is a combination of both: in each round, every node selects a random neighbor to perform a PUSH if it has the message, or a PULL if it does not have it.

Simplicity and effectiveness are among the many virtues of rumor spreading, making it an ideal candidate for quick dissemination of information inside a network. The algorithm has been extensively studied, since these three strategies were introduced by Demers et al. [12]. Naturally, one of the most studied questions concerns the completion time, i.e., how many rounds it takes to disseminate the information to all nodes in the graph, assuming a worst-case source.

*Supported in part by the ERC Starting Grant DMAP 680153, by a Google Focused Research Award and by the SIR Grant RBSI14Q743.

[†]Supported in part by the ANR NDFusion grant.

[‡]This author was kindly supported by faculty awards of Yahoo! Research, IBM and Google.

Typically this question has been approached on a case by case basis, by giving careful estimates for particular classes of graphs. For instance, when the network is the complete graph we know that PUSH informs every vertex within $\log_2 n + \ln n + \mathcal{O}(1)$ rounds almost surely [20, 33]. Other classes of graphs for which rumor spreading has been analyzed include hypercubes, constant degree graphs, and a large variety of random graphs (see Section 2 for more details).

Our approach here is quite different. Rather than focusing on special classes of graphs, we want to identify a set of necessary and/or sufficient conditions for rumor spreading to be fast in any given network satisfying such conditions, where by “fast” we mean that the completion time is poly-logarithmic in the size of the network with high probability. In this work, we make a step in this direction by providing a general sufficient condition for rumor spreading to be fast: *high conductance*.

The first result of this paper is to show that if a network has n nodes and conductance ϕ then, with high probability, the completion time of PUSH-PULL is $\mathcal{O}(\log n/\phi)$. This bound is optimal, for there are graphs of conductance ϕ and diameter $\Theta(\log n/\phi)$. The proof of this bound is based on a decoupling of the operations of PUSH and PULL. We show that, once the rumor reaches a vertex of maximum degree, every remaining vertex is informed within $\mathcal{O}(\log n/\phi)$ many rounds with high probability by means of PULL alone. There is a clear symmetry between PUSH and PULL that can be exploited to show that the probability of spreading the message from a vertex u to another vertex v within t rounds via PULL is the same as that of spreading the message in the opposite direction, from v to u , within the same t rounds, via PUSH. Therefore, the argument above also shows that, starting from any source, a maximum degree vertex will be informed within $\mathcal{O}(\log n/\phi)$ many rounds with high probability by using PUSH alone. Taken together, these two bounds for PUSH and PULL give the optimal $\mathcal{O}(\log n/\phi)$ bound for PUSH-PULL.

The second result of this paper is an alternative proof showing that the completion time of PUSH-PULL is bounded by a polynomial function in $\log n/\phi$. Unlike the previous approach, which is based on a decoupling of PUSH and PULL, this approach analyzes the PUSH-PULL process as a whole. It is based on what we consider to be an interesting and, at the outset, surprising connection between rumor spreading and the spectral sparsification theory. This connection is used to show that a small number of PUSH-PULL rounds (polynomial in $\log n/\phi$) expands the volume of the set of informed nodes by a factor of $(1 + \Theta(\phi^2))$, with high probability. In turn, this implies that with high probability, starting from any source, after a number of rounds of PUSH-PULL that is poly-logarithmic in n and polynomial in ϕ^{-1} , every node is informed. The resulting bound on the completion time is sub-optimal, but we believe that the proof remains interesting. In a broad sense, it draws a connection between the theory of spectral sparsification and the speed with which diffusion processes make progress in a network. This could potentially lead to further insights beyond the present work and is worth exploring.

The above bounds do not hold for neither PUSH nor PULL in isolation. Perhaps the simplest counterexample is a star of n nodes. The star has maximum conductance, but PUSH takes at least $n - 1$ rounds to complete, regardless of the source, while a standard Coupon Collector’s argument shows that $\Theta(n \log n)$ many rounds are needed with high probability. Moreover, a basic expectation argument shows that when a leaf is the source, PULL requires n many rounds, on average. See also [8] for examples of preferential attachment graphs with high conductance in which the two strategies are proven to be slow.

The final result of the paper is to show a simple sufficient condition for the optimal $\mathcal{O}(\log n/\phi)$ bound to be attained with high probability by PUSH and by PULL alone: for each edge in the graph, the degrees of its two endpoints differ by at most a constant factor.

One of the motivations for our work comes from social networks. Loosely stated, we are looking for a theorem of the form “*Rumor spreading is fast in social networks.*” There is however no

universally accepted definition of social network. One way to circumvent this difficulty is to look for properties that are enjoyed by social networks and see whether they imply that rumor spreading is fast. The empirical evidence available suggests that social networks have high conductance [29]. Furthermore, this is true for preferential attachment models of social networks such as the one proposed by Mihail et al. [30]. Thus, our results can be considered as a possible explanation of why epidemic processes resembling rumor spreading spread virally in social networks.

We should note that the PUSH-PULL process has some sociological justification. PUSH corresponds to communicating in an active fashion, e.g., texting a message with a mobile phone or sending an email to a specific recipient. PULL on the other hand corresponds to, e.g., reading a tweet or a post on the wall of a Facebook account. Thus, PUSH-PULL seems to be a reasonable first approximation to model the spread of viral information in the age of the Internet.

The work presented here has been the first to show a connection between rumor spreading and graph expansion properties that, intuitively, should ensure fast dissemination.¹ This line of research has generated considerable follow-up work on the relationship between expansion properties and rumor spreading which we now know runs very deep. Following our work, optimal bounds for rumor spreading in terms of vertex expansion have been obtained [22, 23, 35], and extensions of the bounds with conductance and vertex expansions have been proposed to dynamic graph settings [24]. Other works have looked for general bounds for the completion time of rumor spreading without any dependence on conductance, showing that $\mathcal{O}(D + \text{polylog}(n))$ rounds are sufficient, where D is the diameter of the graph (and hence a trivial lower bound on the completion time) [5, 26]. Still, these results rely on conductance, in the sense that they show that every graph can be partitioned into regions each of which is of high conductance, and within which rumors spread fast for the reasons elucidated in this work.

The rest of the paper is organized as follows. Section 2 discusses the relevant literature. Section 3 introduces some common concepts and definitions to be used in the rest of the paper. Section 4 presents the optimal $\mathcal{O}(\log n/\phi)$ bound for PUSH-PULL. Section 5 presents a sufficient condition for this bound to hold when PUSH and PULL operate by themselves. Finally, Section 6 shows the connection between rumor spreading and graph sparsification.

2 Related Work

The literature on rumor spreading is huge and we confine ourselves to what appears to be more relevant to the present work.

Clearly, at least diameter-many rounds of rumor spreading are needed for the message to reach all nodes. On the other hand, Feige et al. [18] have shown that $\mathcal{O}(n \log n)$ rounds are sufficient with high probability for any connected graph with n nodes. The problem has been studied on a number of graph classes, such as cliques, hypercubes, bounded-degree graphs, and various types of random graphs (see, for instance, [18, 20, 33]). More recently, there has been a lot of work on “quasi-regular” expanders, i.e., expander graphs for which the ratio between the maximum and minimum degree is constant. It has been shown in different settings [3, 14, 15, 34] that $\mathcal{O}(\log n)$ rounds are sufficient for the message to spread throughout the graph. Our work can be viewed as an extension of these studies to graphs of arbitrary degree distribution. Observe that many real world graphs (e.g., Facebook, the Internet, etc.) have a very skewed degree distribution, i.e., the

¹This paper is the result of a series of intermediate steps [9, 10, 21]. Paper [10] was the first to show the connection between rumor spreading and conductance. By using a different approach, [9] followed up with an almost tight solution to the problem. Finally, [21] improved the latter approach introducing new ideas, obtaining a tight solution.

ratio between the maximum and minimum degrees is very high. In most graph models for social networks this ratio is polynomial in the order n of the graph.

Mihail et al. [30] have studied the edge expansion and conductance of graphs that are very similar to preferential attachment graphs, and showed that their edge expansion and conductance are constant. For “genuine” preferential-attachment graphs, Chierichetti et al. [8] showed that rumor spreading finishes in $\mathcal{O}(\log^2 n)$ rounds. A subsequent paper by Doerr et al. [13] improved that, by showing a tight bound of $\mathcal{O}(\log n)$ rounds and moreover, if one is interested in informing only $n - o(n)$ nodes, then one has to wait for just $\mathcal{O}(\log n / \log \log n)$ many rounds.

Fountoulakis et al. [19] studied another class of graphs with power-law degree distribution, for which they showed that $\mathcal{O}(\log \log n)$ rounds suffice with high probability to inform $n - o(n)$ nodes.

Boyd et al. [4], and Mosk-Aoyama and Shah [32] have shown that high conductance implies that *non-uniform* (over neighbors) rumor spreading is fast. By non-uniform we mean that, for every ordered pair of neighbors i and j , node i will select j with probability p_{ij} for the rumor spreading step (in general, $p_{ij} \neq p_{ji}$). This result does not extend to the case of uniform probabilities studied in this paper. In our setting (but not in theirs), the existence of a non-uniform distribution that makes rumor spreading fast is a rather trivial matter. A graph of conductance ϕ has diameter bounded by $\mathcal{O}(\log n / \phi)$. Observe that in a synchronous network, it is possible to elect a leader in $\mathcal{O}(\log n / \phi)$ many rounds and set up a BFS tree originating from it. Assigning probability 1 to the edge between a node and its parent gives the desired non-uniform probability distribution.

Giakkoupis et al. [25] have shown an $\mathcal{O}(\log n / \phi)$ bound for a “low-randomness” variant of PUSH on regular graphs. In their algorithm, the random choices of nodes are not independent, but rely instead on a pseudorandom generator whose random seed is attached to the message.

Another standard measure of graph expansion that has been studied in connection with rumor spreading is vertex expansion [22, 23, 35]. The main result of this line of work is a tight bound of $\mathcal{O}(\log n \cdot \log \Delta / \alpha)$ on the number of rounds for PUSH-PULL to spread a message with high probability, for any graph with vertex expansion α and maximum degree Δ . This bound is incomparable to the $\mathcal{O}(\log n / \phi)$ bound, as the vertex expansion is incomparable to conductance for general graphs.

An interesting application of the $\mathcal{O}(\log n / \phi)$ bound has been in the design of a new class of information dissemination protocols [5, 6, 26]. 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 a round. These new protocols achieve fast information spreading for broad families of networks, and some of these protocols achieve for all graphs time bounds that are close (within polylogarithmic factors or polylogarithmic terms) to the network diameter.

For the message complexity of PUSH-PULL, we observe first that it has been determined precisely only for very special classes of graphs (namely, cliques [28] and Erdős-Rényi random graphs [17]). Apart from that, given the generality of the graph class we consider, it is not possible to obtain a non-trivial upper bound on the number of messages — where the trivial upper bound is the number of rounds times n . Consider for instance the “lollipop graph.” Fix $\omega(n^{-1}) < \phi < \mathcal{O}(\log^{-1} n)$, and suppose that we have a path of length ϕ^{-1} connected to a clique of size $n - \phi^{-1} = \Theta(n)$. This graph has conductance $\mathcal{O}(\phi)$. Let the source be an arbitrary node in the clique. After $\Theta(\log n)$ rounds each node in the clique will have the information with high probability. Further it will take at least ϕ^{-1} steps for the information to be sent to each node in the path. So, at least $n - \phi^{-1} = \Theta(n)$ messages are pushed (by the nodes in the clique) in each round, for at least $\phi^{-1} - \Theta(\log n) = \Theta(\phi^{-1})$ rounds. Thus, the total number of messages sent will be $\Omega(n\phi^{-1})$. Observing that the number of rounds is $\Theta(\phi^{-1} + \log n) = \Theta(\phi^{-1})$, we obtain that the number of rounds times n is (asymptotically) less than or equal to the number of transmitted messages.

We also note that one cannot give fault-tolerance guarantees (that is, the ability of the protocol

to resist to node and/or edge deletions) based only on conductance. A star has high conductance, but failure of the central node destroys connectivity.

3 Preliminaries

In this section we introduce some, mostly standard, notation and provide some basic technical lemmas to be used in the analysis.

Let $G = (V, E)$ be an undirected graph with $n = |V|$ vertices. For a vertex $v \in V$, $\Gamma_G(v)$ denotes the set of v 's neighbors in G , and $\deg_G(v) := |\Gamma_G(v)|$ is the degree of v . By Δ_G we denote the maximum degree of G . For a set of vertices $S \subseteq V$, we let

$$\Gamma_G(S) := \bigcup_{v \in S} \Gamma_G(v) - S.$$

We define the *volume* of set $S \subseteq V$ to be the sum of the degrees of the vertices contained in S ,

$$\text{vol}_G(S) := \sum_{v \in S} \deg_G(v).$$

Observe that $\text{vol}_G(V) = 2 \cdot |E|$. Given two disjoint sets of vertices $A, B \subseteq V$, we let $\text{cut}_G(A, B)$ be the set of edges with one endpoint in A and the other in B .

Definition 1 (Conductance [27]). *The conductance of a non-empty vertex set S in graph G is*

$$\phi_G(S) := \frac{|\text{cut}_G(S, V - S)|}{\text{vol}_G(S)}.$$

The conductance of G is the minimum $\phi_G(S)$, over all sets S for which $\text{vol}(S) \leq \text{vol}_G(V)/2$,

$$\phi_G := \min_{S \subseteq V, 0 < \text{vol}(S) \leq \text{vol}(V)/2} \frac{|\text{cut}_G(S, V - S)|}{\text{vol}_G(S)}.$$

In all the notations above we will omit G whenever the graph is clear from the context.

We say that an event occurs *with high probability*, abbreviated as *w.h.p.*, if it happens with probability $1 - n^{-\Omega(1)}$.

Let $B(k, p)$ denote a binomial random variable with parameters k and p , i.e., $B(k, p)$ is the number of successes among k independent trials each of which succeeds with probability p . We will use the next standard lemma, which can be found in [1, Lemma 3.1].

Lemma 2. *Let Z_1, \dots, Z_k be random variables with values in some arbitrary domain, and let X_1, \dots, X_k be binary random variables such that $X_i = X_i(Z_1, \dots, Z_i)$, for $1 \leq i \leq k$. Let $X := \sum_{i=1}^k X_i$. If $\Pr(X_i = 1 \mid Z_1, \dots, Z_{i-1}) \geq p$, for all $1 \leq i \leq k$, then for any t ,*

$$\Pr(X < t) \leq \Pr(B(k, p) < t).$$

The following lemma gives simple bounds on the number of coupons collected in the Coupon Collector's process, within a given number of trials. We denote that number by $C(k, t)$, where k is the total number of coupons, and t is the number of trials.

Lemma 3. Consider the standard Coupon Collector's process, where we are given an urn of k different coupons, and we draw coupons from the urn, equally likely with replacement. Let $C(k, t)$ denote the number of distinct coupons drawn in t trials. Then for any $t \geq 0$,

$$\Pr(C(k, t) < k) \leq ke^{-t/k}, \quad (1)$$

and for any $\lambda < k$,

$$\Pr(C(k, t) \leq \lambda) \leq \Pr(B(t, 1 - \lambda/k) \leq \lambda). \quad (2)$$

Proof. The probability that a given coupon is not selected in t trials is

$$(1 - 1/k)^t \leq e^{-t/k},$$

and by a union bound, the probability at least one coupon is not selected, i.e., $C(k, t) < k$, is at most $ke^{-t/k}$. This proves (1).

Let X_i , $1 \leq i \leq t$, be a binary random variable which is 1 if a new coupon is selected in the i th trial or if the total number of distinct coupons selected in the previous $i - 1$ trials is larger than λ ; otherwise, $X_i = 0$. Then

$$\Pr(C(k, t) \leq \lambda) = \Pr\left(\sum_{i=1}^t X_i \leq \lambda\right).$$

The probability that $X_i = 1$ is at least $1 - \frac{\lambda}{k}$, independently of the outcome of the previous trials, because either more than λ distinct coupons have been selected in the first $i - 1$ trials and then $X_i = 1$, or at most λ coupons have been selected, and thus the probability a new coupon is selected in the i th trial is at least $\frac{k-\lambda}{k}$. Therefore, we can apply Lemma 2 to obtain

$$\Pr\left(\sum_{i=1}^t X_i \leq \lambda\right) \leq \Pr(B(t, 1 - \lambda/k) \leq \lambda).$$

Combining the last two equations above gives (2). \square

Finally, we will use the following Chernoff bounds for binomial random variables, which can be found, e.g., in [16, 31].

Lemma 4 (Chernoff bounds).

(a) For $t \leq k/2$,

$$\Pr(B(k, 1/2) \leq t) \leq e^{-2(k/2-t)^2/k}.$$

(b) For $0 < \epsilon \leq 1$,

$$\Pr(B(k, p) \geq (1 + \epsilon)np) \leq e^{-\epsilon^2 np/3}.$$

4 Optimal Bound for Push-Pull

In this section, we prove a general upper bound of $\mathcal{O}(\log n/\phi)$ rounds for PUSH-PULL on any n -node graph with conductance ϕ . Further, we argue that this bound is best possible, if we have no additional information for the graph other than ϕ and n .

The formal statement of the upper bound is as follows.

Theorem 5. *Let $G = (V, E)$ be a connected n -node graph with conductance ϕ . Then $\mathcal{O}(\log n/\phi)$ rounds of PUSH-PULL suffice w.h.p. to spread to all nodes of G a message originated at an arbitrary source node.*

The proof of Theorem 5 relies on the next key lemma, which states that a message spreads from a node of maximum degree to all other nodes in $\mathcal{O}(\log n/\phi)$ rounds w.h.p. This statement is true even if only PULL is used. Recall that Δ denotes the maximum degree of G .

Lemma 6. *Using PULL, a message originated at a node of degree Δ spreads to all other nodes in $\mathcal{O}(\log n/\phi)$ rounds with probability at least $1 - n^{-c}$, for any constant $c > 0$.*

The proof of Lemma 6 is based on an analysis of the growth of the volume of the set of informed nodes, and is given in Section 4.1.

We will also use the following basic lemma, which underlines the symmetries in the PUSH and PULL strategies.

Lemma 7. *Let $u \xrightarrow{t}_{\text{push}} v$ denote the event that a message originated at vertex u reaches vertex v in at most t rounds when PUSH is used, and define $u \xrightarrow{t}_{\text{pull}} v$ similarly. For any $u, v \in V$, and $t \geq 0$, we have $\Pr(u \xrightarrow{t}_{\text{push}} v) = \Pr(v \xrightarrow{t}_{\text{pull}} u)$.*

Proof. W.l.o.g. we can assume that in each round, every node (informed or not) chooses a random neighbor. Of course only informed nodes actually establish a connection to the neighbor they choose in the PUSH strategy, and only uninformed nodes establish a connection in the PULL strategy. Let Ω_t be the space of possible choices for all nodes over t rounds, i.e., each $\omega \in \Omega_t$ specifies for every node $u \in V$ and round $i \in \{1 \dots t\}$, the neighbor that u chooses in round i . Observe that all $\omega \in \Omega_t$ have the same probability: $\Pr(\omega) = \prod_{u \in V} (\deg(u))^{-t}$. Further, each $\omega \in \Omega_t$ completely determines the spread of a message in t rounds, given the source node and the rumor spreading strategy. Let $\Omega_t(u, v, \text{push})$ be the set of all $\omega \in \Omega_t$ for which a message originated at u reaches v within t rounds when PUSH is used, and define $\Omega_t(u, v, \text{pull})$ similarly. Since all $\omega \in \Omega_t$ have the same probability, it suffices to show that $|\Omega_t(u, v, \text{push})| = |\Omega_t(v, u, \text{pull})|$ in order to prove $\Pr(u \xrightarrow{t}_{\text{push}} v) = \Pr(v \xrightarrow{t}_{\text{pull}} u)$.

For $\omega \in \Omega_t$, let $rev(\omega) \in \Omega_t$ contain the same collection of choices as ω but in reverse round order, i.e., if node u chooses node v in round i in ω , then u chooses v in round $t - i$ in $rev(\omega)$. We now argue that $\omega \in \Omega_t(u, v, \text{push})$ iff $rev(\omega) \in \Omega_t(v, u, \text{pull})$. A *push path* for $\omega \in \Omega_t$ is a list of nodes u_0, \dots, u_t such that for any two consecutive nodes u_{i-1} and u_i , either (i) $u_{i-1} = u_i$, or (ii) u_{i-1} chooses u_i in round i . A *pull path* is defined similarly except that (ii) is replaced by: (ii') u_i chooses u_{i-1} in round i . For a given $\omega \in \Omega_t$, PUSH (or PULL) informs node v within t rounds of a message originated at u iff there is a push path (resp. pull path) from u to v . Further, u_0, u_1, \dots, u_t , with $u_0 = u$ and $u_t = v$, is a push path for ω iff u_t, u_{t-1}, \dots, u_0 is a pull path for $rev(\omega)$. From the last two observations it follows that $\omega \in \Omega_t(u, v, \text{push})$ iff $rev(\omega) \in \Omega_t(v, u, \text{pull})$. From this and the fact that for distinct $\omega, \omega' \in \Omega_t$ we have $rev(\omega) \neq rev(\omega')$, we obtain $|\Omega_t(u, v, \text{push})| = |\Omega_t(v, u, \text{pull})|$. \square

Theorem 5 follows easily from the two lemmas above: From Lemma 6 we have that a message originating from an arbitrary source node u needs $\mathcal{O}(\log n/\phi)$ rounds of PULL to spread to all nodes w.h.p., after it has reached a node v of degree Δ . Further, by combining Lemma 6 with Lemma 7, we obtain that $\mathcal{O}(\log n/\phi)$ rounds suffice w.h.p. for u 's message to reach the max-degree node v using PUSH.

Proof of Theorem 5 (assuming Lemma 6 holds). Let u be an arbitrary node, and v be a node with $\deg(v) = \Delta$. By Lemma 6, there is some $t = \mathcal{O}(\log n/\phi)$ such that t rounds of PULL

suffice with probability at least $1 - n^{-c}$ to spread a message from v to all other nodes, and thus to node u . From this and Lemma 7, it follows that t rounds of PUSH suffice to spread a message from u to v with probability at least $1 - n^{-c}$; thus the same is true when PUSH-PULL is used instead of just PUSH. Once v has learned u 's message, Lemma 6 implies, as mentioned above, that all nodes will learn the message with probability at least $1 - n^{-c}$ after at most t additional rounds of PULL (and thus of PUSH-PULL). A union bound then gives that with probability at least $1 - 2n^{-c}$, v learns u 's message within t rounds and then all remaining nodes learn the message within t additional rounds. Thus all nodes get informed after a total number of at most $2t = \mathcal{O}(\log n/\phi)$ rounds w.h.p. \square

Below, we give the proof of Lemma 6, in Section 4.1, and we argue about the optimality of Theorem 5, in Section 4.2.

4.1 Analysis of Pull (Proof of Lemma 6)

In this section, we show that $\mathcal{O}(\log n/\phi)$ rounds of PULL suffice w.h.p. to spread a message to all nodes, if the source node has degree Δ .

The proof analyses the growth rate of the volume of informed vertices. Let I_t be the set of informed nodes after the first t rounds, and let $U_t = V - I_t$ be the set of uninformed nodes at that time. Roughly speaking, we argue that while $\text{vol}(I_t)$ is at most some constant fraction of the total volume $\text{vol}(V)$, $\text{vol}(I_t)$ increases by a factor of at least $1 + \Omega(\phi)$ per round on average; and after that, $\text{vol}(U_t)$ decreases by a factor of at most $1 - \Omega(\phi)$. It follows that $\mathcal{O}(\log(\text{vol}(V))/\phi) = \mathcal{O}(\log n/\phi)$ rounds suffice for $\text{vol}(I_t)$ to increase to a constant fraction of $\text{vol}(V)$, and then $\mathcal{O}(\log n/\phi)$ additional rounds suffice for $\text{vol}(U_t)$ to drop to zero.

To convey the main intuition let us compute the expected increase of $\text{vol}(I_t)$ in round $t + 1$ (see Lemma 8(a) for the formal argument): Suppose that $\text{vol}(I_t) \leq \text{vol}(V)/2$, and for each uninformed node $v \in U_t$ let $\text{deg}'(v)$ be the number of its informed neighbors. The contribution of v to the expected increase of $\text{vol}(I_t)$ is then $(\text{deg}'(v)/\text{deg}(v)) \cdot \text{deg}(v) = \text{deg}'(v)$, as v chooses some of its informed neighbors in round $t + 1$ (and pulls the message from it) with probability $\text{deg}'(v)/\text{deg}(v)$. It follows

$$\mathbf{E}[\text{vol}(I_{t+1}) \mid I_t] - \text{vol}(I_t) = \sum_{v \in U_t} \text{deg}'(v) = |\text{cut}(I_t, U_t)| \geq \phi \cdot \text{vol}(I_t),$$

by the definition of conductance. Thus, $\mathbf{E}[\text{vol}(I_{t+1}) \mid I_t] \geq (1 + \phi) \text{vol}(I_t)$. Similarly, we obtain $\mathbf{E}[\text{vol}(U_{t+1}) \mid U_t] \leq (1 - \phi) \text{vol}(U_t)$, when $\text{vol}(I_t) \geq \text{vol}(V)/2$. These simple expectation bounds, together with Markov's Inequality, suffice to obtain an $\mathcal{O}(\log n/\phi)$ bound w.h.p. on the number of rounds until all nodes get informed, from the time when a constant fraction of the total volume has been informed (Lemma 11). To bound the time to inform that constant fraction, however, more work is required. For that we bound also the variance of the increase of $\text{vol}(I_t)$ (Lemma 8(b)), and use the second moment method to show that the volume doubles in $\mathcal{O}(1/\phi)$ rounds with constant probability (Lemma 9). From this last result and Chernoff Bounds, we obtain the desired high-probability bound on the time to inform a constant fraction of the total volume (Lemma 10).

Next we give the detailed analysis.

Consider the volume L of nodes that get informed in a round by pulling the message through a given set D of edges between informed and uninformed nodes (D does not necessarily contain all edges between informed and uninformed nodes, and is fixed before the round starts). The next lemma states that the expectation of L is equal to the number of edges contained in D , and L 's variance is bounded by the number of edges times the maximum degree of their endpoints in U_t .

Lemma 8. Fix an arbitrary subset D of $\text{cut}(I_t, U_t)$ before round $t + 1$, and let L be the volume of the set of nodes that get informed in round $t + 1$ by pulling the message through edges from D . Then

(a) $\mathbf{E}[L] = |D|$.

(b) $\mathbf{Var}[L] \leq |D| \cdot \max_{v \in N} \deg(v)$, where N is the set of nodes $v \in U_t$ incident to some edge in D .

Proof. For each node $v \in N$, let $\deg_D(v)$ be the number of edges from set D that are incident to v . Clearly, $\sum_{v \in N} \deg_D(v) = |D|$. For each $v \in N$, let X_v be a binary random variable that is 1 if v gets informed in round $t + 1$ by pulling the message through some edge from D , and is 0 otherwise. Then, $\Pr(X_v = 1) = \deg_D(v) / \deg(v)$. We have $L = \sum_{v \in N} (X_v \cdot \deg(v))$, and thus

$$\begin{aligned} \mathbf{E}[L] &= \sum_{v \in N} \mathbf{E}[X_v \cdot \deg(v)] = \sum_{v \in N} \Pr(X_v = 1) \cdot \deg(v) \\ &= \sum_{v \in N} (\deg_D(v) / \deg(v)) \cdot \deg(v) \\ &= \sum_{v \in N} \deg_D(v) \\ &= |D|. \end{aligned}$$

Next we bound the variance of L . Since $L = \sum_{v \in N} (X_v \cdot \deg(v))$ and the random variables X_v are independent,

$$\mathbf{Var}[L] = \sum_{v \in N} \mathbf{Var}[X_v \cdot \deg(v)] = \sum_{v \in N} \mathbf{Var}[X_v] \cdot (\deg(v))^2.$$

Using also that $\mathbf{Var}[X_v] = \Pr(X_v = 1) \cdot \Pr(X_v = 0) \leq \Pr(X_v = 1) = \deg_D(v) / \deg(v)$, we obtain

$$\mathbf{Var}[L] \leq \sum_{v \in N} (\deg_D(v) \cdot \deg(v)) \leq \left(\sum_{v \in N} \deg_D(v) \right) \cdot \max_{v \in N} \deg(v) = |D| \cdot \max_{v \in N} \deg(v). \quad \square$$

In the following, we first bound the number of rounds to inform a constant fraction of the total volume, namely a $1/3$ fraction, and then bound the number of rounds to inform the remaining nodes.

Informing a Constant Fraction of the Volume

We prove that $\mathcal{O}(\ln n / \phi)$ rounds of PULL suffice w.h.p. to inform a subset of nodes with volume at least $\text{vol}(V)/3$, when the message starts from a node of degree Δ .

Our key claim is that if $\Delta \leq \text{vol}(I_t) \leq \text{vol}(V)/3$, then $\mathcal{O}(1/\phi)$ rounds suffice with constant probability to double the volume of informed nodes. The proof proceeds as follows. Before each round $t + k > t$, we fix an edge set $D_k \subseteq \text{cut}(I_{t+k-1}, U_{t+k-1})$ of size $|D_k| = \phi \cdot \text{vol}(I_t)$; this set can be arbitrary but must be fixed before round $t + k$ starts. Then we compute the volume L_k of the set of nodes that get informed in round $t + k$ by pulling the message through edges from D_k . Clearly, $\sum_{1 \leq i \leq k} L_i$ is a lower bound on $\text{vol}(I_{t+k}) - \text{vol}(I_t)$, and thus to prove the claim it suffices to show that $\sum_{1 \leq i \leq k} L_i \geq \text{vol}(I_t)$ holds with constant probability. To show that, we use Lemma 8 to compute the expectation of $\sum_{1 \leq i \leq k} L_i$ and to bound its variance, and then apply Chebyshev's Inequality.

Lemma 9 below proves a slightly more general statement, which does not assume that volume $\text{vol}(I_t)$ is at least Δ , and gives a lower bound on the probability that the volume doubles within a given number k of steps, as a function of k and the ratio $\Delta / \text{vol}(I_t)$.

Lemma 9. *If $\text{vol}(I_t) \leq \text{vol}(V)/3$ then for any $k > 1/\phi$,*

$$\Pr(\text{vol}(I_{t+k}) \geq 2 \text{vol}(I_t) \mid I_t) > 1 - \frac{k\phi}{(k\phi - 1)^2} \cdot \frac{\Delta}{\text{vol}(I_t)}.$$

Proof. Throughout the proof we assume that I_t is fixed and $\text{vol}(I_t) \leq \text{vol}(V)/3$. Let

$$\tau = \min \{k: \text{vol}(I_{t+k}) \geq 2 \text{vol}(I_t)\}.$$

For each $1 \leq k \leq \tau$, we define random variable L_k as follows. Let D_k denote the set of the first $M := \lceil \phi \cdot \text{vol}(I_t) \rceil$ edges in $\text{cut}(I_{t+k-1}, U_{t+k-1})$, according to some fixed total order over E .² (We will see in (3) that $|\text{cut}(I_{t+k-1}, U_{t+k-1})| \geq M$.) Then L_k is the volume of the set of nodes that get informed in round $t+k$ by pulling the message through edges from D_k . Further, for each $k > \tau$, we let $L_k = M$.

We argue now that for any t' with $t \leq t' < t + \tau$,

$$|\text{cut}(I_{t'}, U_{t'})| \geq M. \quad (3)$$

If $\text{vol}(I_{t'}) \leq \text{vol}(V)/2$ then $\phi(I_{t'}) \geq \phi$. Also $\text{vol}(I_{t'}) \geq \text{vol}(I_t)$. Then

$$|\text{cut}(I_{t'}, U_{t'})| = \phi(I_{t'}) \cdot \text{vol}(I_{t'}) \geq \phi \cdot \text{vol}(I_t).$$

If $\text{vol}(I_{t'}) > \text{vol}(V)/2$ then $\text{vol}(U_{t'}) < \text{vol}(V)/2$, and thus $\phi(U_{t'}) \geq \phi$. Also, $\text{vol}(I_{t'}) < 2 \text{vol}(I_t)$ and thus $\text{vol}(U_{t'}) > \text{vol}(V) - 2 \text{vol}(I_t) \geq \text{vol}(I_t)$, as $\text{vol}(I_t) \leq \text{vol}(V)/3$. Then

$$|\text{cut}(I_{t'}, U_{t'})| = \phi(U_{t'}) \cdot \text{vol}(U_{t'}) \geq \phi \cdot \text{vol}(I_t).$$

For each $k \geq 0$, let $Y_k = \sum_{1 \leq i \leq k} L_i$. Then for $k \leq \tau$, Y_k is a lower bound on the volume of nodes informed in rounds $t+1$ up to $t+k$, i.e., $Y_k \leq \text{vol}(I_{t+k}) - \text{vol}(I_t)$. It follows that if $\text{vol}(I_{t+k}) < 2 \text{vol}(I_t)$, then $k < \tau$ and $Y_k \leq \text{vol}(I_{t+k}) - \text{vol}(I_t) < \text{vol}(I_t)$. Therefore,

$$\Pr(\text{vol}(I_{t+k}) < 2 \text{vol}(I_t)) \leq \Pr(Y_k < \text{vol}(I_t)). \quad (4)$$

Next we bound the probability on the right side, that $Y_k < \text{vol}(I_t)$. For that, we compute the expectation and variance of Y_k , and then apply Chebyshev's Inequality.

If $1 \leq k \leq \tau$, then from Lemma 8(a) it follows that

$$\mathbf{E}[L_k \mid I_{t+k-1}] = |D_k| = M.$$

This equation is also trivially true if $k > \tau$, as $L_k = M$ in this case. It follows that for all $k \geq 1$,

$$\mathbf{E}[L_k] = \mathbf{E}[\mathbf{E}[L_k \mid I_{t+k-1}]] = M,$$

and by linearity of expectation,

$$\mathbf{E}[Y_k] = kM. \quad (5)$$

For the variance of Y_k we have

$$\mathbf{Var}[Y_k] = \mathbf{Var}[Y_{k-1} + L_k] = \mathbf{Var}[Y_{k-1}] + \mathbf{Var}[L_k] + 2 \mathbf{Cov}[Y_{k-1}, L_k].$$

²In fact, it is not important how set $D_k \subseteq \text{cut}(I_{t+k-1}, U_{t+k-1})$ is chosen, as long as it is fixed before round $t+k$, and has size M .

The covariance term is zero, as

$$\begin{aligned} & \mathbf{E} [(Y_{k-1} - \mathbf{E}[Y_{k-1}]) \cdot (L_k - \mathbf{E}[L_k]) \mid I_t \dots I_{t+k-1}] \\ &= (Y_{k-1} - \mathbf{E}[Y_{k-1}]) \cdot (\mathbf{E}[L_k \mid I_{t+k-1}] - \mathbf{E}[L_k]) \\ &= (Y_{k-1} - \mathbf{E}[Y_{k-1}]) \cdot (M - M) = 0, \end{aligned}$$

and thus $\mathbf{Cov}[Y_{k-1}, L_k] = \mathbf{E} [\mathbf{E} [(Y_{k-1} - \mathbf{E}[Y_{k-1}]) \cdot (L_k - \mathbf{E}[L_k]) \mid I_t \dots I_{t+k-1}]] = 0$. Substituting this to the formula for $\mathbf{Var}[Y_k]$ above yields

$$\mathbf{Var}[Y_k] = \mathbf{Var}[Y_{k-1}] + \mathbf{Var}[L_k].$$

We now bound $\mathbf{Var}[L_k]$. If $1 \leq k \leq \tau$, then from Lemma 8(b) it follows

$$\mathbf{Var}[L_k \mid I_{t+k-1}] \leq |D_k| \cdot \max_{v \in N_k} \deg(v) \leq M\Delta, \quad (6)$$

where N_k is the set of nodes $v \in U_{t+k-1}$ incident to some edge in D_k . The above inequality is also trivially true if $k > \tau$, as $\mathbf{Var}[L_k \mid I_{t+k-1}] = 0$ in this case. It follows that $\mathbf{Var}[L_k] \leq M\Delta$, for all $k \geq 1$. Substituting this to the equation for $\mathbf{Var}[Y_k]$ above gives $\mathbf{Var}[Y_k] \leq \mathbf{Var}[Y_{k-1}] + M\Delta$, and applying this inequality repeatedly yields

$$\mathbf{Var}[Y_k] \leq kM\Delta. \quad (7)$$

We use now Chebyshev's Inequality to bound the probability that $Y_k < \text{vol}(I_t)$. We assume that $k > 1/\phi$. Then from (5), $\mathbf{E}[Y_k] = kM > M/\phi \geq \text{vol}(I_t)$. We have

$$\Pr(Y_k < \text{vol}(I_t)) = \Pr(\mathbf{E}[Y_k] - Y_k > \mathbf{E}[Y_k] - \text{vol}(I_t)) < \frac{\mathbf{Var}[Y_k]}{(\mathbf{E}[Y_k] - \text{vol}(I_t))^2}.$$

Substituting the value of $\mathbf{E}[Y_k]$ from (5) and the bound for $\mathbf{Var}[Y_k]$ from (7), we obtain

$$\begin{aligned} \Pr(Y_k < \text{vol}(I_t)) &< \frac{kM\Delta}{(kM - \text{vol}(I_t))^2} \\ &= \frac{\Delta}{\text{vol}(I_t)} \cdot \frac{kM/\text{vol}(I_t)}{(kM/\text{vol}(I_t) - 1)^2} \\ &\leq \frac{\Delta}{\text{vol}(I_t)} \cdot \frac{k\phi}{(k\phi - 1)^2}, \end{aligned}$$

where the last inequality holds because $kM/\text{vol}(I_t) \geq k\phi$, and the function $f(x) = x/(x-1)^2$ is decreasing for $x > 1$. Combining the above bound on $\Pr(Y_k < \text{vol}(I_t))$ with (4) completes the proof of Lemma 9. \square

Using Lemma 9 and Chernoff Bounds, we show next our main claim, that $\mathcal{O}(\ln n/\phi)$ rounds suffice w.h.p. to inform a subset of nodes with volume at least $\text{vol}(V)/3$.

Lemma 10. *If the source of the message has degree Δ , then $\Pr(\text{vol}(I_t) \geq \text{vol}(V)/3) \geq 1 - n^{-c}$, for $t = \lceil 2c \ln n + 4 \log n \rceil \cdot \lceil 4/\phi \rceil$.*

Proof. We divide rumor spreading into phases of $k := \lceil 4/\phi \rceil$ rounds each, and consider whether the volume of informed vertices doubles in these phases. For each $i \geq 1$, let X_i be a binary random

variable that is 1 if $\text{vol}(I_{ik}) \geq 2 \text{vol}(I_{(i-1)k})$ or $\text{vol}(I_{(i-1)k}) \geq \text{vol}(V)/3$, and is 0 otherwise. From Lemma 9 it follows

$$\Pr(X_i = 1 \mid I_{(i-1)k}) > 1 - \frac{k\phi}{(k\phi - 1)^2} \cdot \frac{\Delta}{\text{vol}(I_{(i-1)k})} \geq 1 - \frac{4}{(4-1)^2} \cdot \frac{\Delta}{\Delta} = 5/9 > 1/2. \quad (8)$$

We bound the probability that $\text{vol}(I_t) < \text{vol}(V)/3$ in terms of the sum of X_i : If $\text{vol}(I_{ik}) < \text{vol}(V)/3$, then fewer than $\log(\text{vol}(V)/(3\Delta))$ among the random variables $X_1 \dots X_i$ are 1, because, otherwise, the volume of informed nodes would double at least $\log(\text{vol}(V)/(3\Delta))$ times in the first ik rounds, and thus increase from Δ to at least $2^{\log(\text{vol}(V)/(3\Delta))} \Delta = \text{vol}(V)/3$. Therefore,

$$\Pr(\text{vol}(I_{ik}) < \text{vol}(V)/3) \leq \Pr\left(\sum_{j \leq i} X_j < \log(\text{vol}(V)/(3\Delta))\right) \leq \Pr\left(\sum_{j \leq i} X_j < \log n\right),$$

as $\text{vol}(V) \leq n\Delta$. To bound the sum of the (non-independent) random variables X_i , we use the fact we showed in (8), that for all $i \geq 1$, $\Pr(X_i = 1 \mid I_{(i-1)k}) \geq 1/2$, and apply Lemma 2 to obtain $\Pr(\sum_{j \leq i} X_j < \log n) \leq \Pr(B(i, 1/2) < \log n)$. Then Lemma 4(a) gives for $i \geq 2 \log n$,

$$\Pr\left(\sum_{j \leq i} X_j < \log n\right) \leq e^{-2(i/2 - \log n)^2/i} \leq e^{-(i/2 - 2 \log n)}.$$

Letting $i \geq 2c \ln n + 4 \log n$, the above bound becomes at most $e^{-c \ln n} = n^{-c}$. Therefore, we have that $\Pr(\text{vol}(I_{ik}) < \text{vol}(V)/3) \leq n^{-c}$, for $i = \lceil 2c \ln n + 4 \log n \rceil$. \square

Informing the Remaining Nodes

We show that $\mathcal{O}(\ln n/\phi)$ rounds of PULL suffice w.h.p. to inform all the remaining nodes, after a set of nodes with volume at least $\text{vol}(V)/3$ has been informed.

Lemma 11. *If $\text{vol}(I_t) \geq \text{vol}(V)/3$, then $\Pr(U_{t+k} = \emptyset \mid I_t) \geq 1 - n^{-c}$, for $k = \lceil 2(c+2) \ln n/\phi \rceil$.*

Proof. In the following we assume I_t is fixed and $\text{vol}(I_t) \geq \text{vol}(V)/3$. We first argue that for $t' \geq t$,

$$|\text{cut}(I_{t'}, U_{t'})| \geq \phi \cdot \text{vol}(U_{t'})/2.$$

If $\text{vol}(U_{t'}) \leq \text{vol}(V)/2$, then $\phi(U_{t'}) \geq \phi$ and $|\text{cut}(I_{t'}, U_{t'})| = \phi(U_{t'}) \cdot \text{vol}(U_{t'}) \geq \phi \cdot \text{vol}(U_{t'})$. If $\text{vol}(U_{t'}) > \text{vol}(V)/2$, then $\phi(I_{t'}) \geq \phi$ and $|\text{cut}(I_{t'}, U_{t'})| = \phi(I_{t'}) \cdot \text{vol}(I_{t'}) \geq \phi \cdot \text{vol}(I_{t'}) \geq \phi \cdot \text{vol}(U_{t'})/2$, as $\text{vol}(I_{t'}) \geq \text{vol}(V)/3$.

Lemma 8(a), for $D = \text{cut}(I_{t'}, U_{t'})$, gives $\mathbf{E}[\text{vol}(U_{t'} - U_{t'+1}) \mid I_{t'}] = |\text{cut}(I_{t'}, U_{t'})|$. Substituting the bound for $|\text{cut}(I_{t'}, U_{t'})|$ from above and rearranging yields

$$\mathbf{E}[\text{vol}(U_{t'+1}) \mid I_{t'}] \leq (1 - \phi/2) \cdot \text{vol}(U_{t'}).$$

Taking the expectation on both sides gives $\mathbf{E}[\text{vol}(U_{t'+1})] \leq (1 - \phi/2) \cdot \mathbf{E}[\text{vol}(U_{t'})]$, and applying this inequality repeatedly yields, for $t' = t + k$,

$$\mathbf{E}[\text{vol}(U_{t'})] \leq (1 - \phi/2)^k \cdot \text{vol}(U_t) \leq e^{-k\phi/2} \cdot n^2.$$

For $k \geq 2(c+2) \ln n/\phi$, we get $\mathbf{E}[\text{vol}(U_{t'})] \leq n^{-c}$, and then Markov's Inequality gives

$$\Pr(U_{t'} \neq \emptyset) = \Pr(\text{vol}(U_{t'}) \geq 1) \leq \mathbf{E}[\text{vol}(U_{t'})]/1 \leq n^{-c}. \quad \square$$

Lemma 6 follows now immediately from Lemmas 10 and 11, and a union bound.

4.2 Tightness of the Upper Bound

We observe that the upper bound of Theorem 5 is tight. Indeed, there exist graphs on n nodes, with conductance ϕ , that have diameter $\Omega(\log n/\phi)$. We now present one example of such a graph. Let $\phi > 1/n^{1-\epsilon}$, and take any graph on $\Theta(\phi n)$ nodes with constant degree and constant conductance (e.g., a random cubic graph on $\Theta(\phi n)$ nodes.) Then, replace each edge of the graph with a path of length $\Theta(\phi^{-1})$. The resulting graph will have conductance $\Theta(\phi)$, and will contain $\Theta(n)$ nodes. Moreover, the diameter of the new graph will be equal to the diameter of the original graph times $\Theta(\phi^{-1})$ — since the original graph had constant degree its diameter is at least logarithmic in its order, that is at least $\Omega(\log(\phi n))$. Therefore, the new graph will have diameter at least $\Omega(\log(\phi n)/\phi)$. Since $\phi > 1/n^{1-\epsilon}$, the lower bound reduces to $\Omega(\log n/\phi)$.

5 A Sufficient Condition for Push and Pull

The $\mathcal{O}(\log n/\phi)$ bound we have shown for PUSH-PULL in Theorem 5 does not hold if only one of the PUSH or PULL strategies is used. As already mentioned, for the n -vertex star, which has conductance $\phi = 1$, at least $n - 1$ rounds of PUSH are needed to inform all nodes, regardless of the source. If PULL is used instead, and the message originates at a leaf of the star, then n rounds are needed in expectation ($n - 1$ expected rounds until the center of the star pulls the message from the source, and 1 additional round for all leaves to pull the message from the center). Observe, however, that if the message originates at the center, then all nodes get informed in just a single round of PULL. As shown in Lemma 6, for any graph $G = (V, E)$ there is some node $u \in V$ (namely, any max-degree node u) such that a message originated at u spreads to all nodes within $\mathcal{O}(\log n/\phi)$ rounds of PULL w.h.p. The same is not true for PUSH, as the example of the star shows.

The next theorem identifies a simple graph property which guarantees that the $\mathcal{O}(\log n/\phi)$ bound holds for PUSH and PULL, regardless of the source of the message. The property states that for any edge of the graph, the ratio of the degrees of its two endpoints is bounded by some constant.

Theorem 12. *Let $G = (V, E)$ be a connected n -node graph with conductance ϕ and the property that for any edge $uv \in E$, $\deg(u)/\deg(v) = \Theta(1)$. Then $\mathcal{O}(\log n/\phi)$ rounds of PUSH or PULL suffice w.h.p. to spread to all nodes of G a message originated at an arbitrary source node.*

Proof. The analysis for PULL is the same as that for Lemma 6, except for a few small modifications we point out now. In place of Lemma 9, we use the following statement.

Lemma 13. *Suppose that for every edge $uv \in E$, we have $\deg(u)/\deg(v) \in [1/\alpha, \alpha]$ for some $\alpha \geq 1$. If $\text{vol}(I_t) \leq \text{vol}(V)/3$ then for any $k > 1/\phi$,*

$$\Pr(\text{vol}(I_{t+k}) \geq 2 \text{vol}(I_t) \mid I_t) > 1 - \frac{k\phi}{(k\phi - 1)^2} \cdot \alpha.$$

The proof of Lemma 13 is identical to that of Lemma 9, except that, in (6), we do not bound $\max_{v \in N_k} \deg(v)$ by Δ , but use instead that

$$\max_{v \in N_k} \deg(v) \leq \alpha \cdot \text{vol}(I_t).$$

This inequality holds because each node $v \in N_k$ has degree bounded by $\alpha \cdot \max_{u \in I_{t+k-1}} \deg(u)$, and each node $u \in I_{t+k-1}$ has degree at most $\text{vol}(I_t)$: if $u \in I_t$ then $\deg(u) \leq \text{vol}(I_t)$, and if $u \in I_{t+k-1} - I_t$ then $\deg(u) \leq \text{vol}(I_{t+k-1}) - \text{vol}(I_t) < \text{vol}(I_t)$, since $\text{vol}(I_{t+k-1}) < 2 \text{vol}(I_t)$.

Using Lemma 13, we show, as in Lemma 10, that $\mathcal{O}(\log n/\phi)$ rounds suffice w.h.p. to inform a 1/3 fraction of the total volume, but now we do not need the assumption that the source node has degree Δ . Precisely, we show that $\Pr(\text{vol}(I_t) \geq \text{vol}(V)/3) \geq 1 - n^{-c}$ for $t = \lceil 2c \ln n + 4 \log n \rceil \cdot \lceil 4\alpha/\phi \rceil$. From this and Lemma 11, the bound of Theorem 12 follows for PULL.

For PUSH, the $\mathcal{O}(\log n/\phi)$ bound is obtained by combining the above $\mathcal{O}(\log n/\phi)$ bound for PULL, with Lemma 7 on the symmetry between the two strategies. \square

6 Rumor Spreading and Graph Sparsification

In this section, we give an alternative proof that PUSH-PULL finishes in a number of rounds that is at most polynomial in $\log n/\phi$. The proof relies on an interesting connection between rumor spreading and the Spielman-Teng sparsification process [36, 37].

Theorem 14. *Let $G = (V, E)$ be a connected n -node graph with conductance ϕ . Then $\mathcal{O}(\log^4 n/\phi^6)$ rounds of PUSH-PULL suffice w.h.p. to spread to all nodes in G a message originated at an arbitrary source node.*

The proof of Theorem 14 relates the *volume expansion* of the Spielman-Teng sparsification process, denoted ST, to that of an edge sampling process, PP, which selects edges similarly to PUSH-PULL. Roughly, the volume expansion of an edge sampling process on G is the minimum, over all $S \subseteq V$, of the ratio $\text{vol}_G(S')/\text{vol}_G(S)$, where S' is the set of vertices in $V - S$ that are reachable from S via a sampled edge (Definition 15). Process ST selects each edge uv independently with probability $\Upsilon/\min\{\deg(u), \deg(v)\}$, where $\Upsilon = \Theta(\log^2 n/\phi^4)$. Process PP on the other hand, chooses for each node u a number of $\Psi = \Theta(\Upsilon \cdot \log n)$ edges incident to u at random with replacement, imitating the random choices of nodes in Ψ rounds of PUSH-PULL.

It has been shown that the subgraph sampled by ST almost preserves certain spectral parameters of the original graph G . We use this result to show that the volume expansion of ST is $\Omega(\phi^2)$ (Lemma 21). We relate ST and PP by showing that process PP stochastically dominates ST, in the sense that there is a coupling of the two processes for which the edges selected by ST are w.h.p. a subset of those selected by PP (Lemma 23); this is our key technical claim. Once we have established this dominance, we infer that the volume expansion of PP is at least that of ST, i.e., $\Omega(\phi^2)$. It follows that Ψ rounds of PUSH-PULL increase the volume of informed nodes by a factor of $1 + \Omega(\phi^2)$, and thus $\mathcal{O}(\Psi \cdot \log n/\phi^2) = \mathcal{O}(\log^4 n/\phi^6)$ rounds suffice to inform half of the total volume. A symmetric argument gives that $\mathcal{O}(\log^4 n/\phi^6)$ additional rounds suffice to inform the remaining graph.

Below, we first set some definitions and notation, in Section 6.1. Then we review Spielman-Teng sparsification and analyze its volume expansion in Section 6.2. Last, we relate that process to rumor spreading and prove Theorem 14 in Section 6.3.

6.1 Definitions

We start with some terminology on edge sampling processes.

Definition 15 (Edge Sampling & Volume Expansion). *Let f be a (randomized) process selecting edges in graph $G = (V, E)$. We denote by $f(G)$ the spanning subgraph of G whose edge set consists of the edges selected by f . Process f may also associate a weight with each edge it selects, in which case $f(G)$ is a weighted graph. The volume expansion of f on G , denoted $\beta_{f,G}$, is the minimum*

over all non-empty sets $S \subset V$, of the ratio³

$$\frac{\text{vol}_G(\Gamma_{f(G)}(S))}{\min\{\text{vol}_G(S), \text{vol}_G(V - S)\}}.$$

Definition 15 above states that graph $f(G)$ obtained by process f may have weights associated to its edges. Next we formalize various notions of weight we will use, and also extend the standard definition of conductance to take weights into account.

Definition 16 (Weights). *Let $G = (V, E, w)$ be a weighted graph where each edge $e \in E$ has weight w_e . The weight of a set of edges $E' \subseteq E$ is $w_G(E') := \sum_{e \in E'} w_e$; the weight of a vertex $u \in V$ is $w_G(u) := \sum_{v \in \Gamma(u)} w_{uv}$; and the weight of a set of vertices $S \subseteq V$ is $w_G(S) := \sum_{u \in S} w_G(u)$.*

Definition 17 (Weighted Conductance). *Consider a weighted graph $G = (V, E, w)$. The weighted conductance of a non-empty vertex set S in graph G is*

$$\varphi_G(S) := \frac{w_G(\text{cut}_G(S, V - S))}{w_G(S)}.$$

The weighted conductance φ_G of G is the minimum $\varphi_G(S)$ over all S for which $w_G(S) \leq w_G(V)/2$.

In the notations above, we will omit subscript G whenever the graph is clear from the context.

6.2 Spectral Sparsification

Next we recall the Spielman-Teng spectral sparsification procedure.

Definition 18 (Spielman-Teng Sparsification). *Given a graph $G = (V, E)$, the Spielman-Teng sparsification process [36, 37], denoted ST , is an edge sampling process which selects each edge $uv \in E$ independently with probability*

$$p_{uv} := \min \left\{ 1, \frac{\Upsilon}{\min\{\deg(u), \deg(v)\}} \right\},$$

where $\Upsilon > 0$ is a parameter of the process. The process also assigns to each selected edge e a weight $w_e := p_e^{-1}$.

Note that since the weight of each edge in $\text{ST}(G)$ is inversely proportional to the edge selection probability, the expected weight of each vertex in $\text{ST}(G)$ is the same as the degree of that vertex in G .

It was shown in [37, Theorem 6.1] that if $\Upsilon = \Omega(\log^2 n / \phi_G^4)$, then with high probability, $\text{ST}(G)$ is an $(1 + \epsilon)$ -spectral approximation of graph G , i.e., it has the same Laplacian quadratic form as G within a factor of $1 + \epsilon$. This implies the following result, stating that $\text{ST}(G)$ approximately preserves cut weights, as well as the conductance of G — the latter within a quadratic exponent.

Lemma 19 (Properties of ST). *Let $0 < \epsilon < 1/2$ and $c \geq 1$, and let $G = (V, E)$ be an n -node graph with conductance ϕ . Suppose that we apply the ST process on G with parameter*

$$\Upsilon \geq \left(\frac{24c \log(3n)}{\epsilon \phi^2} \right)^2,$$

and let $\tilde{G} := \text{ST}(G)$. Then with probability at least $1 - n^{-c}$, the following hold.

³Recall that $\Gamma_G(S) = \bigcup_{u \in S} \Gamma_G(u) - S$.

(a) For any $S \subseteq V$,

$$\frac{1}{1+\epsilon} \cdot |\text{cut}_G(S, V-S)| \leq w_{\tilde{G}}(\text{cut}_{\tilde{G}}(S, V-S)) \leq (1+\epsilon) \cdot |\text{cut}_G(S, V-S)|. \quad (9)$$

(b) The weighted conductance $\varphi_{\tilde{G}}$ of \tilde{G} satisfies

$$\varphi_{\tilde{G}} \geq \frac{\phi_G^2}{4(1+\epsilon)}. \quad (10)$$

Remark 20. As already mentioned, [37, Theorem 6.1] is more general than our Lemma 19, proving that \tilde{G} is an $(1+\epsilon)$ -spectral approximation of G . This immediately implies the weaker notion of $(1+\epsilon)$ -cut approximation stated in (9) [2]. To obtain (10) we use the fact that the $(1+\epsilon)$ -spectral approximation implies that the eigenvalues of the Laplacians of G and \tilde{G} differ by at most a factor of $(1+\epsilon)$. Then we apply Cheeger's Inequality [7, 11], stating that for any unweighted or weighted graph with (weighted) conductance ϕ the smallest non-zero eigenvalue λ of its normalized Laplacian satisfies $\phi^2/2 \leq \lambda \leq 2\phi$.

Using Lemma 19, we show now that the volume expansion of ST is $\Omega(\phi_G^2)$ w.h.p.

Lemma 21 (Volume Expansion of ST). *Under the same assumptions as in Lemma 19, with probability at least $1 - n^{-c}$, the volume expansion of ST on G is*

$$\beta_{\text{ST},G} \geq \frac{\phi_G^2}{4(1+\epsilon)^3}. \quad (11)$$

Proof. Lemma 19(a), applied for a singleton set $S = \{u\}$, implies that for any $u \in V$,

$$\frac{1}{1+\epsilon} \cdot w_{\tilde{G}}(u) \leq \deg_G(u) \leq (1+\epsilon) \cdot w_{\tilde{G}}(u).$$

From this, it follows that for any $T \subseteq V$, the volume of T in G : $\text{vol}_G(T) = \sum_{u \in T} \deg_G(u)$, satisfies

$$\frac{1}{1+\epsilon} \cdot w_{\tilde{G}}(T) \leq \text{vol}_G(T) \leq (1+\epsilon) \cdot w_{\tilde{G}}(T). \quad (12)$$

We now proceed to derive (11). Applying (12) for $T = \Gamma_{\tilde{G}}(S)$, we obtain

$$\begin{aligned} \text{vol}_G(\Gamma_{\tilde{G}}(S)) &\geq \frac{1}{1+\epsilon} \cdot w_{\tilde{G}}(\Gamma_{\tilde{G}}(S)) \\ &\geq \frac{1}{1+\epsilon} \cdot w_{\tilde{G}}(\text{cut}_{\tilde{G}}(S, V-S)), \end{aligned} \quad (13)$$

where the last inequality holds because $\text{cut}_{\tilde{G}}(S, V-S)$ is just a subset of the edges incident to vertices in $\Gamma_{\tilde{G}}(S)$. From the definition of weighted conductance,

$$\begin{aligned} w_{\tilde{G}}(\text{cut}_{\tilde{G}}(S, V-S)) &\geq \varphi_{\tilde{G}} \cdot \min\{w_{\tilde{G}}(S), w_{\tilde{G}}(V-S)\} \\ &\geq \varphi_{\tilde{G}} \cdot \min\left\{\frac{\text{vol}_G(S)}{1+\epsilon}, \frac{\text{vol}_G(V-S)}{1+\epsilon}\right\}, \quad \text{by (12)} \\ &\geq \frac{\phi_G^2}{4(1+\epsilon)^2} \cdot \min\{\text{vol}_G(S), \text{vol}_G(V-S)\}, \quad \text{by (10)}. \end{aligned}$$

Substituting that to (13) and using the definition of $\beta_{\text{ST},G}$, we obtain (11). \square

6.3 Relating Sparsification to Rumor Spreading (Proof of Theorem 14)

We consider now the following edge sampling process, PP, which imitates the random choices of nodes in PUSH-PULL.

Definition 22 (Process PP). *Let PP be an edge sampling process which selects edges from a given graph $G = (V, E)$ as follows. Let E_u be the set of all edges $e \in E$ incident to vertex u , and let $\Psi > 0$ be an integer parameter of the process. For each vertex $u \in V$, process PP draws Ψ edges from E_u independently at random with replacement; let F_u be the set of edges drawn for u . The total set of edges selected by PP is the union of all sets F_u , $u \in V$.*

Note that graph $\text{PP}(G)$ is unweighted, unlike graph $\text{ST}(G)$.

The next lemma is our main technical claim, showing that w.h.p. the volume expansion of PP is at least equal to that of ST, if $\Psi = \Omega(\Upsilon \cdot \log n)$. The proof shows a coupling of PP and ST under which $\text{ST}(G)$ is a subgraph of $\text{PP}(G)$ w.h.p. For this coupling we consider a variant of ST, denoted ST_2 , which makes two (independent) attempts to select each edge uv : one attempt succeeds with probability $\Upsilon / \deg(u)$, and the other with probability $\Upsilon / \deg(v)$ ⁴ (thus, the total success probability is at least as high as for a single attempt with the probability p_{uv} from Definition 18). We associate the first attempt with vertex u , and the second with v . This allows us to couple ST_2 and PP by considering independently for each vertex u , the selections of the two processes: in ST_2 we select each edge incident to u with probability $\Upsilon / \deg(u)$, while in PP we make $\Psi = \Omega(\Upsilon \cdot \log n)$ independent choices selecting random edges incident to u with replacement; let L_u and F_u be the respective sets of selected edges. We are looking for a coupling that ensures $L_u \subseteq F_u$ w.h.p. We observe that given $|L_u| = k$, the set L_u is just a random k -subset of the edges incident to u , and the same is true for F_u , given $|F_u| = k$. It follows that the desired coupling exists, if $|L_u|$ is stochastically dominated by $|F_u|$. We have that $|L_u|$ is distributed like binomial random variable $B(\deg(u), \Upsilon / \deg(u))$, while $|F_u|$ is distributed like Coupon Collector's random variable $C(\deg(u), \Psi)$ (see Lemma 3). We then use Lemma 3 to show the desired stochastic dominance of $B(\deg(u), \Upsilon / \deg(u))$ by $C(\deg(u), \Psi)$.

Lemma 23 (Volume Expansion of PP). *Let $0 < \epsilon < 1/2$ and $c \geq 1$, and let $G = (V, E)$ be an n -node graph with conductance ϕ . Suppose that we apply the PP process on G with parameter*

$$\Psi \geq \left(\frac{24c \log(3n)}{\epsilon \phi^2} \right)^2 \cdot (4c + 8) \ln n. \quad (14)$$

Then with probability at least $1 - 2n^{-c}$, the volume expansion of PP on G is

$$\beta_{\text{PP}, G} \geq \frac{\phi_G^2}{4(1 + \epsilon)^3}.$$

Proof. We show that there exists a coupling of processes PP and ST, where PP has parameter Ψ satisfying (14) and ST has parameter

$$\Upsilon = \left(\frac{24c \log(3n)}{\epsilon \phi^2} \right)^2,$$

such that with probability at least $1 - n^{-c}$, graph $\tilde{G} := \text{ST}(G)$ is a subgraph of $\hat{G} := \text{PP}(G)$. This implies that with probability at least $1 - n^{-c}$, for any set $S \subseteq V$,

$$\text{vol}_G(\Gamma_{\hat{G}}(S)) \geq \text{vol}_G(\Gamma_{\tilde{G}}(S)).$$

⁴Formally, the probabilities are $\min\{1, \Upsilon / \deg(u)\}$ and $\min\{1, \Upsilon / \deg(v)\}$, respectively, but for simplicity we omit the ‘‘min’’ in this informal overview.

Combining that with Lemma 21, and using a union bound, proves Lemma 23. The rest of the proof demonstrates the coupling we promised above.

Coupling of PP and ST: Let ST' be a process identical to ST except that the probability of selecting edge uv is

$$p'_{uv} := \min \left\{ 1, \frac{\Upsilon}{\max\{\deg(u), \deg(v)\}} \right\}.$$

(Notice that we use $\max\{\deg(u), \deg(v)\}$ instead of \min .) Consider now a process ST_2 which executes independently ST and ST' on G , and selects all edges picked by at least one of ST and ST' . We will show a coupling of PP and ST_2 for which $ST_2(G)$ is a subgraph of $PP(G)$ with probability $1 - n^{-c}$. This immediately implies a coupling of PP and ST with the same property, since there is a trivial coupling of ST_2 and ST for which $ST(G)$ is a subgraph of $ST_2(G)$ (with probability 1).

There is a convenient equivalent way to describe ST_2 , along the lines of Definition 22 of PP : For each vertex $u \in V$, process ST_2 chooses each edge $e \in E_u$ independently with probability

$$\min \left\{ 1, \frac{\Upsilon}{\deg(u)} \right\};$$

let L_u denote the set of edges chosen for u . Then the total set of edges selected by ST_2 is the union of all L_u , $u \in V$. This description is equivalent to the original definition of ST_2 , because each edge $uv \in E$, where $\deg(u) \leq \deg(v)$, is added to L_u with probability p_{uv} and is independently added to L_v with probability p'_{uv} .

Observe that the random sets L_u , $u \in V$, are mutually independent, similar to the sets F_u in the definition of PP . For each $u \in V$, the cardinality of set L_u follows the binomial distribution $B(\deg(u), \min\{1, \frac{\Upsilon}{\deg(u)}\})$. Moreover, for any $0 \leq k \leq \deg(u)$, conditionally on $|L_u| = k$, we have that L_u is a uniformly random k -subset of E_u . For PP on the other hand, we have that $|F_u|$ has the same distribution as $C(\deg(u), \Psi)$, in the statement of Lemma 3, i.e., the number of coupons collected in Ψ trials in the Coupon Collector's process with $\deg(u)$ coupons in total. Similar to L_u , we have that F_u is a uniformly random k -subset of E_u , given $|F_u| = k$.

We use the following coupling of the sets L_u and F_u . For each $u \in V$, we draw a sample ℓ_u from distribution $B(\deg(u), \min\{1, \frac{\Upsilon}{\deg(u)}\})$ and a sample f_u from the distribution of $C(\deg(u), \Psi)$, and independently we choose a random permutation π_u over the edges in E_u .⁵ Then we let L_u be the set of the first ℓ_u edges in permutation π_u , and let F_u be the set of the first f_u edges in π_u . From the discussion above, it follows that the marginal distributions of L_u and F_u are the correct ones.

To show that $ST_2(G)$ is a subgraph of $PP(G)$, i.e., $\bigcup_{u \in V} L_u \subseteq \bigcup_{u \in V} F_u$, it suffices to show that $L_u \subseteq F_u$ for all $u \in V$. And this is equivalent, under the coupling above, to showing that $\ell_u \leq f_u$ for all $u \in V$. We now show that $\ell_u \leq f_u$, for any $u \in V$, w.h.p.

For each vertex $u \in V$, we distinguish two cases depending on $\deg(u)$.

- *Case $\Upsilon/\deg(u) \geq 1/4$:* We apply Equation (1) of Lemma 3, for $k = \deg(u)$ and $t = \Psi$, to get

$$\begin{aligned} \Pr(C(\deg(u), \Psi) < \deg(u)) &\leq \deg(u) \cdot e^{-\Psi/\deg(u)} \\ &\leq n \cdot e^{-\Upsilon \cdot (4c+8) \ln n / \deg(u)} \\ &\leq n \cdot e^{-(4c+8) \ln n / 4}, \quad \text{by the case hypothesis} \\ &= n^{-c-1}. \end{aligned}$$

It follows that $\Pr(f_u < \deg(u)) \leq n^{-c-1}$, and since $\ell_u \leq \deg(u)$, we have

$$\Pr(f_u \geq \ell_u) \geq 1 - n^{-c-1}.$$

⁵Samples ℓ_u and f_u do not need to be independent.

- *Case $\Upsilon/\deg(u) < 1/4$:* Let

$$\lambda = \min\{\deg(u)/2, \Psi/4\}.$$

We apply Equation (2) of Lemma 3, for $k = \deg(u)$, $t = \Psi$ and the above value of λ , to get

$$\begin{aligned} \Pr(C(\deg(u), \Psi) \leq \lambda) &\leq \Pr(B(\Psi, 1 - \lambda/\deg(u)) \leq \lambda) \\ &\leq \Pr(B(\Psi, 1/2) \leq \lambda) \\ &\leq \Pr(B(\Psi, 1/2) \leq \Psi/4). \end{aligned}$$

Using Lemma 4(a) to bound the probability in the last line above, gives

$$\Pr(C(\deg(u), \Psi) \leq \lambda) \leq e^{-2(\Psi/2 - \Psi/4)^2/\Psi} = e^{-\Psi/8}.$$

Therefore,

$$\Pr(f_u \leq \lambda) \leq e^{-\Psi/8}. \quad (15)$$

Next we bound the probability that $\ell_u \geq \lambda$. We have $\lambda \geq 2\Upsilon$, because $2\Upsilon \leq \deg(u)/2$ by the case hypothesis, and $2\Upsilon \leq \Psi/4$ for the chosen values of Υ and Ψ . Then,

$$\begin{aligned} \Pr(B(\deg(u), \Upsilon/\deg(u)) \geq \lambda) &\leq \Pr(B(\deg(u), \Upsilon/\deg(u)) \geq 2\Upsilon), \\ &\leq e^{-\Upsilon/3}, \end{aligned}$$

by Lemma 4(b). Therefore,

$$\Pr(\ell_u \geq \lambda) \leq e^{-\Upsilon/3}.$$

Combining that with (15), and applying a union bound, gives

$$\Pr(f_u > \lambda > \ell_u) \geq 1 - e^{-\Psi/8} - e^{-\Upsilon/3} > 1 - n^{-c-1}.$$

Therefore, in both cases above it holds

$$\Pr(f_u \geq \ell_u) \geq 1 - n^{-c-1}.$$

Then, by a union bound, the probability that $f_u \geq \ell_u$ for all $u \in V$, is at least $1 - n^{-c}$. This completes the proof of Lemma 23. \square

The next claim follows from Lemma 23, and the analogy between the random choices made by PP and PUSH-PULL.

Lemma 24. *Suppose that a set of nodes $S \subseteq V$ in graph $G = (V, E)$ know a message at a given time. Then after applying r rounds of PUSH-PULL, where*

$$r := \left(\frac{24c \log(3n)}{\epsilon \phi^2} \right)^2 \cdot (4c + 8) \ln n, \quad (16)$$

$0 < \epsilon < 1/2$, and $c \geq 1$, the total volume of informed nodes has increased by at least

$$\frac{\phi_G^2}{4(1 + \epsilon)^3} \cdot \min\{\text{vol}(S), \text{vol}(V - S)\},$$

with probability at least $1 - 2n^{-c}$.

Proof. For each $u \in V$, let V_u denote the set of neighbors that u selects during the r rounds of PUSH-PULL. The edge sets $\{uv: v \in V_u\}$, $u \in V$, are distributed identically to the sets F_u selected by PP, when PP is applied to G with parameter $\Psi = r$. It follows that, if S' is the set of vertices $u \in V - S$ that either select or are selected by at least one neighbor from S during PUSH-PULL, then $\text{vol}_G(S')$ has the same distribution as $\text{vol}_G(\Gamma_{\text{PP}(G)}(S))$. By the definition of volume expansion,

$$\text{vol}_G(\Gamma_{\text{PP}(G)}(S)) \geq \beta_{\text{PP},G} \cdot \min\{\text{vol}_G(S), \text{vol}_G(V - S)\},$$

and applying Lemma 23, gives that

$$\text{vol}_G(\Gamma_{\text{PP}(G)}(S)) \geq \frac{\phi_G^2}{4(1+\epsilon)^3} \cdot \min\{\text{vol}_G(S), \text{vol}_G(V - S)\},$$

with probability at least $1 - n^{-c}$. The claim now follows, because all vertices $u \in S'$ get informed during PUSH-PULL, since they communicate with at least one informed neighbor from S . \square

We are now ready to prove our main result.

Proof of Theorem 14. In preparation to apply Lemma 24, we divide rumor spreading into phases of r rounds each, where r is as in (16), i.e.,

$$r := \left(\frac{24c \log(3n)}{\epsilon \phi^2} \right)^2 \cdot (4c + 8) \ln n.$$

Recall that I_t denotes the set of informed nodes after the first t rounds and $U_t = V - I_t$. Let

$$\rho := \frac{\phi^2}{4(1+\epsilon)^3}.$$

For each $i \geq 1$, let X_i be a binary random variable that is 1 if

$$\text{vol}(I_{(i-1)r} - I_{ir}) \geq \rho \cdot \min\{\text{vol}(I_{(i-1)r}), \text{vol}(U_{(i-1)r})\};$$

otherwise, $X_i = 0$. Lemma 24 implies that for each $i \geq 1$, $\Pr(X_i = 1) \geq 1 - 2n^{-c}$. Moreover, it is easy to see that $I_{ir} \neq V$ (i.e., not all nodes are informed after ir rounds) implies

$$\sum_{1 \leq j \leq i} X_j < 2 \ln_{1+\rho}(\text{vol}(V)/2).$$

By a union bound, the probability that $X_j = 1$ for all $j \leq i$ is at least $1 - in^{-c}$. It follows that for $i = 2 \ln_{1+\rho}(\text{vol}(V)/2)$, we have $I_{ir} = V$ with probability at least

$$1 - \ln_{1+\rho}(\text{vol}(V)/2) \cdot n^{-c} = 1 - \Omega(n^{-c} \log n / \rho) = 1 - \Omega(n^{-c+4} \log n),$$

where for the last equation we used that $\phi = \Omega(1/n^2)$ for any connected n -node graph. It follows that, with the above probability, the number of rounds before all nodes get informed is at most

$$r \cdot 2 \ln_{1+\rho}(\text{vol}(V)/2) = \mathcal{O}(\log^4 n / \rho \phi^4) = \mathcal{O}(\log^4 n / \phi^6). \quad \square$$

References

- [1] Y. Azar, A. Z. Broder, A. R. Karlin, and E. Upfal. Balanced allocations. *SIAM Journal on Computing*, 29(1):180–200, 1999.
- [2] J. Batson, D. A. Spielman, N. Srivastava, and S.-H. Teng. Spectral sparsification of graphs: Theory and algorithms. *Communications of the ACM*, 56(8):87–94, 2013.
- [3] P. Berenbrink, R. Elsässer, and T. Friedetzky. Efficient randomised broadcasting in random regular networks with applications in peer-to-peer systems. In *Proc. 27th ACM Symposium on Principles of Distributed Computing (PODC)*, pages 155–164, 2008.
- [4] S. P. Boyd, A. Ghosh, B. Prabhakar, and D. Shah. Gossip algorithms: Design, analysis and applications. *IEEE Transactions on Information Theory*, 52:1653–1664, 2006.
- [5] K. Censor-Hillel, B. Haeupler, J. A. Kelner, and P. Maymounkov. Global computation in a poorly connected world: Fast rumor spreading with no dependence on conductance. In *Proc. 44th ACM Symposium on Theory of Computing (STOC)*, pages 961–970, 2012.
- [6] K. Censor-Hillel and H. Shachnai. Fast information spreading in graphs with large weak conductance. In *Proc. 22nd ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 440–448, 2011.
- [7] J. Cheeger. A lower bound for the smallest eigenvalue of the laplacian. *Problems in Analysis*, pages 195–199, 1970.
- [8] F. Chierichetti, S. Lattanzi, and A. Panconesi. Rumor spreading in social networks. In *Proc. 36th International Colloquium on Automata, Languages and Programming (ICALP)*, pages 375–386, 2009.
- [9] F. Chierichetti, S. Lattanzi, and A. Panconesi. Almost tight bounds for rumour spreading with conductance. In *Proc. 42nd ACM Symposium on Theory of Computing (STOC)*, pages 399–408, 2010.
- [10] F. Chierichetti, S. Lattanzi, and A. Panconesi. Rumour spreading and graph conductance. In *Proc. 21st ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1657–1663, 2010.
- [11] F. R. K. Chung. *Spectral Graph Theory*. CBMS Regional Conference Series in Mathematics. American Mathematical Society, 1997.
- [12] A. J. Demers, D. H. Greene, C. Hauser, W. Irish, J. Larson, S. Shenker, H. E. Sturgis, D. C. Swinehart, and D. B. Terry. Epidemic algorithms for replicated database maintenance. In *Proc. 6th ACM Symposium on Principles of Distributed Computing (PODC)*, 1987.
- [13] B. Doerr, M. Fouz, and T. Friedrich. Social networks spread rumors in sublogarithmic time. In *Proc. 43rd ACM Symposium on Theory of Computing (STOC)*, pages 21–30, 2011.
- [14] B. Doerr, T. Friedrich, and T. Sauerwald. Quasirandom rumor spreading. In *Proc. 19th ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1964–1991, 2008.
- [15] B. Doerr, T. Friedrich, and T. Sauerwald. Quasirandom rumor spreading: Expanders, push vs. pull, and robustness. In *Proc. 36th International Colloquium on Automata, Languages and Programming (ICALP)*, pages 366–377, 2009.

- [16] D. Dubhashi and A. Panconesi. *Concentration of Measure for the Analysis of Randomized Algorithms*. Cambridge University Press, 2009.
- [17] R. Elsässer. On the communication complexity of randomized broadcasting in random-like graphs. In *Proc. 18th ACM Symposium on Parallel Algorithms and Architectures (SPAA)*, pages 148–157, 2006.
- [18] U. Feige, D. Peleg, P. Raghavan, and E. Upfal. Randomized broadcast in networks. *Random Structures and Algorithms*, 1(4):447–460, 1990.
- [19] N. Fountoulakis, K. Panagiotou, and T. Sauerwald. Ultra-fast rumor spreading in social networks. In *Proc. 23rd ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1642–1660, 2012.
- [20] A. Frieze and G. Grimmett. The shortest-path problem for graphs with random arc-lengths. *Discrete Applied Mathematics*, 1(10):57–77, 1985.
- [21] G. Giakkoupis. Tight bounds for rumor spreading in graphs of a given conductance. In *Proc. 28th International Symposium on Theoretical Aspects of Computer Science (STACS)*, pages 57–68, 2011.
- [22] G. Giakkoupis. Tight bounds for rumor spreading with vertex expansion. In *Proc. 25th ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 801–815, 2014.
- [23] G. Giakkoupis and T. Sauerwald. Rumor spreading and vertex expansion. In *Proc. 23rd ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1623–1641, 2012.
- [24] G. Giakkoupis, T. Sauerwald, and A. Stauffer. Randomized rumor spreading in dynamic graphs. In *Proc. 41st International Colloquium on Automata, Languages and Programming (ICALP)*, pages 495–507, 2014.
- [25] G. Giakkoupis, T. Sauerwald, H. Sun, and P. Woelfel. Low randomness rumor spreading via hashing. In *Proc. 29th International Symposium on Theoretical Aspects of Computer Science (STACS)*, pages 314–425, 2012.
- [26] B. Haeupler. Simple, fast and deterministic gossip and rumor spreading. In *Proc. 24th ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 705–716, 2013.
- [27] M. Jerrum and A. Sinclair. Approximating the permanent. *SIAM Journal on Computing*, 18(6):1149–1178, 1989.
- [28] R. Karp, C. Schindelhauer, S. Shenker, and B. Vöcking. Randomized rumor spreading. In *Proc. 41st IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 565–574, 2000.
- [29] J. Leskovec, K. J. Lang, A. Dasgupta, and M. W. Mahoney. Statistical properties of community structure in large social and information networks. In *Proc. 17th International Conference on World Wide Web (WWW)*, pages 695–704, 2008.
- [30] M. Mihail, C. H. Papadimitriou, and A. Saberi. On certain connectivity properties of the Internet topology. *Journal of Computer and System Sciences*, 72(2):239–251, 2006.
- [31] M. Mitzenmacher and E. Upfal. *Probability and Computing: Randomization and Probabilistic Techniques in Algorithms and Data Analysis*. Cambridge University Press, 2005.

- [32] D. Mosk-Aoyama and D. Shah. Fast distributed algorithms for computing separable functions. *IEEE Transactions on Information Theory*, 54(7):2997–3007, 2008.
- [33] B. Pittel. On spreading a rumor. *SIAM Journal on Applied Mathematics*, 47:213–223, 1987.
- [34] T. Sauerwald. On mixing and edge expansion properties in randomized broadcasting. *Algorithmica*, 56(1):51–88, 2010.
- [35] T. Sauerwald and A. Stauffer. Rumor spreading and vertex expansion on regular graphs. In *Proc. 22nd ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 462–475, 2011.
- [36] D. A. Spielman and S. H. Teng. Nearly-linear time algorithms for graph partitioning, graph sparsification, and solving linear systems. In *Proc. 36th ACM Symposium on Theory of Computing (STOC)*, pages 81–90, 2004.
- [37] D. A. Spielman and S. H. Teng. Spectral sparsification of graphs. *SIAM Journal on Computing*, 40(4):981–1025, 2011.