

Tight bounds for rumor spreading in graphs of a given conductance*

George Giakkoupis
LIAFA, Université Paris 7
ggiak@liafa.jussieu.fr

Abstract

We study the connection between the rate at which a rumor spreads throughout a graph and the *conductance* of the graph—a standard measure of a graph’s expansion properties. We show that for any n -node graph with conductance ϕ , the classical PUSH-PULL algorithm distributes a rumor to all nodes of the graph in $O(\phi^{-1} \log n)$ rounds with high probability (w.h.p.). This bound improves a recent result of Chierichetti, Lattanzi, and Panconesi [6], and it is tight in the sense that there exist graphs where $\Omega(\phi^{-1} \log n)$ rounds of the PUSH-PULL algorithm are required to distribute a rumor w.h.p.

We also explore the PUSH and the PULL algorithms, and derive conditions that are both necessary and sufficient for the above upper bound to hold for those algorithms as well. An interesting finding is that every graph contains a node such that the PULL algorithm takes $O(\phi^{-1} \log n)$ rounds w.h.p. to distribute a rumor started at that node. In contrast, there are graphs where the PUSH algorithm requires significantly more rounds for any start node.

1. Introduction

Gossip-based algorithms have become a prominent paradigm for designing simple, efficient, and robust protocols for disseminating information in large networks. Perhaps the most basic and most well-studied example of a gossip-based information-dissemination algorithm is the, so-called, *rumor-spreading model*. The algorithm proceeds in a sequence of synchronous rounds. Initially, in round 0, an arbitrary *start* node receives a piece of information, called the *rumor*. This rumor is then spread iteratively to other nodes: In each round, every *informed* node (i.e., every node that received the rumor in a previous round) chooses a random neighbor to which it transmits the rumor. This is the *PUSH* version of the rumor-spreading model. The *PULL* version is symmetric: In each round, every *uninformed* node chooses a random neighbor, and if that neighbor knows the rumor it transmits it to the uninformed node. Finally, the *PUSH-PULL* algorithm is the combination of both strategies: In each round, every node chooses a random neighbor to transmit the rumor to, if the node knows the rumor, or to request the rumor from, otherwise.

The above three rumor-spreading algorithms were proposed in [8], in the context of maintaining distributed replicated database systems. Subsequently, these algorithms (and variations of them) have been used in various applications, such as failure detection [27], resource discovery [21], and

*This work was supported by the ANR project PROSE, the European project EULER, and the INRIA project GANG.

data aggregation [3]. Also, their performance has been studied theoretically for several classes of networks (see the Related Work Section).

In this paper, we investigate the relationship between the performance of rumor spreading in an arbitrary connected network (represented by an undirected graph), and the expansion properties of the network. More precisely, we look at the *broadcast time* of the above rumor-spreading algorithms, i.e., the number of rounds until all nodes get informed—a primary measure for rumor spreading. And we study its connection to the *conductance* of the network—one of the most studied measures of graph expansion. Roughly speaking, the conductance of a connected graph is a value ϕ in the range $0 < \phi \leq 1$, which is large for graphs that are well connected (e.g., the complete graph), and small for graphs that are not (e.g., graphs with communication bottlenecks).

A connection between broadcast time and conductance has been observed in several works, e.g., in [1, 10, 23, 25], where upper bounds on the broadcast time were obtained for various graph topologies based, essentially, on lower bounds on the conductance. In [5], Chierichetti, Lattanzi, and Panconesi posed the question whether rumor spreading is fast in *all* graphs with high conductance. For the PUSH and the PULL algorithms the answer is negative; as observed in [6], a star with n vertices has constant conductance but the expected broadcast time for a random start node is $\Omega(n)$ rounds. For the PUSH-PULL algorithm, however, the answer to the above question is positive. In [7], it was shown that for any graph and any start node, the broadcast time of the PUSH-PULL algorithm is $O(\phi^{-6} \log^4 n)$ rounds, with high probability (*w.h.p.*).¹ It was also noted in [7] that this result suggests a justification as to why rumors spread quickly among humans, since experimental studies have shown that social networks have high conductance. The above bound was subsequently improved to $O((\log \phi^{-1})^2 \phi^{-1} \log n)$ rounds w.h.p., in [6]. Further, it was shown there that this bound is by at most a $(\log \phi^{-1})^2$ -factor larger than the optimal bound. More precisely, it was shown that for any $\phi \geq 1/n^{1-\epsilon}$, there are n -node graphs with conductance at least ϕ and diameter $\Omega(\phi^{-1} \log n)$. Finally, the authors of [6] provided a sufficient condition for their upper bound to hold for the PUSH and the PULL algorithms as well. This condition states that for any edge, the ratio of the degrees of its two endpoints is bounded by a constant.

Two other important measures of a graph’s expansion properties are *edge* and *vertex expansion*. The authors of [5] described a graph with constant edge expansion in which the expected broadcast time of the PUSH-PULL algorithm for a random start node is $\Omega(\sqrt{n})$. The question whether high vertex expansion yields fast rumor spreading (also posed in [5]) is largely open; in a very recent work [26], it was shown that for *regular* graphs this is true.

Our Contributions. We saw that an upper bound of $O((\log \phi^{-1})^2 \phi^{-1} \log n)$ rounds w.h.p. is known for the broadcast time of the PUSH-PULL algorithm in *any* graph; and $\Omega(\phi^{-1} \log n)$ rounds are required for *some* graph with n nodes and conductance ϕ , for any n and $\phi \geq 1/n^{1-\epsilon}$. Our first contribution is the following result, which closes the gap between these two bounds.

Theorem 1.1. *For any graph on n vertices and any start vertex, the broadcast time of the PUSH-PULL algorithm is $O(\phi^{-1} \log n)$ rounds w.h.p.*

We also show that Theorem 1.1 is tight for $\phi = \Omega(1/n)$ —not just for $\phi \geq 1/n^{1-\epsilon}$ as it was previously known. Clearly, the theorem is not tight for $\phi = o(1/n)$, since the broadcast time of the PUSH algorithm is known to be $O(n \log n)$ rounds w.h.p. for any graph [16].

The proof of Theorem 1.1 is based on an analysis of the PUSH and the PULL algorithms. We show that in any graph, the broadcast time of the PULL algorithm is $O(\phi^{-1} \log n)$ rounds w.h.p., *if the start node has degree Δ* , the maximum degree of the graph. Also, based on the symmetry

¹ By “with high probability” we mean with probability $1 - O(n^{-c})$, for an arbitrary constant $c > 0$.

between the PULL and the PUSH algorithms, we show that for any start node, the PUSH algorithm takes $O(\phi^{-1} \log n)$ rounds w.h.p. to inform a node of degree Δ . Therefore, w.h.p. the PUSH-PULL algorithm takes $O(\phi^{-1} \log n)$ rounds to inform a node of degree Δ , and $O(\phi^{-1} \log n)$ additional rounds to inform the remaining nodes.

Our analysis is different than previous approaches. Specifically, the proof in [7] is based on a connection between rumor spreading and a spectral sparsification process; and the proof in [6] analyzes the PUSH-PULL process directly. Still, our analysis uses some ideas from [6].

Recall that high conductance does not always yield short broadcast times for the PUSH and the PULL algorithms. Our second contribution is that we derive conditions guaranteeing a broadcast time of $O(\phi^{-1} \log n)$ rounds w.h.p. for those algorithms. As mentioned above, in the proof of Theorem 1.1 we show that one such condition for the PULL algorithm is that the start node have degree Δ . We extend this result as follows. Let δ denote the minimum degree of the graph.

Theorem 1.2. (a) *For any graph on n vertices and any start vertex with degree $\Omega(\Delta(\phi + \delta^{-1}))$, the broadcast time of the PULL algorithm is $O(\phi^{-1} \log n)$ rounds w.h.p. (b) If, in particular, $\Delta = O(1/\phi)$ then the above bound on the broadcast time holds for any start vertex.*

Further, we show that the conditions specified in Theorem 1.2 are optimal, in the sense that for any given ϕ , δ , Δ , d with $\Delta = \omega(1/\phi)$ and $d = o(\Delta(\phi + \delta^{-1}))$, there is a graph with those ϕ , δ , Δ , and with a start node of degree d such that the broadcast time of the PULL algorithm is $\omega(\phi^{-1} \log n)$ with non-negligible probability (i.e., with probability $n^{-o(1)}$).

Note that Theorem 1.2(a) does not hold for the PUSH algorithm: a star on n vertices has constant conductance, but the broadcast time of the PUSH algorithm is at least $n - 1$.

From Theorem 1.2(a) it follows that if $\delta = \Omega(\Delta(\phi + \delta^{-1}))$ then the broadcast time of the PULL algorithm is $O(\phi^{-1} \log n)$ w.h.p. for *all* start nodes. This, and Theorem 1.2(b), are also true for the PUSH algorithm, by the symmetry argument used in the proof of Theorem 1.1.

Finally, we also tighten the result of [6] for the PUSH and the PULL algorithms. We show that if, for any edge, the ratio of the degrees of its endpoints is bounded, then the broadcast time of those algorithms is $O(\phi^{-1} \log n)$ rounds w.h.p., for any start node.

Related Work. The broadcast time of the PUSH algorithm has been analyzed for various graph topologies, including the complete graph [19, 24], the hypercube and random graphs [16], star and Cayley graphs [12, 13], regular graphs [15], and random regular graphs [17].

Besides the broadcast time, another performance measure of interest is the total number of transmissions of the rumor. Fewer transmissions are typically achieved using the PUSH-PULL algorithm. The broadcast time and the number of transmissions of the PUSH-PULL algorithm (and variations of it) have been analyzed for the complete graph [22], random graphs [11, 14], and random regular graphs [1]. The problem of minimizing the total communication complexity (i.e., the total number of bits transmitted) was studied in [18] for the complete graph.

A quasi-random variant of the rumor-spreading model was proposed in [9], as a means to reduce the amount of randomness. In the quasi-random model, each node has a (cyclic) list of its neighbors in which it just chooses a random starting position—instead of choosing a new random position in each round. This model was shown to be at least as efficient as the classical rumor-spreading model for several families of graphs [9, 10]. The problem of further reducing the amount of randomness was studied in [20].

The problem of rumor spreading in arbitrary graphs and its connection to the graph's expansion properties were also studied in [3, 23], in the context of gossip-based data aggregation. In both papers, the data-aggregation protocols proposed employ generalizations of the PUSH-PULL

algorithm with non-uniform selection probabilities: in each round, node v chooses its neighbor u with probability $p_{v,u}$. Under certain symmetry conditions for the matrix of $p_{v,u}$, upper bounds on the broadcast time were established, as a function of certain measures of this matrix that resemble graph conductance. These results, however, are not directly comparable to our results. In particular, as observed in [6], there are graphs with high conductance for which the above approaches yield large bounds for the broadcast time.

The problem of *partial* rumor spreading, where it suffices that the rumor be spread to a constant fraction of the nodes, was studied in [4]. There, a refinement of graph conductance, called *weak conductance*, was introduced, and it was shown that high weak conductance always implies fast partial rumor spreading (using the PUSH-PULL algorithm), even if the (standard) conductance is small.

Paper organization. We begin with some definitions and notations, in Section 2. Section 3, which constitutes the largest part of the paper, contains the analysis of the PULL algorithm, including the proof of Theorem 1.2. In Section 4, we provide a result on the symmetry between the PUSH and PULL algorithms, which allows us to derive the properties of the PUSH algorithm from the analysis of the PULL algorithm. Finally, in Section 5, we analyze the PUSH-PULL algorithm and prove Theorem 1.1 using results from Sections 3 and 4.

2. Preliminaries

We consider an arbitrary connected network, represented by an undirected graph $G = (V, E)$. The degree of a vertex $v \in V$ is denoted $d(v)$. By Δ we denote the maximum degree of G , $\Delta = \max_{v \in V} d(v)$, and by δ we denote the minimum degree. The *volume* of a subset of vertices $S \subseteq V$ is the sum of the degrees of the vertices in S , $\text{vol}(S) = \sum_{v \in S} d(v)$. Note that $\text{vol}(V) = 2|E|$. By $\text{cut}(S, V - S)$ we denote the set of edges crossing the partition $\{S, V - S\}$ of V , i.e., $\text{cut}(S, V - S) = \{\{v, u\} \in E : v \in S, u \in V - S\}$. The *conductance* ϕ of G is defined as

$$\phi = \min_{S \subseteq V, \text{vol}(S) \leq |E|} \frac{|\text{cut}(S, V - S)|}{\text{vol}(S)}.$$

It is easy to see that $0 < \phi \leq 1$. (It is $\phi \neq 0$ because graph G is connected.) Also,

Observation 2.1. For any $S \subseteq V$, $|\text{cut}(S, V - S)| \geq \lceil \phi \cdot \min\{\text{vol}(S), \text{vol}(V - S)\} \rceil$.

We will denote by S_i the set of informed vertices at the *end* of round i of the rumor-spreading algorithm, and by U_i the set of uninformed vertices at that time, $U_i = V - S_i$. S_0 and U_0 denote the corresponding sets initially. To simplify notation, we will assume that S_0 can be any non-empty subset of vertices—we do not require that $|S_0| = 1$.

3. PULL Algorithm

In Section 3.1, we establish a general upper bound on the broadcast time of the PULL algorithm, for any initial set of informed vertices. In Section 3.2, we build upon and refine this result to derive conditions that guarantee broadcast times of $O(\phi^{-1} \log n)$ rounds. More precisely, we prove Theorem 1.2 and demonstrate its optimality, and we show that a condition proposed in [6] also achieves the above broadcast time.

3.1. An Upper Bound on the Broadcast Time

The main result of this section is the following high-probability bound on the broadcast time for an arbitrary initial set of informed vertices. Recall that Δ is the maximum degree of G .

Lemma 3.1. *For any initial set of informed vertices $S_0 \subseteq V$ and any fixed $\beta > 0$, all vertices get informed in at most $50(\beta + 2) \log n (\phi^{-1} + \Delta / \lceil \phi \text{vol}(S_0) \rceil)$ rounds of the PULL algorithm, with probability $1 - O(n^{-\beta})$.*

Note that if $\Delta / \lceil \phi \text{vol}(S_0) \rceil = O(1/\phi)$ then the broadcast time is $O(\phi^{-1} \log n)$ w.h.p.

To prove Lemma 3.1 we divide the execution of the algorithm into three phases: The first phase lasts until the total volume of informed vertices becomes at least Δ ; the second lasts until this volume exceeds $|E|$, i.e., it exceeds one half of the total volume of the graph; and the third lasts until all vertices get informed. We measure progress in the first two phases by the increase in the volume of informed vertices; and in the third phase by the decrease in the volume of uninformed vertices. For each phase, the next lemma gives upper bounds on the number of rounds until “significant” progress is made with constant probability.

Lemma 3.2.

- (a) *If $\text{vol}(S_0) < \Delta$ then $\Pr(\text{vol}(S_i) \geq \Delta) \geq 1/2$, for $i \geq 4\Delta / \lceil \phi \text{vol}(S_0) \rceil$.*
- (b) *If $\Delta \leq \text{vol}(S_0) \leq |E|$ then $\Pr(\text{vol}(S_i) \geq \min\{2 \text{vol}(S_0), |E| + 1\}) \geq 1/2$, for $i \geq 4/\phi$.*
- (c) *If $\text{vol}(S_0) > |E|$ then $\Pr(\text{vol}(U_i) \leq \text{vol}(U_0)/2) \geq 1/2$, for $i \geq 6/\phi$.*

The proof of Lemma 3.2 proceeds as follows. Consider part (a)—for parts (b) and (c) the reasoning is similar. Consider round i . At the beginning of the round there are at least $\phi \text{vol}(S_{i-1}) \geq \phi \text{vol}(S_0)$ edges between informed and uninformed vertices. We fix $\lceil \phi \text{vol}(S_0) \rceil$ of these edges arbitrarily before round i is executed, and then count the total volume L_i of the vertices that get informed in round i due to the rumor being transmitted *through those edges*. Clearly, L_i is a lower bound on the total volume of the vertices informed in round i . Thus, to prove (a) it suffices to show that $\sum_{k \leq i} L_k \geq \Delta - \text{vol}(S_0)$ with probability at least $1/2$. By employing a martingale argument we compute the expectation and the variance of $\sum_{k \leq i} L_k$, and then we bound $\sum_{k \leq i} L_k$ using Chebyshev’s inequality.

The approach used to prove Lemma 3.2 is at the heart of our analysis, and it is also used to prove analogous results in Section 3.2.

Proof of Lemma 3.2. (a) Let L_1, L_2, \dots be a sequence of random variables with L_i , for $i \geq 1$, be defined as follows. We distinguish two cases:

- If $\text{vol}(S_{i-1}) \leq |E|$, then, by Observation 2.1, $|\text{cut}(S_{i-1}, U_{i-1})| \geq \lceil \phi \text{vol}(S_{i-1}) \rceil \geq \lceil \phi \text{vol}(S_0) \rceil$. Let E_i be an arbitrary subset of $\text{cut}(S_{i-1}, U_{i-1})$ consisting of $M = \lceil \phi \text{vol}(S_0) \rceil$ edges. Set E_i is (arbitrarily) fixed at the *beginning* of round i —before the round is executed. Then L_i is the total volume of the vertices that get informed in round i as a result of the rumor being transmitted through edges in E_i . Formally, for each vertex $u \in U_{i-1}$, let $L_{i,u}$ be the 0/1 random variable with $L_{i,u} = 1$ if and only if in round i vertex u receives the rumor through some edge in E_i . Then, $L_i = \sum_{u \in U_{i-1}} L_{i,u} d(u)$.
- If $\text{vol}(S_{i-1}) > |E|$, then $L_i = M$.

We will show the following results for the expectation and the variance of the sum of L_i .

Claim 3.3. $\mathbf{E}[\sum_{k \leq i} L_k] = iM$ and $\mathbf{Var}(\sum_{k \leq i} L_k) \leq iM\Delta$.

Using this claim, the lemma follows by Chebyshev's inequality: Let $\mu = \mathbf{E}[\sum_{k \leq i} L_k] = iM$. Note that for $i \geq 4\Delta/M$, $\mu > \Delta$. So,

$$\Pr\left(\sum_{k \leq i} L_k < \Delta\right) \leq \Pr\left(\left|\sum_{k \leq i} L_k - \mu\right| > \mu - \Delta\right) \leq \frac{\mathbf{Var}(\sum_{k \leq i} L_k)}{(\mu - \Delta)^2} \leq \frac{iM\Delta}{(iM - \Delta)^2} < 1/2, \quad (3.1)$$

for $i \geq 4\Delta/M$. Note that if $\text{vol}(S_i) < \Delta$ then $\sum_{k \leq i} L_k < \Delta$, because $\sum_{k \leq i} L_k$ cannot be larger than the total volume of all vertices informed since round 1 and thus $\sum_{k \leq i} L_k \leq \text{vol}(S_i) - \text{vol}(S_0) < \Delta$. Hence, $\Pr(\text{vol}(S_i) < \Delta) \leq \Pr(\sum_{k \leq i} L_k < \Delta) < 1/2$, for $i \geq 4\Delta/M$.

To complete the proof of part (a) it remains to show Claim 3.3, which we do next.

Expectation of the Sum of L_i : For $i \geq 0$, define $\mathcal{L}_i = \sum_{k \leq i} (L_k - M)$. Let \mathcal{F}_i be the σ -algebra generated by all the choices of the algorithm in the first i rounds. It is easy to see that the sequence $\mathcal{L}_0, \mathcal{L}_1, \dots$ is a martingale with respect to the filter $\mathcal{F}_0 \subseteq \mathcal{F}_1 \subseteq \dots$:

- If $\text{vol}(S_{i-1}) \leq |E|$,

$$\begin{aligned} \mathbf{E}[\mathcal{L}_i - \mathcal{L}_{i-1} \mid \mathcal{F}_{i-1}] &= \mathbf{E}[L_i - M \mid \mathcal{F}_{i-1}] = \mathbf{E}\left[\sum_{u \in U_{i-1}} L_{i,u} d(u) \mid \mathcal{F}_{i-1}\right] - M \\ &= \sum_{u \in U_{i-1}} \mathbf{E}[L_{i,u} \mid \mathcal{F}_{i-1}] \cdot d(u) - M, \end{aligned}$$

where the last relation holds because U_{i-1} is \mathcal{F}_{i-1} -measurable. For any $u \in U_{i-1}$,

$$\mathbf{E}[L_{i,u} \mid \mathcal{F}_{i-1}] = \Pr(L_{i,u} = 1 \mid \mathcal{F}_{i-1}) = g_i(u)/d(u), \quad (3.2)$$

where $g_i(u)$ is the number of edges in E_i that are incident to u . Note that

$$\sum_{u \in U_{i-1}} g_i(u) = |E_i| = M, \quad (3.3)$$

since each edge in E_i is incident to exactly one $u \in U_{i-1}$. Combining the above yields $\mathbf{E}[\mathcal{L}_i - \mathcal{L}_{i-1} \mid \mathcal{F}_{i-1}] = \sum_{u \in U_{i-1}} g_i(u) - M = M - M = 0$.

- If $\text{vol}(S_{i-1}) > |E|$, then $\mathcal{L}_i - \mathcal{L}_{i-1} = L_i - M = M - M = 0$.

So, in both cases, $\mathbf{E}[\mathcal{L}_i - \mathcal{L}_{i-1} \mid \mathcal{F}_{i-1}] = 0$, which yields $\mathbf{E}[\mathcal{L}_i] = \mathbf{E}[\mathcal{L}_0] = 0$. Substituting to this the definition of \mathcal{L}_i , we obtain the desired formula for the expectation, $\mathbf{E}[\sum_{k \leq i} L_k] = iM$.

Variance of the Sum of L_i :

$$\begin{aligned} \mathbf{E}[\mathcal{L}_i^2 \mid \mathcal{F}_{i-1}] &= \mathbf{E}[(\mathcal{L}_i - \mathcal{L}_{i-1}) + \mathcal{L}_{i-1}]^2 \mid \mathcal{F}_{i-1}] \\ &= \mathbf{E}[(\mathcal{L}_i - \mathcal{L}_{i-1})^2 \mid \mathcal{F}_{i-1}] + \mathcal{L}_{i-1}^2 + 2\mathbf{E}[\mathcal{L}_i - \mathcal{L}_{i-1} \mid \mathcal{F}_{i-1}] \cdot \mathcal{L}_{i-1} \\ &= \mathbf{E}[(L_i - M)^2 \mid \mathcal{F}_{i-1}] + \mathcal{L}_{i-1}^2, \end{aligned} \quad (3.4)$$

since $\mathbf{E}[\mathcal{L}_i - \mathcal{L}_{i-1} \mid \mathcal{F}_{i-1}] = 0$. We bound $\mathbf{E}[(L_i - M)^2 \mid \mathcal{F}_{i-1}]$ as follows:

- If $\text{vol}(S_{i-1}) \leq |E|$, then, by the definition of L_i and Equation (3.3),

$$\begin{aligned} \mathbf{E}[(L_i - M)^2 | \mathcal{F}_{i-1}] &= \mathbf{E}\left[\left(\sum_{u \in U_{i-1}} (L_{i,u}d(u) - g_i(u))\right)^2 \middle| \mathcal{F}_{i-1}\right] \\ &= \sum_{u \in U_{i-1}} \mathbf{E}[(L_{i,u}d(u) - g_i(u))^2 | \mathcal{F}_{i-1}], \end{aligned}$$

where the last relation holds because $\mathbf{E}[(L_{i,u}d(u) - g_i(u))(L_{i,u'}d(u') - g_i(u')) | \mathcal{F}_{i-1}] = 0$, for any $u, u' \in U_{i-1}$ with $u \neq u'$. This last statement is true because, by (3.2), $\mathbf{E}[L_{i,u}d(u) - g_i(u) | \mathcal{F}_{i-1}] = 0$, and because the random variables $L_{i,u}$, $u \in U_{i-1}$, are mutually independent conditionally on \mathcal{F}_{i-1} . Using again (3.2) and (3.3), we get

$$\begin{aligned} \sum_{u \in U_{i-1}} \mathbf{E}[(L_{i,u}d(u) - g_i(u))^2 | \mathcal{F}_{i-1}] &= \sum_{u \in U_{i-1}} (\mathbf{E}[L_{i,u}^2(d(u))^2 | \mathcal{F}_{i-1}] - (g_i(u))^2) \\ &\leq \sum_{u \in U_{i-1}} \mathbf{E}[L_{i,u}^2(d(u))^2 | \mathcal{F}_{i-1}] = \sum_{u \in U_{i-1}} \mathbf{E}[L_{i,u} | \mathcal{F}_{i-1}] \cdot (d(u))^2 \\ &= \sum_{u \in U_{i-1}} g_i(u)d(u) \leq \sum_{u \in U_{i-1}} g_i(u)\Delta = M\Delta. \end{aligned}$$

Therefore, $\mathbf{E}[(L_i - M)^2 | \mathcal{F}_{i-1}] \leq M\Delta$.

- If $\text{vol}(S_{i-1}) > |E|$, the last inequality is still true, since $L_i = M$.

By applying the above to (3.4) yields $\mathbf{E}[\mathcal{L}_i^2 | \mathcal{F}_{i-1}] \leq M\Delta + \mathcal{L}_{i-1}^2$, and recursively we obtain $\mathbf{E}[\mathcal{L}_i^2] \leq iM\Delta$. The desired bound for $\mathbf{Var}(\sum_{k \leq i} L_k)$ then follows by observing that $\mathbf{Var}(\sum_{k \leq i} L_k) = \mathbf{E}[\mathcal{L}_i^2]$. This completes the proof of Claim 3.3, and of Lemma 3.2(a).

(b) We consider the same sequence of random variables L_1, L_2, \dots as in part (a). Similarly to (3.1), by using Claim 3.3 and Chebyshev's inequality we obtain that

$$\Pr\left(\sum_{k \leq i} L_k < \text{vol}(S_0)\right) \leq \frac{iM\Delta}{(iM - \text{vol}(S_0))^2} \leq \frac{iM \text{vol}(S_0)}{(iM - \text{vol}(S_0))^2} < 1/2,$$

for $i \geq 4 \text{vol}(S_0)/M$, and thus, for $i \geq 4/\phi$. Part (b) then follows by observing that if $\text{vol}(S_i) < \min\{2 \text{vol}(S_0), |E| + 1\}$ then $\sum_{k \leq i} L_k \leq \text{vol}(S_i) - \text{vol}(S_0) < \text{vol}(S_0)$, and thus, $\Pr(\text{vol}(S_i) < \min\{2 \text{vol}(S_0), |E| + 1\}) \leq \Pr(\sum_{k \leq i} L_k < \text{vol}(S_0))$.

(c) Unlike in parts (a) and (b), the set of uninformed vertices has now a smaller volume than the set of informed vertices. So, by Observation 2.1, $|\text{cut}(S_i, U_i)| \geq \lceil \phi \text{vol}(U_i) \rceil$. We consider the sequence L_1, L_2, \dots of random variables, with L_i defined as follows:

- If $\text{vol}(U_{i-1}) > \text{vol}(U_0)/2$, we let E_i be an arbitrary subset of $\text{cut}(S_{i-1}, U_{i-1})$ consisting of $M = \lceil \phi \text{vol}(U_0)/2 \rceil$ edges. (E_i is fixed at the beginning of round i .) As before, L_i is the total volume of the vertices that get informed in round i as a result of the rumor being transmitted through edges in E_i .
- If $\text{vol}(U_{i-1}) \leq \text{vol}(U_0)/2$, then $L_i = M$.

Similarly to Claim 3.3, we can show that $\mathbf{E}[\sum_{k \leq i} L_k] = iM$ and $\mathbf{Var}(\sum_{k \leq i} L_k) \leq iM \text{vol}(U_0)$. For the latter we use the fact that the degree of any uninformed vertex is at most $\text{vol}(U_0)$. As before, by Chebyshev's inequality, we can show that $\Pr(\sum_{k \leq i} L_k < \text{vol}(U_0)/2) < 1/2$, for $i \geq 6/\phi$. Part (c) then follows by observing that if $\text{vol}(U_i) > \text{vol}(U_0)/2$ then $\sum_{k \leq i} L_k \leq \text{vol}(U_0) - \text{vol}(U_i) < \text{vol}(U_0)/2$. \square

Using the bounds of Lemma 3.2, Lemma 3.1 follows easily:

Proof of Lemma 3.1. By Lemma 3.2(a), if $\text{vol}(S_i) < \Delta$ then, with probability $1/2$, it takes at most $\lceil 4\Delta/\lceil \phi \text{vol}(S_i) \rceil \rceil \leq 5\Delta/\lceil \phi \text{vol}(S_i) \rceil$ additional rounds until the total volume of informed vertices becomes at least Δ . Thus, if $\text{vol}(S_0) < \Delta$, the probability that $\text{vol}(S_t) < \Delta$ for $t = 2\beta \ln n \cdot (5\Delta/\lceil \phi \text{vol}(S_0) \rceil)$ is at most $(1 - 1/2)^{2\beta \ln n} \leq e^{-2\beta \ln n/2} = n^{-\beta}$.

By Lemma 3.2(b), if $\Delta \leq \text{vol}(S_i) \leq |E|$ then, with probability $1/2$, it takes at most $\lceil 4/\phi \rceil$ rounds until the total volume of informed vertices is increased to at least $\min\{2\text{vol}(S_i), |E| + 1\}$. Now, divide the execution of the algorithm into phases of $\lceil 4/\phi \rceil$ rounds each, starting from the end of the first round i with $\text{vol}(S_i) \geq \Delta$. A phase is *successful* if the total volume of informed vertices at the end of the phase is at least $\min\{2\text{vol}(S), |E| + 1\}$, where S is the set of informed vertices at the beginning of the phase. (Note that if $\text{vol}(S) \geq |E| + 1$ then the phase is always successful.) Then, for any k , the probability that the k -th phase is successful is at least $1/2$, regardless of the outcome of the previous $k - 1$ phases. From this (and a simple coupling argument), the number of successful phases among the first k phases is (stochastically) greater or equal to the binomial random variable $B(k, 1/2)$. So, by Chernoff bounds, the probability that fewer than $m = \log |E|$ of the first $k = (2\beta + 4)m$ phases are successful is at most equal to

$$\Pr(B(k, 1/2) < m) = \Pr(k/2 - B(k, 1/2) > k/2 - m) \leq e^{-2(k/2 - m)^2/k} \leq e^{-\beta m} = O(n^{-\beta}),$$

since $|E| \geq n - 1$. And since at most m successful phases are required until the total volume of informed vertices exceeds $|E|$, it follows that with probability $1 - O(n^{-\beta})$ the number of rounds required is at most $k \lceil 4/\phi \rceil = (2\beta + 4) \log(|E|) \lceil 4/\phi \rceil \leq (2\beta + 4)(2 \log n)(5/\phi)$.

Finally, by Lemma 3.2(c), if $\text{vol}(S_i) > |E|$ then, with probability $1/2$, it takes at most $\lceil 6/\phi \rceil$ rounds until the total volume of uninformed vertices is halved. By similar reasoning as before, we can show that once the volume of informed vertices has exceeded $|E|$, then $(2\beta + 4)(2 \log n)(7/\phi)$ rounds suffice to inform all nodes with probability $1 - O(n^{-\beta})$.

Combing all the above and applying the union bound, we obtain that with probability $1 - O(n^{-\beta})$ all vertices get informed within $50(\beta + 2) \log n (\phi^{-1} + \Delta/\lceil \phi \text{vol}(S_0) \rceil)$ rounds. \square

3.2. Conditions for Rumor Spreading in $O(\phi^{-1} \log n)$ Rounds

3.2.1. Derivation of Theorem 1.2

Lemma 3.1 implies that if $\Delta/\lceil \phi \text{vol}(S_0) \rceil = O(1/\phi)$, the broadcast time is $O(\phi^{-1} \log n)$ rounds w.h.p. Theorem 1.2(b) follows then directly, since $\Delta/\lceil \phi \text{vol}(S_0) \rceil \leq \Delta$, for any S_0 . Also, the following weaker version of Theorem 1.2(a) is immediate, because if the degree of the start vertex is $\Omega(\Delta)$ then $\text{vol}(S_0) = \Omega(\Delta)$ and $\Delta/\lceil \phi \text{vol}(S_0) \rceil \leq \Delta/\phi \text{vol}(S_0) = O(1/\phi)$.

Corollary 3.4. *For any start vertex of degree $\Omega(\Delta)$, the broadcast time of the PULL algorithm is $O(\phi^{-1} \log n)$ rounds w.h.p.*

This result is weaker than Theorem 1.2(a) because $\phi + \delta^{-1} = O(1)$. However, it will suffice for the purposes of proving Theorem 1.1 (in Section 5).

Next we describe the proof of Theorem 1.2(a). Recall that Lemma 3.1, on which the proof of Corollary 3.4 was based, assumes that S_0 may be any subset of vertices. Under this assumption, the size of $\text{cut}(S_0, U_0)$ can be as small as $\lceil \phi \text{vol}(S_0) \rceil$. However, if S_0 consists of a single vertex, then $|\text{cut}(S_0, U_0)| = \text{vol}(S_0)$, which can be significantly larger than $\lceil \phi \text{vol}(S_0) \rceil$. This observation is a key ingredient in our proof.

We begin by observing that if S_0 consists of a single vertex, then the size of $\text{cut}(S_i, U_i)$ remains $\Omega(\text{vol}(S_0))$ until $\text{vol}(S_i)$ increases to at least $\Omega(\delta \text{vol}(S_0))$. More precisely, suppose that $S_0 = \{v\}$; so, $\text{vol}(S_0) = |\text{cut}(S_0, U_0)| = d(v)$. Then, $|\text{cut}(S_i, U_i)| \geq \text{vol}(S_0) - |S_i| + 1$, because all the $\text{vol}(S_0)$ edges of the start vertex v are initially incident to uniformed vertices; and each new vertex that gets informed is incident to at most one of those edges. Also, clearly, $\text{vol}(S_i) \geq |S_i| \cdot \delta$, thus, $|S_i| \leq \text{vol}(S_i)/\delta$. Therefore, $|\text{cut}(S_i, U_i)| \geq \text{vol}(S_0) - \text{vol}(S_i)/\delta$. So,

Observation 3.5. *If $|S_0| = 1$ and $\text{vol}(S_i) \leq \delta \text{vol}(S_0)/2$ then $|\text{cut}(S_i, U_i)| \geq \text{vol}(S_0)/2$.*

We use this result in the proof of the next lemma, which is similar to Lemma 3.2(a).

Lemma 3.6. *Let $D = \min\{\Delta, \delta \text{vol}(S_0)/2\}$. If $|S_0| = 1$ then $\Pr(\text{vol}(S_{j+i}) \geq D \mid S_j) \geq 1/2$, for $i \geq 8\Delta/\text{vol}(S_0)$.*

Proof. Fix the set S_j arbitrarily. As in the proof of Lemma 3.2(a), we consider a sequence L_1, L_2, \dots of random variables, where L_i is as follows:

- If $\text{vol}(S_{j+i-1}) \leq D$, let E_i be an arbitrary subset of $\text{cut}(S_{j+i-1}, U_{j+i-1})$ of size $M = \lceil \text{vol}(S_0)/2 \rceil$, fixed before round $j+i$. (By Observation 3.5, $|\text{cut}(S_{j+i-1}, U_{j+i-1})| \geq M$.) Then L_i is the total volume of the vertices informed in round $j+i$ through edges in E_i .
- If $\text{vol}(S_{j+i-1}) > D$, then $L_i = M$.

Similarly to Claim 3.3, $\mathbf{E}[\sum_{k \leq i} L_k] = iM$ and $\mathbf{Var}(\sum_{k \leq i} L_k) \leq iM\Delta$. And, similarly to (3.1), $\Pr(\sum_{k \leq i} L_k < D) \leq iM\Delta/(iM - D)^2 < 1/2$, for $i \geq 2(\Delta + D)/M$. Since $\text{vol}(S_{j+i}) < D$ implies $\sum_{k \leq i} L_k < D$, and since $2(\Delta + D)/M \leq 8\Delta/\text{vol}(S_0)$, the lemma follows. \square

We can now derive Theorem 1.2(a) similarly to Lemma 3.1.

Proof of Theorem 1.2(a). Let $d = \text{vol}(S_0)$ be the degree of the start vertex. By Lemma 3.6, the probability that the total volume of informed vertices is smaller than $D = \min\{\Delta, \delta d/2\}$ after $c \ln n \lceil 8\Delta/d \rceil$ rounds is at most $(1 - 1/2)^{c \ln n} \leq n^{-c/2}$. The above number of rounds is $O(\phi^{-1} \ln n)$, since $d = \Omega(\Delta(\phi + \delta^{-1})) = \Omega(\phi\Delta)$. Thus, w.h.p., it takes $O(\phi^{-1} \ln n)$ rounds until the total volume of informed vertices becomes at least D .

Since $d = \Omega(\Delta(\phi + \delta^{-1})) = \Omega(\Delta/\delta)$, we have $D = \Omega(\Delta)$. Thus, by Lemma 3.1, once the total volume of informed vertices is at least D , it takes $O(\log n(\phi^{-1} + \Delta/\lceil \phi D \rceil)) = O(\phi^{-1} \ln n)$ additional rounds until all vertices get informed w.h.p. \square

The following direct corollary of Theorem 1.2(a) gives a condition for rumor spreading in $O(\phi^{-1} \log n)$ rounds for *any* start vertex.

Corollary 3.7. *If $\delta = \Omega(\Delta(\phi + \delta^{-1}))$, or, equivalently, $\delta = \Omega(\phi\Delta + \sqrt{\Delta})$ then, for any start vertex, the broadcast time of the PULL algorithm is $O(\phi^{-1} \log n)$ rounds w.h.p.*

3.2.2. Optimality of Theorem 1.2

The conditions described in Theorem 1.2, that the degree of the start vertex be $d = \Omega(\Delta(\phi + \delta^{-1}))$ or the maximum degree be $\Delta = O(1/\phi)$, are optimal in the following sense.

Theorem 3.8. *For any ϕ, δ, Δ, d with $\delta \leq d = o(\Delta(\phi + \delta^{-1}))$ and $\Delta = \omega(1/\phi)$, there exists an infinite sequence of graphs G_1, G_2, \dots such that G_n has $\Theta(n)$ vertices, conductance $\Theta(\phi)$, and maximum (minimum) degree $\Theta(\Delta)$ ($\Theta(\delta)$), and it contains a start vertex of degree $\Theta(d)$ such that $\omega(\phi^{-1} \log n)$ rounds of the PULL algorithm are required to inform all vertices w.h.p.*

Proof. First we consider the case of $d = o(\phi\Delta)$. Construct the following graph: Take a Δ -regular graph R_Δ on n vertices with edge expansion $\xi = \Theta(\Delta)$. Such a graph exists since the edge expansion of a random Δ -regular graph is $\Theta(\Delta)$ w.h.p. [2]. The conductance of R_Δ is obviously $\xi/\Delta = \Theta(1)$. Add a vertex s of degree d and a vertex v_{\min} of degree δ , choosing their neighbors arbitrarily among the vertices of R_Δ . Vertex s will be the start vertex, while v_{\min} is added just to have minimum degree δ . Next we add a component to achieve conductance ϕ : Take the complete graph on Δ vertices K_Δ . Let A be an arbitrary subset of the vertices of R_Δ of size $|A| = \lfloor \phi\Delta \rfloor$. (It is $|A| > 0$ since $1 \leq d = o(\phi\Delta)$.) Draw edges between each vertex of K_Δ and each vertex in A . It is not hard to see that the resulting graph has the desired number of vertices, maximum and minimum degrees, and conductance. Also, since $d = o(\phi\Delta)$, the probability that no neighbor of s receives the rumor from s in $k = \lfloor \phi^{-1} \ln n \cdot (2/3) \sqrt{\phi\Delta/d} \rfloor = \omega(\phi^{-1} \ln n)$ rounds is at least

$$(1 - 1/\Delta)^{kd} \geq e^{-3kd/2\Delta} \geq e^{-\ln n \sqrt{d/\phi\Delta}} = n^{-o(1)},$$

where for the first inequality we used the fact that $1 - x \geq e^{-3x/2}$, for $0 \leq x \leq 1/2$. So, with probability $n^{-o(1)}$, no vertex learns a rumor started at s in $O(\phi^{-1} \ln n)$ rounds.

Next we consider the complementary case, $d = \Omega(\phi\Delta)$. Since $d = o(\Delta(\phi + \delta^{-1}))$, we have $\phi = o(1/\delta)$ and $d = o(\Delta/\delta)$. Consider the following graph: Take the graph we constructed before and remove vertex s together with its incident edges. Take also $\lceil d/\delta \rceil$ copies of K_δ . Add a vertex s' of degree $\Theta(d)$ with neighbors the vertices of the $\lceil d/\delta \rceil$ δ -cliques, plus the elements of an arbitrary subset B of the vertices of R_Δ , with $|B| = \lceil \phi d \delta \rceil$. (It is $|B| = O(d)$ since $\phi = o(1/\delta)$ as we saw above.) It is not hard to see that the resulting graph has the desired number of vertices, maximum and minimum degrees, and conductance. Also, with probability $n^{-o(1)}$, no vertex in B learns a rumor started at s' in $O(\phi^{-1} \ln n)$ rounds: Since $d = o(\Delta/\delta)$ and $\Delta = \omega(1/\phi)$, we have $|B| \leq \phi d \delta + 1 = o(\phi\Delta) + 1 = o(\phi\Delta)$. Thus, the probability that no neighbor of s' in B receives the rumor from s' in $k = \lfloor \phi^{-1} \ln n \cdot (2/3) \sqrt{\phi\Delta/|B|} \rfloor = \omega(\phi^{-1} \ln n)$ rounds is at least $(1 - 1/\Delta)^{k|B|} \geq e^{-3k|B|/2\Delta} \geq e^{-\ln n \sqrt{|B|/\phi\Delta}} = n^{-o(1)}$. \square

3.2.3. Bounded Ratio of the Degrees of Adjacent Vertices

It was shown in [6] that if the ratio of the degrees of any two adjacent vertices is bounded by a constant, then the broadcast time of the PULL algorithm is $O((\log \phi^{-1})^2 \phi^{-1} \log n)$ rounds w.h.p., for any start vertex. By similar reasoning as in the proofs of Lemma 3.1 and Theorem 1.2(a), we can show that, in fact, the above condition yields a broadcast time of $O(\phi^{-1} \log n)$ rounds. The proof can be found in the Appendix.

Theorem 3.9. *If, for every edge $\{v, u\}$, $d(v)/d(u) = \Theta(1)$ then, the broadcast time of the PULL algorithm is $O(\phi^{-1} \log n)$ rounds w.h.p., for any start vertex.*

4. PUSH Algorithm

The analysis of the PUSH algorithm can be reduced to that of the PULL algorithm, by exploiting a symmetry between the two algorithms, described in the following result. This result is similar to Lemma 3 in [6]. Its proof can be found in the Appendix.

Lemma 4.1. *Let $\mathcal{E}_{\text{PUSH}}(v, u, t)$ denote the event that the PUSH algorithm spreads to vertex u a rumor started at vertex v in at most t rounds; and let $\mathcal{E}_{\text{PULL}}(v, u, t)$ be defined similarly. Then, $\Pr(\mathcal{E}_{\text{PUSH}}(v, u, t)) = \Pr(\mathcal{E}_{\text{PULL}}(u, v, t))$.*

Suppose that for any vertex u , the PULL algorithm distributes a rumor started at u to all vertices in at most t rounds with probability at least $1 - q$. Then, by Lemma 4.1, for any vertex v , the PUSH algorithm spreads to a given u a rumor started at v in at most t rounds with probability at least $1 - q$; and, by the union bound, if $q \leq 1/(n-1)$, the rumor started at v is spread to *all* vertices in at most t rounds with probability at least $1 - (n-1)q$. Thus, if the broadcast time of the PULL algorithm is $O(\phi^{-1} \log n)$ rounds w.h.p. for any start vertex, then the same is true for the PUSH algorithm, as well. Hence, the conditions described in Section 3 guaranteeing a broadcast time of $O(\phi^{-1} \log n)$ rounds w.h.p. for any start vertex, apply to the PUSH algorithm as well; specifically, Theorem 1.2(b), Corollary 3.7, and Theorem 3.9. Finally, Theorem 3.8 is also true for the PUSH algorithm for $d = \delta$. (For, otherwise, by the same reasoning as above, with the roles of the PUSH and the PULL algorithms switched, we would contradict Theorem 3.8.)

5. PUSH-PULL Algorithm

We prove Theorem 1.1, which gives a bound of $O(\phi^{-1} \log n)$ rounds w.h.p. on the broadcast time of the PUSH-PULL algorithm, and argue that this bound is tight.

Proof of Theorem 1.1. Fix a vertex v and let v_{\max} be a vertex of maximum degree. By Corollary 3.4, we have that: (A) The PULL algorithm distributes a rumor from v_{\max} to all other vertices (and thus to v) in $O(\phi^{-1} \log n)$ rounds w.h.p. Combining this with Lemma 4.1 yields: (B) The PUSH algorithm spreads to v_{\max} a rumor started at v in $O(\phi^{-1} \log n)$ rounds w.h.p. The theorem now follows easily: Statement (B) implies (a fortiori) that the PUSH-PULL algorithm spreads to v_{\max} a rumor started at v in $O(\phi^{-1} \log n)$ rounds w.h.p.; and, once v_{\max} is informed, Statement (A) implies that from v_{\max} the PUSH-PULL algorithm spreads the rumor to all vertices in $O(\phi^{-1} \log n)$ additional rounds w.h.p. \square

The following result was shown in [6].

Lemma 5.1. *For any $\phi \geq 1/n^{1-\epsilon}$, for a fixed $\epsilon > 0$, there exists an infinite sequence of graphs G_1, G_2, \dots such that G_n has $\Theta(n)$ vertices, conductance $\Theta(\phi)$, and diameter $\Omega(\phi^{-1} \log n)$.*

From this, it is immediate that rumor spreading requires $\Omega(\phi^{-1} \log n)$ rounds, if $\phi \geq 1/n^{1-\epsilon}$. Thus, the bound of Theorem 1.1 is asymptotically tight for $\phi \geq 1/n^{1-\epsilon}$. The next result shows this is in fact true for all $\phi = \Omega(1/n)$.

Lemma 5.2. *For any ϕ with $2/(n+2) \leq \phi \leq 1/2$, there exists an infinite sequence of graphs G_1, G_2, \dots such that G_n has n vertices and conductance $\Theta(\phi)$, and, for any start vertex, $\Omega(\phi^{-1} \log n)$ rounds of the PUSH-PULL algorithm are required to inform all vertices w.h.p.*

Proof. Consider the n -vertex graph obtained by taking two stars, one with $\lceil\phi^{-1}\rceil$ vertices and another with $n - \lceil\phi^{-1}\rceil$ vertices, and connecting their centers with an edge. It is easy to see that the resulting graph has conductance $\Theta(\phi)$. We now show that for any start vertex and any constant $c > 0$, at least $c \ln n / 3\phi$ rounds are required to inform all vertices with probability $1 - n^{-c}$. Let v and v' be the centers of the two stars, where v is the center of the star containing the start vertex. Let j be the round when v gets informed. (If v is the start vertex then $j = 0$.) The probability that v' is not informed by the end of round $j + i$, which happens if the rumor is not transmitted from v to v' via a PUSH or PULL operation in any of the rounds $j + 1, \dots, j + i$, is clearly

$$(1 - 1/\lceil\phi^{-1}\rceil)^i (1 - 1/(n - \lceil\phi^{-1}\rceil))^i \geq (1 - 1/\lceil\phi^{-1}\rceil)^{2i} \geq (1 - \phi)^{2i} \geq e^{-3i\phi},$$

where for the first inequality we used the fact that $\phi \geq 2/(n + 2)$, that for the last the fact that $1 - x \geq e^{-3x/2}$, for $0 \leq x \leq 1/2$. For $i < c \ln n / 3\phi$, it is $e^{-3i\phi} > n^{-c}$. Thus, at least $c \ln n / 3\phi$ rounds are required to inform all vertices with probability $1 - n^{-c}$. \square

References

- [1] Petra Berenbrink, Robert Elsässer, and Tom Friedetzky. Efficient randomised broadcasting in random regular networks with applications in peer-to-peer systems. In *Proc. 27th ACM Symp. on Principles of Distributed Computing (PODC)*, pages 155–164, 2008.
- [2] Béla Bollobás. The isoperimetric number of random regular graphs. *Eur. J. Comb.*, 9(3):241–244, 1988.
- [3] Stephen Boyd, Arpita Ghosh, Balaji Prabhakar, and Devavrat Shah. Gossip algorithms: design, analysis and applications. In *Proc. 24th IEEE INFOCOM*, pages 1653–1664, 2005.
- [4] Keren Censor-Hillel and Hadas Shachnai. Partial information spreading with application to distributed maximum coverage. In *Proc. 29th ACM Symp. on Principles of Distributed Computing (PODC)*, pages 161–170, 2010.
- [5] Flavio Chierichetti, Silvio Lattanzi, and Alessandro Panconesi. Rumor spreading in social networks. In *Proc. 36th Intl. Colloq. on Automata, Languages and Programming (ICALP)*, pages 375–386, 2009.
- [6] Flavio Chierichetti, Silvio Lattanzi, and Alessandro Panconesi. Almost tight bounds for rumour spreading with conductance. In *Proc. 42nd ACM Symp. on Theory of Computing (STOC)*, pages 399–408, 2010.
- [7] Flavio Chierichetti, Silvio Lattanzi, and Alessandro Panconesi. Rumour spreading and graph conductance. In *Proc. 21st ACM-SIAM Symp. on Discrete Algorithms (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 ACM Symp. on Principles of Distributed Computing (PODC)*, pages 1–12, 1987.
- [9] Benjamin Doerr, Tobias Friedrich, and Thomas Sauerwald. Quasirandom rumor spreading. In *Proc. 19th ACM-SIAM Symp. on Discrete Algorithms (SODA)*, pages 773–781, 2008.
- [10] Benjamin Doerr, Tobias Friedrich, and Thomas Sauerwald. Quasirandom rumor spreading: Expanders, push vs. pull, and robustness. In *Proc. 36th Intl. Colloq. on Automata, Languages and Programming (ICALP)*, pages 366–377, 2009.

- [11] Robert Elsässer. On the communication complexity of randomized broadcasting in random-like graphs. In *Proc. 18th ACM Symp. on Parallelism in Algorithms and Architectures (SPAA)*, pages 148–157, 2006.
- [12] Robert Elsässer, Ulf Lorenz, and Thomas Sauerwald. On randomized broadcasting in star graphs. *Discrete Appl. Math.*, 157(1):126–139, 2009.
- [13] Robert Elsässer and Thomas Sauerwald. Broadcasting vs. mixing and information dissemination on Cayley graphs. In *Proc. 24th Symp. on Theoretical Aspects of Computer Science (STACS)*, pages 163–174, 2007.
- [14] Robert Elsässer and Thomas Sauerwald. The power of memory in randomized broadcasting. In *Proc. 19th ACM-SIAM Symp. on Discrete Algorithms (SODA)*, pages 218–227, 2008.
- [15] Robert Elsässer and Thomas Sauerwald. On the runtime and robustness of randomized broadcasting. *Theoretical Computer Science*, 410(36):3414–3427, 2009.
- [16] Uriel Feige, David Peleg, Prabhakar Raghavan, and Eli Upfal. Randomized broadcast in networks. *Random Structures and Algorithms*, 1(4):447–460, 1990.
- [17] Nikolaos Fountoulakis and Konstantinos Panagiotou. Rumor spreading on random regular graphs and expanders. In *Proc. 14th Intl. Workshop on Randomization and Computation (RANDOM)*, pages 560–573, 2010.
- [18] Pierre Fraigniaud and George Giakkoupis. On the bit communication complexity of randomized rumor spreading. In *Proc. 22nd ACM Symp. on Parallel Algorithms and Architectures (SPAA)*, pages 134–143, 2010.
- [19] Alan Frieze and Geoffrey Grimmett. The shortest-path problem for graphs with random arc-lengths. *Discrete Appl. Math.*, 10:57–77, 1985.
- [20] George Giakkoupis and Philipp Woelfel. On the randomness requirements of rumor spreading. In *Proc. 22nd ACM-SIAM Symp. on Discrete Algorithms (SODA)*, 2011.
- [21] Mor Harchol-Balter, Frank Thomson Leighton, and Daniel Lewin. Resource discovery in distributed networks. In *Proc. 18th ACM Symp. on Principles of Distributed Computing (PODC)*, pages 229–237, 1999.
- [22] Richard Karp, Christian Schindelhauer, Scott Shenker, and Berthold Vöcking. Randomized rumor spreading. In *Proc. 41st IEEE Symp. on Foundations of Computer Science (FOCS)*, pages 565–574, 2000.
- [23] Damon Mosk-Aoyama and Devavrat Shah. Computing separable functions via gossip. In *Proc. 25th ACM Symp. on Principles of Distributed Computing (PODC)*, pages 113–122, 2006.
- [24] Boris Pittel. On spreading a rumor. *SIAM J. Appl. Math.*, 47(1):213–223, 1987.
- [25] Thomas Sauerwald. On mixing and edge expansion properties in randomized broadcasting. In *Proc. 18th Intl. Symp. on Algorithms and Computation (ISAAC)*, pages 196–207, 2007.
- [26] Thomas Sauerwald and Alexandre Stauffer. Rumor spreading and vertex expansion on regular graphs. In *Proc. 22nd ACM-SIAM Symp. on Discrete Algorithms (SODA)*, 2011.
- [27] Robbert van Renesse, Yaron Minsky, and Mark Hayden. A gossip-style failure detection service. In *Proc. 15th IFIP Intl. Conf. on Distributed Systems Platforms (Middleware)*, pages 55–70, 1998.

A. Appendix

A.1. Proof of Theorem 3.9

We begin by showing the following result, which is similar to Lemmata 3.2(a) and 3.6.

Lemma A.1. *Suppose that for any edge $\{v, u\}$, $d(v)/d(u) \in [1/\alpha, \alpha]$. If $\text{vol}(S_0) < \Delta$ then $\Pr(\text{vol}(S_i) \geq \min\{2 \text{vol}(S_0), \Delta\}) \geq 1/2$, for $i \geq 4\alpha/\phi$.*

Proof. We consider a sequence L_1, L_2, \dots of random variables, where L_i is defined as follows:

1. If $\text{vol}(S_{i-1}) < \Delta$, we let E_i be an arbitrary subset of $\text{cut}(S_{i-1}, U_{i-1})$ consisting of $M = \lceil \phi \text{vol}(S_0) \rceil$ edges. Then L_i is the total volume of the vertices that get informed in round i as a result of the rumor being transmitted through edges in E_i .
2. If $\text{vol}(S_{i-1}) \geq \Delta$, then $L_i = M$.

Similarly to Claim 3.3, we can show that $\mathbf{E}[\sum_{k \leq i} L_k] = iM$ and $\mathbf{Var}(\sum_{k \leq i} L_k) \leq iM(\alpha \text{vol}(S_0))$. To obtain the latter result we use the fact that if $\text{vol}(S_j) < 2 \text{vol}(S_0)$, then every vertex $v \in S_j$ has degree at most $\text{vol}(S_0)$, and thus, every neighbor of v in U_j has degree at most $\alpha \text{vol}(S_0)$. Next, similarly to (3.1), we have

$$\Pr\left(\sum_{k \leq i} L_k < \text{vol}(S_0)\right) \leq \frac{iM\alpha \text{vol}(S_0)}{(iM - \text{vol}(S_0))^2} \leq \frac{iM\alpha \text{vol}(S_0)}{(iM - \alpha \text{vol}(S_0))^2} < 1/2,$$

for $i \geq 4\alpha \text{vol}(S_0)/M$, and thus, for $i \geq 4\alpha/\phi$. Finally, since $\text{vol}(S_i) < \min\{2 \text{vol}(S_0), \Delta\}$ implies $\sum_{k \leq i} L_k \leq \text{vol}(S_i) - \text{vol}(S_0) < \text{vol}(S_0)$, the lemma follows. \square

Theorem 3.9 follows now easily: By Lemma A.1 and Chernoff bounds, it takes w.h.p. $O(\phi^{-1} \log n)$ rounds until the total volume of informed vertices becomes at least Δ . And, by Lemma 3.1, it takes $O(\phi^{-1} \ln n)$ additional rounds until all vertices get informed w.h.p.

A.2. Proof of Lemma 4.1

Consider the space of all possible executions of the first t rounds of the PUSH or the PULL algorithm; each execution specifies the vertex that v calls in round i , for every vertex v and $i = 1, \dots, t$. We call these executions t -executions. The probability that a given t -execution occurs is the same for all t -executions and it is equal to $\prod_{v \in V} (d(v))^{-t}$. For $\mathcal{A} \in \{\text{PUSH}, \text{PULL}\}$, let $\Omega_{\mathcal{A}}(v, u, t)$ be the subset of t -executions that constitute the event $\mathcal{E}_{\mathcal{A}}(v, u, t)$. Then, $\Pr(\mathcal{E}_{\mathcal{A}}(v, u, t)) = |\Omega_{\mathcal{A}}(v, u, t)| \cdot \prod_{v \in V} (d(v))^{-t}$. Thus, to complete the proof it suffices to show that $|\Omega_{\text{PUSH}}(v, u, t)| = |\Omega_{\text{PULL}}(u, v, t)|$.

For a t -execution ω let ω' be the “inverse” t -execution consisting of the same sequence of rounds as ω but executed in reverse order; i.e., if vertex v calls vertex u in round i in ω , then v calls u in round $t - i$ in ω' . We now prove that $\omega \in \Omega_{\text{PUSH}}(v, u, t)$ if and only if $\omega' \in \Omega_{\text{PULL}}(u, v, t)$. From this, and the fact that distinct t -executions have distinct inverse executions, it follows that $|\Omega_{\text{PUSH}}(v, u, t)| = |\Omega_{\text{PULL}}(u, v, t)|$.

A *PUSH-path* for a t -execution is a sequence of vertices u_0, u_1, \dots, u_t such that, for any two consecutive vertices u_{i-1}, u_i , either (i) $u_{i-1} = u_i$, or (ii) u_{i-1} calls u_i in round i . A *PULL-path* is defined similarly except that condition (ii) now states that u_i calls u_{i-1} in round i . Clearly, for a given t -execution, the PUSH (PULL) algorithm spreads a rumor from v to u if and only if there exists a PUSH-path (PULL-path) from v to u . Also, $u_0 = v, u_1, \dots, u_t = u$ is a PUSH-path for ω if and only if u_t, u_{t-1}, \dots, u_0 is a PULL-path for ω' . Combining these two facts yields $\omega \in \Omega_{\text{PUSH}}(v, u, t)$ if and only if $\omega' \in \Omega_{\text{PULL}}(u, v, t)$.