

Randomized Rumor Spreading in Dynamic Graphs

George Giakkoupis¹, Thomas Sauerwald², and Alexandre Stauffer³

¹ INRIA Rennes, France

² Computer Laboratory, University of Cambridge, UK

³ Department of Mathematical Sciences, University of Bath, UK

Abstract. We consider the well-studied rumor spreading model in which nodes contact a random neighbor in each round in order to push or pull the rumor. Unlike most previous works which focus on static topologies, we look at a dynamic graph model where an adversary is allowed to rewire the connections between vertices before each round, giving rise to a sequence of graphs, G_1, G_2, \dots . Our first result is a bound on the rumor spreading time in terms of the conductance of those graphs. We show that if the degree of each node does not change much during the protocol (that is, by at most a constant factor), then the spread completes within t rounds for some t such that the sum of conductances of the graphs G_1 up to G_t is $O(\log n)$. This result holds even against an *adaptive* adversary whose decisions in a round may depend on the set of informed vertices before the round, and implies the known tight bound with conductance for static graphs. Next we show that for the alternative expansion measure of vertex expansion, the situation is different. An adaptive adversary can delay the spread of rumor significantly even if graphs are regular and have high expansion, unlike in the static graph case where high expansion is known to guarantee fast rumor spreading. However, if the adversary is *oblivious*, i.e., the graph sequence is decided before the protocol begins, then we show that a bound close to the one for the static case holds for any sequence of regular graphs.

1 Introduction

Randomized rumor spreading is a popular epidemic protocol for disseminating information in large distributed networks. The protocol proceeds in a sequence of synchronous rounds. Initially, in round 0, an arbitrary node has a piece of information, the *rumor*. This rumor is then spread iteratively to the other nodes: In each round, every informed node (i.e., every node that learned the rumor in a previous round) chooses a random neighbor to which it sends the rumor. This is the so-called PUSH protocol. The PULL protocol is symmetric: In each round, every *uninformed* node chooses a random neighbor, and if this neighbor knows the rumor it transmits it to the uninformed node. Finally, the PUSH-PULL protocol is the combination of both strategies: In each round, every node chooses a random neighbor to send the rumor to, if the node is informed, or to request the rumor from, otherwise.

Several aspects of rumor spreading have been analyzed, including its running time (i.e., the number of rounds until all nodes get informed), the corresponding number of messages, and the amount of randomness needed. The running time is arguably the most fundamental and well-studied of those aspects. In particular, it has been shown that just a logarithmic number of rounds suffice to spread a rumor with high probability (*w.h.p.*) on several topologies, from basic communication networks, such as complete graphs, hypercubes and random graphs [27,11,19], to more complex structures, such as preferential attachment graphs or power-law random graphs modeling social networks [9,12]. Recently, a number of studies have extended this line of work by establishing bounds on the running time of rumor spreading in terms of expansion parameters of the underlying graph, namely conductance [24,4,14,3] and vertex expansion [29,16,15]. This connection between rumor spreading and graph theory is also relevant for understanding social and other real networks, as studies have indicated that such networks have good expansion properties [23,5].

A limitation of the above results is that they require the graph to be fixed throughout the execution of the rumor spreading protocol, whereas many of the prevalent topologies, such as peer-to-peer or wireless networks, are inherently dynamic. In particular, the structure of these networks may change more quickly than a rumor spreads.

In this paper we analyze the running time of randomized rumor spreading in a dynamic setting, given by a sequence of graphs G_1, G_2, \dots with the same vertex set of size n , but possibly distinct edge sets. In this setting, at each round t , a vertex contacts a random neighbor *in graph* G_t in order to push or pull the rumor. We assume that the edge set for each round is determined by an adversary. The adversary can be either *adaptive*, i.e., it decides the edge set for each round at the beginning of the round, knowing the set of informed vertices at the time, or *oblivious*, i.e., it fixes the complete sequence of graphs before rumor spreading starts, knowing just the source of the rumor.

Our first result is an upper bound on the running time of PUSH-PULL in terms of the conductances of graphs G_1, G_2, \dots . Suppose that during the execution of the protocol, we have for each node that the ratio of its maximum over minimum degree is bounded by some $\rho \geq 1$. We show that rumor spreading then completes within t rounds for some t such that the sum of conductances of the graphs G_1 up to G_t is $O(\rho \log n)$. Moreover, this bound holds even if the adversary is adaptive.

Theorem 1. *Let G_1, G_2, \dots be a sequence of graphs determined by an adaptive adversary such that for each vertex u , its degree on each graph G_t is at least $\delta_u > 0$ and at most Δ_u . The degree bounds δ_u, Δ_u may be different for each u , and are fixed (in advance) by the adversary. Let $\rho = \max_u (\Delta_u / \delta_u)$. Also for each $t \geq 1$, let ϕ_t be the conductance of graph G_t .*

For any constant $\beta > 0$, there exists a constant $b > 0$ so that the following bound holds for PUSH-PULL. Let τ be the first round for which $\sum_{t=1}^{\tau} \phi_t \geq b\rho \log n$; $\tau = \infty$ if not such round exists. Then with probability $1 - n^{-\beta}$, either $\tau = \infty$ or all nodes have been informed by the end of round τ .

If all graphs G_1, G_2, \dots are the same and have conductance ϕ , then $\rho = 1$ and the bound of Theorem 1 implies the optimal bound for static graphs established in [14], stating that $O(\log(n)/\phi)$ rounds suffice w.h.p. to spread a rumor on any graph with conductance ϕ .

The dependence of the bound in Theorem 1 on ρ is not an artifact of our analysis. For example, even an oblivious adversary can construct a sequence of graphs in which every graph has constant conductance (but the degrees of nodes change widely and thus ρ is large) so that PUSH-PULL needs a linear number of rounds to inform all nodes (see Proposition 1, in Sect. 5).

We point out that Theorem 1 is shown for a more general setting, where multiple edges and self-loops are allowed. Moreover, it holds even if we define ϕ_t to be the conductance in G_t of the *set of informed vertices* before round t . This can be much larger than the conductance of G_t , which is the minimum conductance of any set of vertices in G_t .

For static graphs, conductance and vertex expansion yield very similar types of bounds on the running time of rumor spreading. For dynamic graphs, however, the situation is different. In particular, we reveal a separation between the adaptive and oblivious adversary models for the case of vertex expansion, which is not observed for conductance: An adaptive adversary can construct a sequence of regular graphs with constant vertex expansion so that PUSH-PULL takes a polynomial number of rounds (see Proposition 2)—by Theorem 1, this is not possible if conductance is considered in place of vertex expansion. For an oblivious adversary, on the other hand, we show that a bound similar to the one in Theorem 1 holds with vertex expansion, for any sequence of regular graphs.

Theorem 2. *Let G_1, G_2, \dots be a sequence of d -regular graphs determined by an oblivious adversary. For each $t \geq 1$, let α_t be the vertex expansion of G_t . Then, for any constant $\beta > 0$, there exists a constant $c > 0$ so that, if there exists a round t with $\sum_{s=1}^t \alpha_s \geq c \cdot \log^4 n \log^2 d$, then PUSH-PULL informs all nodes within t rounds with probability at least $1 - n^{-\beta}$.*

If all graphs G_1, G_2, \dots are the same and have vertex expansion α , then the bound of Theorem 2 matches within a polylogarithmic factor the optimal bound for static graphs from [15], which states that $O(\log(n) \log(d)/\alpha)$ rounds suffice w.h.p. to spread a rumor on any graph (even non-regular) with vertex expansion α and maximum degree d . Whether Theorem 2 extends also to non-regular graph sequences is an interesting open problem.

Our proofs are non-trivial extensions of previous analyses for static graphs, in particular from [14] and [29]. The dynamic setting, and also the adaptivity of the adversary for the case of conductance, add new challenges to the problem. The proof of Theorem 1 is based on a new martingale argument which exposes the outcome of each round gradually, one vertex at a time. The proof of Theorem 2 has to overcome the problem that the standard symmetry argument relating push and pull no longer holds, and this breaks key arguments used in existing proofs of bounds with vertex expansion. Moreover, tighter proofs from [16,15] employ potential functions based on the boundary of the informed nodes, which may largely fluctuate in the dynamic setting.

Related Work. There have been several studies on information spreading processes in dynamic graphs. Perhaps the closest one to our work is a recent work by Clementi et al. [6] about PUSH in a random edge-Markovian model, yielding a dynamic variant of the Erdős-Rényi random graph. Motivated by the increasing importance of wireless networks, works [20,26,22,25] analyzed the dynamics of information dissemination among moving objects in d -dimensional grids. There are also several analyses on the *flooding* process—a variant of rumor spreading where *every* neighbor of an informed node becomes informed in a round [8,7,2]. We note that in all these works, the graph dynamics are governed by a random process, whereas in our model the dynamics are controlled by an adversary.

Avin et al. [1] analyzed the cover time of random walks on dynamic graphs specified by an oblivious adversary. They constructed graphs in which a simple random walk has an exponential cover time, but also proved that a lazy random walk has a polynomial cover time for any sequence of connected graphs.

Kuhn et al. [21] introduced the so-called *k-token dissemination problem* in a synchronous setting with a dynamically changing network. They considered a worst-case scenario, where the communication links are chosen by an adversary, and nodes do not know who their neighbors are for the current round before they send their messages. In contrast to our model, the connectivity (expansion) assumptions are weaker and correspondingly the time complexity bounds are much larger, i.e., at least polynomial in n [21,17,10,18,28]. Georgiou et al. [13] considered the complexity of asynchronous gossip in a fault-prone distributed setting. While their model is quite different from ours, they also exhibited a separation between an adaptive and oblivious adversary.

2 Model

We consider the standard PUSH, PULL, and PUSH-PULL rumor spreading protocols. We will denote by I_t the set of informed vertices after the first t rounds of the protocol, and by U_t the set $V \setminus I_t$ of uninformed vertices. In particular, I_0 is the singleton set containing just the source, which is an arbitrarily chosen node.

The dynamic graph model we consider is an infinite sequence of graphs G_1, G_2, \dots on the same set of n vertices, but possibly with different edge sets. In each round t of rumor spreading, if a vertex must choose a neighbor to push the rumor to or pull the rumor from, then it chooses a random one among its neighbors *in graph G_t* (if it has neighbors in G_t).

For each $t \geq 1$, $G_t = (V, E_t)$, where $V = \{1, 2, \dots, n\}$, and the edge set E_t is determined by an *adversary*. We distinguish between two adversarial models, *adaptive* and *oblivious*. An *adaptive adversary* decides the edge set E_t knowing the outcome of all rounds before round t ; precisely, E_t is a function of I_0, \dots, I_{t-1} . An *oblivious adversary*, on the other hand, has to specify the entire graph sequence in advance; precisely, E_t is just a function of I_0 . An adversary can be either deterministic or randomized, where in the latter case the sequence of E_t is also a function of a random bit string.

We recall now the definitions of two standard graph expansion parameters we use. For a graph $G = (V, E)$, the *conductance* of a non-empty vertex set $S \subset V$ and the conductance of graph G are defined respectively as

$$\phi(S) = \frac{|E(S, V \setminus S)|}{\min\{\text{vol}(S), \text{vol}(V \setminus S)\}} \quad \text{and} \quad \phi(G) = \min_{S \subset V, S \neq \emptyset} \phi(S),$$

where $E(S, V \setminus S)$ is the set of edges with one endpoint in each of the sets S and $V \setminus S$; and $\text{vol}(S) = \sum_{u \in S} \deg(u)$ is the *volume* of S , with $\deg(u)$ denoting the degree of u . The *vertex expansion* of S and G are respectively

$$\alpha(S) = \frac{|\partial S|}{\min\{|S|, |V \setminus S|\}} \quad \text{and} \quad \alpha(G) = \min_{S \subset V, S \neq \emptyset} \alpha(S),$$

where $\partial S := N(S) \setminus S$ is the set of vertices outside S that are adjacent to some vertex in S . For any graph G , both $\phi(G)$ and $\alpha(G)$ are between 0 and 1, with high values indicating that the graph is well connected. If G is disconnected then $\phi(G) = \alpha(G) = 0$.

In the following, when we write $\text{vol}(I_t)$, ∂I_t , $\phi(I_t)$, etc. (and similarly for U_t), we will assume that these quantities refer to graph G_{t+1} (and not to G_t), unless mentioned otherwise. This is convenient since I_t is the set of informed vertices at the beginning of round $t + 1$.

Theorem 1 holds even if the graph sequence consists of *multigraphs*, with parallel edges and self-loops (as long as the degree of each vertex is at most polynomially large in n). To compute the conductance of a multigraph, parallel edges are counted with respect to their multiplicity and every self-loop counts as a single edge. Further, when a vertex must pick a random neighbor, this is done proportional to the multiplicity of a (parallel) edge or self-loop. If a self-loop is chosen then no communication takes place.

3 Proof of the Bound with Conductance (Theorem 1)

We observe that it suffices to consider just deterministic adversaries, since in case of a randomized adversary, we can just expose all its randomness (i.e., its random bit string) before the protocol starts, and then proceed deterministically.

Not all vertices are guaranteed to get informed eventually, as the adversary may permanently disconnect the network. However, it is not difficult to show that one can always modify a (deterministic) adversary, in such a way that this does not happen and it suffices to consider the modified adversary in the analysis. Thus, we will assume that the expected number of rounds until all nodes are informed is finite, i.e.,

$$\mathbf{E}[\min\{t: I_t = V\}] < \infty. \tag{1}$$

For each round $t \geq 1$, let $\Phi_t = \phi_1 + \dots + \phi_t$ be the sum of the conductances of the graphs in the first t rounds. We must show that for $\tau = \inf\{t: \Phi_t \geq b\rho \log n\}$, we have w.h.p. that all nodes have been informed within τ rounds.

For each set $S \subseteq V$ of vertices, we define the *min-volume* of S as $\text{vol}^*(S) = \sum_{i \in S} \delta_i$. Thus $\text{vol}^*(S)$ is a lower bound for the volume of S on any of the graphs G_t . Our proof is based on an analysis of the growth of $\text{vol}^*(I_t)$.

We will show the following lemma stating (a) if the min-volume of informed vertices is smaller than a constant fraction of the total min-volume, then the sum of conductances ϕ_t until the min-volume doubles is bounded in expectation by $O(\rho)$; and (b) if the min-volume of informed vertices is larger, then the sum of ϕ_t until the min-volume of uninformed vertices halves is bounded by $O(\rho)$.

Lemma 1. *There is a fixed constant $c > 0$ such that for any round $t \geq 1$,*

- (a) *If $\text{vol}^*(I_t) \leq \text{vol}^*(V)/3$ and $\tau_t = \min\{k: \text{vol}^*(I_k) \geq 2 \text{vol}^*(I_t)\}$, then $\mathbf{E}[\Phi_{\tau_t} - \Phi_t \mid I_t] \leq c\rho$.*
- (b) *If $\text{vol}^*(I_t) > \text{vol}^*(V)/3$ and $\tau_t = \min\{k: \text{vol}^*(U_k) \leq \text{vol}^*(U_t)/2\}$, then $\mathbf{E}[\Phi_{\tau_t} - \Phi_t \mid I_t] \leq c\rho$.*

From this result, the bound of Theorem 1 follows easily: From Lemma 1(a) and Markov's Inequality it follows that when $\text{vol}^*(I_t) \leq \text{vol}^*(V)/3$, we have $\Pr[\Phi_{\tau_t} - \Phi_t \leq 2c\rho \mid I_t] \geq 1/2$, i.e., with probability at least $1/2$, $\text{vol}^*(I_t)$ doubles after Φ_t has increased by at most $2c\rho$. It follows then by Chernoff bounds that $1/3$ of the total min-volume gets informed with probability at least $1 - n^{-\beta}/2$, for any fixed β , after a number t_1 of rounds such that $\Phi_{t_1} = 2c\rho \cdot O(\log(\text{vol}^*(V)/3)) = O(\rho \log n)$, as $\log(\text{vol}^*(V)) = O(\log n)$. A similar argument using Lemma 1(b) shows that if $1/3$ of the total min-volume has been informed by some round t , then an additional t_2 rounds, such that $\Phi_{t+t_2} - \Phi_t = O(\rho \log n)$, suffice to inform the remaining vertices with probability at least $1 - n^{-\beta}/2$. From these two results and the union bound, Theorem 1 follows.

3.1 Proof of Lemma 1

Recall that during the spread, the degree of each vertex $i \in V$ is lower bounded by $\delta_i > 0$ and upper bounded by Δ_i , and $\Delta_i/\delta_i \leq \rho$. Let $\delta = \max_i \delta_i$.

The proof distinguishes three cases, depending on the min-volume of informed vertices initially: (i) $\text{vol}^*(I_t) < \delta$, (ii) $\delta \leq \text{vol}^*(I_t) \leq \text{vol}^*(V)/3$, and (iii) $\text{vol}^*(I_t) > \text{vol}^*(V)/3$. For cases (ii) and (iii) it suffices that we consider only pull operations, while for case (i) we must consider both push and pull. Due to space limitations we only give the proof for case (ii). In this case, $\tau_t = \min\{k \geq t: \text{vol}^*(I_k) \geq 2 \text{vol}^*(I_t)\}$.

Claim 1. *If $\delta \leq \text{vol}^*(I_t) \leq \text{vol}^*(V)/3$ then $\mathbf{E}[\Phi_{\tau_t} - \Phi_t \mid I_t] < 2\rho + 1$.*

Proof. We use a martingale argument that relates the min-volume of vertices informed by pull transmissions to the number of edges between informed and uninformed vertices. In this argument, the outcome of each round is exposed gradually, one vertex at a time.⁴

⁴ The reason we expose one vertex at a time (rather than all at once), is that by stopping this process right after the min-volume of informed vertices has doubled,

Assume I_t is fixed. We divide each round $k > t$ into $|\partial I_{k-1}|$ steps (∂I_{k-1} is the set of uninformed vertices at the beginning of round k that have some informed neighbor). Each of those $|\partial I_{k-1}|$ steps reveals the push and pull transmissions of the rumor in round k to a single vertex $i \in \partial I_{k-1}$. The order in which vertices $i \in \partial I_{k-1}$ are considered can be arbitrary. We look at the sequence of all those steps, from round $t + 1$ until round τ_t . For each step $s = 1, 2, \dots$, let i_s be the vertex considered in step s , let k_s be the round in which step s takes place, d_s be the degree of i_s during round k_s (i.e., i_s 's degree in G_{k_s}), and γ_s be the number of informed neighbors of i_s in G_{k_s} at the beginning of round k_s .

Below we first show that the sum of all γ_s until the step when the min-volume of informed vertices has doubled is bounded in expectation by $\rho(\text{vol}^*(I_t) + \delta)$. Then we bound the corresponding increase in Φ_t in terms of the sum of γ_s , and combine the two results to obtain the claim.

Let X_s be the indicator variable that is 1 if i_s pulls the rumor in step s , and $X_s = 0$ otherwise. Further, let Z_s be the indicator variable that is 1 if i_s gets informed in step s , and 0 otherwise. Note that $Z_s \geq X_s$, since i_s may get informed by a push transmission. Note also that the sequence Z_1, \dots, Z_{s-1} completely determines the evolution of the set of informed vertices in the first $s - 1$ steps, and thus determines i_s , k_s , d_s , and γ_s . In order for i_s to pull the rumor in step s , it must choose one of its γ_s informed neighbors, among its d_s neighbors in total. It follows that

$$\mathbf{E}[X_s \mid Z_1 \dots Z_{s-1}] = \gamma_s/d_s. \quad (2)$$

For each $s \geq 0$, we define $Y_s = \sum_{1 \leq j \leq s} (X_j d_j - \gamma_j)$. The sequence Y_0, Y_1, \dots is a martingale with respect to Z_1, Z_2, \dots , because

$$\mathbf{E}[Y_s \mid Z_1 \dots Z_{s-1}] = Y_{s-1} + \mathbf{E}[X_s \mid Z_1 \dots Z_{s-1}] \cdot d_s - \gamma_s \stackrel{(2)}{=} Y_{s-1}.$$

Let T be the number of steps after round t until the min-volume of informed vertices doubles, i.e., $T = \min\{s : \sum_{1 \leq j \leq s} Z_j \delta_{i_j} \geq \text{vol}^*(I_t)\}$. (Observe, $k_T = \tau_t$.) Since T is a stopping time for Z_1, Z_2, \dots , and we have that $\mathbf{E}[T] < \infty$ ⁵ and the differences $Y_s - Y_{s-1}$ are bounded, it follows from the Optional Stopping Theorem that $\mathbf{E}[Y_T] = \mathbf{E}[Y_0] = 0$. Substituting the definition of Y_T and rearranging gives

$$\mathbf{E} \left[\sum_{1 \leq j \leq T} \gamma_j \right] = \mathbf{E} \left[\sum_{1 \leq j \leq T} X_j d_j \right]. \quad (3)$$

We have $\sum_{1 \leq j \leq T} X_j d_j \leq \rho \sum_{1 \leq j \leq T} Z_j \delta_{i_j}$, because $X_j \leq Z_j$ and $d_j \leq \Delta_{i_j} \leq \rho \delta_{i_j}$; and from T 's definition, $\sum_{1 \leq j \leq T} Z_j \delta_{i_j} < \text{vol}^*(I_t) + \delta$. It follows

$$\mathbf{E} \left[\sum_{1 \leq j \leq T} \gamma_j \right] < \rho(\text{vol}^*(I_t) + \delta). \quad (4)$$

we have the guarantee that the min-volume at that time is by at most $\delta - 1 < \text{vol}^*(I_t)$ larger than $2 \text{vol}^*(I_t)$. (This is used in the line right above Eq.(4), in order to obtain (4).) On the other hand, the min-volume right after the round during which the min-volume doubles may be much larger than $2 \text{vol}^*(I_t)$.

⁵ This is immediate from Eq.(1).

Next we bound $\Phi_{\tau_t} - \Phi_t$ in terms of $\sum_{1 \leq j \leq T} \gamma_j$, and apply the above inequality to bound $\mathbf{E}[\Phi_{\tau_t} - \Phi_t]$. For each round k with $t < k \leq \tau_t$, the conductance ϕ_k of G_k is bounded by the conductance of I_{k-1} in G_k , and thus

$$\phi_k \leq \frac{\sum_{s:k_s=k} \gamma_s}{\min\{\text{vol}(I_{k-1}), \text{vol}(U_{k-1})\}} \leq \frac{\sum_{s:k_s=k} \gamma_s}{\text{vol}^*(I_t)},$$

where the second inequality holds because $\text{vol}(I_{k-1}) \geq \text{vol}^*(I_{k-1}) \geq \text{vol}^*(I_t)$ and

$$\begin{aligned} \text{vol}(U_{k-1}) &\geq \text{vol}^*(U_{k-1}) \geq \text{vol}^*(U_{\tau_t-1}) = \text{vol}^*(V) - \text{vol}^*(I_{\tau_t-1}) \\ &> \text{vol}^*(V) - 2 \text{vol}^*(I_t) \geq \text{vol}^*(I_t), \end{aligned}$$

as $\text{vol}^*(I_t) \leq \text{vol}^*(V)/3$. From the above bound on ϕ_k applied for $t < k < \tau_t$, we obtain

$$\Phi_{\tau_t} - \Phi_t = (\phi_{t+1} + \dots + \phi_{\tau_t-1}) + \phi_{\tau_t} \leq \frac{\sum_{s:k_s < \tau_t} \gamma_s}{\text{vol}^*(I_t)} + 1.$$

From T 's definition, $k_T = \tau_t$, and thus the sum above is $\sum_{s:k_s < \tau_t} \gamma_s < \sum_{s \leq T} \gamma_s$. Thus,

$$\mathbf{E}[\Phi_{\tau_t} - \Phi_t] < \frac{\mathbf{E} \left[\sum_{s \leq T} \gamma_s \right]}{\text{vol}^*(I_t)} + 1 \stackrel{(4)}{<} \frac{\rho(\text{vol}^*(I_t) + \delta)}{\text{vol}^*(I_t)} + 1 \leq 2\rho + 1, \quad (5)$$

as $\delta \leq \text{vol}^*(I_t)$. This completes the proof of Claim 1. \square

The proof for the case of $\text{vol}^*(I_t) > \text{vol}^*(V)/3$ is similar, but considers the set of uninformed vertices instead of the set of informed ones.

Claim 2. *If $\text{vol}^*(I_t) > \text{vol}^*(V)/3$ then $\mathbf{E}[\Phi_{\tau_t} - \Phi_t \mid I_t] < 2\rho + 1$.*

The analysis in the proof of Claim 1 does not carry over to the case of $\text{vol}^*(I_t) < \delta$: The final inequality in (5) does not hold, as the ratio $\delta/\text{vol}^*(I_t)$ may be very large. In fact any analysis that relies only on pull transmissions is bound to fail, for otherwise Theorem 1 would hold even if only PULL were used, which is easily seen to be wrong because of the star graph counter-example. To show the next claim, we extend the approach of Claim 1 by taking into account also push transmissions.

Claim 3. *If $\text{vol}^*(I_t) < \min\{\delta, \text{vol}^*(V)/3\}$ then $\mathbf{E}[\Phi_{\tau_t} - \Phi_t \mid I_t] < 3\rho + 1$.*

4 Proof of the Bound with Vertex Expansion (Theorem 2)

Theorem 2 can be deduced easily from the following result, in which we assume a uniform lower bound on the vertex expansion. (If $d = O(\log^3 n)$ then Theorem 2 follows directly from Theorem 1 since $\alpha \leq d \cdot \Phi$.)

Theorem 3. *Let t be any integer, and G_1, G_2, \dots, G_t be a sequence of d -regular graphs with $d = \Omega(\log^3 n)$ determined by an oblivious adversary, so that for each $1 \leq s \leq t$, G_s has vertex expansion at least $\alpha > 0$. Then, for any constant $\beta > 0$, there exists a constant $\tilde{c} > 0$ such that, if t satisfies $t \cdot \alpha \geq \tilde{c} \cdot \log^3 n \log^2 d$, then PUSH-PULL informs all nodes within t rounds with probability at least $1 - n^{-\beta}$.*

4.1 Proof of Theorem 3

The analysis of Theorem 3 is divided into three phases, according to the number of informed nodes. Before we analyze the different phases, we provide some tools we will use in the analysis. The proofs of all statements in this section are omitted due to space limitations.

The next lemma establishes a (nearly) exponential growth of the number of informed nodes until that number reaches $d/\log n$.

Lemma 2. *Let $u \in V$ be arbitrary with $I_0 = \{u\}$ and let $\kappa := \log_3(d/(2 \log n)) + 440$. Then the following statements hold.*

1. $\Pr \left[I_\kappa \geq \frac{d}{6 \log n} \right] \geq 1/2$.
2. $\Pr \left[I_\kappa \leq 3^{220} \cdot 10^6 \cdot d \right] \geq 1 - n^{-4}$, and more generally, for any round $1 \leq s \leq \kappa$, $\Pr \left[I_s \leq 2 \cdot 10^6 \cdot \log n \cdot 3^s \right] \geq 1 - n^{-4}$.

The first statement of Lemma 2 motivates the following definition.

Definition 1 (Friend). *Let $U \subseteq V$ be any subset. Then a node u is a friend of U in round t if, for κ as in Lemma 2, $\Pr \left[|I_{t+\kappa} \cap U| \geq \frac{d}{12 \log n} \mid I_t = \{u\} \right] \geq 1/4$.*

A similar notion of a friend was defined in [29, Definition 3.1] for a static graph; our definition depends on the sequence of graphs $G_{t+1}, G_{t+2}, \dots, G_{t+\kappa}$, and so in particular, on the choice of t . Applying the first statement of Lemma 2, we have that for every subset $U \subseteq V$, every node $u \in V$, and every round t , node u is a friend of either U or $V \setminus U$ in round t .

Next we consider the situation where I_t is of any size, and half of the nodes in ∂I_t are friends of I_t . We shall prove that an almost constant fraction of nodes in I_t gets informed after κ rounds. It should be noted that this is relatively straightforward in the case of static graphs, as it follows from a standard symmetry argument relating PUSH and PULL (cf. [29, Lemma 3.1]). Here, however, the analysis is considerably more involved, as we are dealing with dynamic graphs.

Lemma 3 (Key Lemma). *Consider a round t with a fixed set of informed nodes I_t , where $1 \leq |I_t| \leq n/2$. Let $S \subseteq V \setminus I_t$ be a set of vertices which are friends of I_t in round $t + 1$. Then, there is a constant $0 < C < 1$ so that, with probability at least $1/16$, at least $C \cdot \frac{|S|}{\log^2 n \log d}$ nodes get informed after κ additional rounds, for κ defined as in Lemma 2.*

Next we analyze the growth of informed nodes in 3 phases: $|I_t| \in [1, \frac{d}{6 \log n}]$, $|I_t| \in [\frac{d}{6 \log n}, \frac{d}{\alpha}]$, and $|I_t| \in [\frac{d}{\alpha}, \frac{n}{2}]$. In the following, κ is defined as in Lemma 2.

Lemma 4 (Phase 1). *Assume that $|I_0| = 1$. Then after $t_1 := \kappa$ rounds, we have $|I_{t_1}| \geq \frac{d}{6 \log n}$ with probability at least $1/2$.*

The above result for the first phase follows immediately from Lemma 2.

Lemma 5 (Phase 2). *Let t_1 be the first round for which $|I_{t_1}| \geq \frac{d}{6 \log n}$. Then, for any constant $\beta > 0$, there exists a constant $c > 0$ so that, for $t_2 := t_1 + \frac{c \log^3 n \log d}{\alpha} \cdot \kappa$, we have $|I_{t_2}| \geq \frac{d}{\alpha}$ with probability at least $1 - n^{-\beta}$.*

The analysis of the second phase is more involved. We only consider every $(\kappa + 1)$ -th round and distinguish between two cases. If half of the nodes in ∂I_t are friends of I_t , then Lemma 3 implies that a large fraction of these nodes becomes informed after $\kappa + 1$ rounds. If half of the nodes in ∂I_t are friends of $V \setminus I_t$, then, if such a node in ∂I_t pulls the rumor in round $t + 1$, then after κ additional rounds $d/(12 \log n)$ nodes get informed in $V \setminus I_t$. Expressing the progress of both cases via a submartingale and applying the Optional Stopping Theorem completes the proof of Lemma 5.

Lemma 6 (Phase 3). *Let t_2 be the first round for which $|I_{t_2}| \geq \frac{d}{\alpha}$. Then, for any constant $\beta > 0$, there exists a constant $c > 0$ so that for $t_3 := t_2 + c \frac{\log^3 n \log d}{\alpha} \cdot \kappa$, we have $|I_{t_3}| \geq \frac{n}{2}$ with probability at least $1 - n^{-\beta}$.*

The analysis of the third phase is somewhat similar to the analysis of the second phase. However, the case where half of the nodes in ∂I_t are friends of $V \setminus I_t$ requires a more careful analysis, since we have to analyze the propagation of the rumor within $V \setminus I_t$ from *several* nodes in ∂I_t in parallel.

From Lemmas 4–6 it follows that, with probability at least $1 - n^{-\beta}$, Phase 1 is completed after $O(\kappa \log n) = O(\log n \log d)$ rounds, and Phases 2 and 3 are completed after $O(\frac{\log^3 n \log^2 d}{\alpha})$ rounds, thus proving the bound of Theorem 3.

5 Counter-Examples

High conductance is not sufficient to guarantee fast rumor spreading in our dynamic graph model. Even an oblivious adversary can construct a sequence of high-conductance graphs (in which the degrees of nodes change widely), so that PUSH-PULL needs a linear number of rounds to inform all nodes.

Proposition 1. *An oblivious adversary can construct a sequence of graphs G_1, G_2, \dots , each of which has conductance 1, so that PUSH-PULL needs $n - 1$ rounds to complete.*

Proof. Instead of analyzing the PUSH-PULL protocol we analyze a different process in which at each round all neighbors of the set of informed nodes become informed (i.e., $I_{t+1} = I_t \cup \partial I_t$). The strategy of the adversary is as follows. In every round t , G_t is a star graph with n vertices. The source of the rumor is a vertex of degree 1, and in each round except for the last one, the informed nodes are vertices of degree 1. At first it may seem that the adversary has to be adaptive to employ this strategy. But, given that the source of the rumor is known, then I_0 is a deterministic set. Given G_1 and I_0 , then I_1 is again a deterministic set, and so on. Therefore, the construction above for the sequence G_1, G_2, \dots can be done in an oblivious manner. With this, our process satisfies $|I_t| = t + 1$, for every $0 \leq t \leq n - 1$. Since the runtime of this process is at most the runtime of the PUSH-PULL protocol, the claim of the proposition follows. \square

An adaptive adversary can significantly delay the spread of the rumor, even if all graphs in the sequence are regular and have high vertex expansion. Thus Theorem 2 does not hold when the adversary is adaptive rather than oblivious.

Proposition 2. *An adaptive adversary can construct a sequence of regular isomorphic graphs G_1, G_2, \dots , each of which has constant vertex expansion, so that PUSH-PULL needs $\Omega(\sqrt{n})$ rounds with probability at least $1/2$.*

Proof. Let G be the Cartesian product of a clique of size \sqrt{n} with a 3-regular expander graph of size \sqrt{n} . This graph has $\sqrt{n} \cdot \sqrt{n}$ vertices and is regular, with all vertices having degree $\sqrt{n} - 1 + 3 = \sqrt{n} + 2$. Observe that G can be seen as a collection of \sqrt{n} cliques of size \sqrt{n} , with every vertex in each clique connected to 3 vertices in other cliques. By [29, Lemma 4.2], G has constant vertex expansion. Every graph in the sequence G_1, G_2, \dots will be isomorphic to G .

In each round $t \geq 1$, the adversary permutes the vertices in a way so that there is at most one clique that is not fully informed or fully uninformed (i.e., whose number of informed nodes is in the interval $[1, \sqrt{n} - 1]$). All other cliques contain either \sqrt{n} informed nodes or none. Consider now a clique in which all \sqrt{n} nodes are informed. The expected number of push transmissions that reach a node outside the clique is bounded by $\sqrt{n} \cdot \frac{3}{\sqrt{n}+2} \leq 3$. Similarly, the expected number of pull transmissions coming from outside is bounded by $3 \cdot \sqrt{n} \cdot \frac{1}{\sqrt{n}+2} \leq 3$. Hence every clique which is completely informed contributes at most 6 to the expected number of newly informed node.

For the single clique which is not fully informed nor fully uninformed, its contribution in expectation is at most $3\sqrt{n}$ newly informed nodes within the same clique and at most 6 newly informed nodes outside the clique. Hence,

$$\mathbf{E}[|I_{t+1}|] \leq |I_t| + (|I_t|/\sqrt{n}) \cdot 6 + 3 \cdot \sqrt{n} + 6 \leq |I_t| + 7 \cdot \sqrt{n},$$

as long as $|I_t| \leq n/2$. Therefore, $\mathbf{E}[|I_{t+\sqrt{n}/14-1}|] \leq 1 + (\sqrt{n}/14 - 1) \cdot 7\sqrt{n} < n/2$. Hence by Markov's inequality, $\Pr[|I_{t+\sqrt{n}/14-1}| \geq 2\mathbf{E}[|I_{t+\sqrt{n}/14-1}|]] \leq 1/2$. \square

References

1. C. Avin, M. Koucký, and Z. Lotker. How to explore a fast-changing world (cover time of a simple random walk on evolving graphs). In *Proc. 35th ICALP*, pages 121–132, 2008.
2. H. Baumann, P. Crescenzi, and P. Fraigniaud. Parsimonious flooding in dynamic graphs. *Distributed Computing*, 24(1):31–44, 2011.
3. K. Censor-Hillel and H. Shachnai. Fast information spreading in graphs with large weak conductance. *SIAM J. Comput.*, 41(6):1451–1465, 2012.
4. F. Chierichetti, S. Lattanzi, and A. Panconesi. Almost tight bounds for rumour spreading with conductance. In *Proc. 42nd STOC*, pages 399–408, 2010.
5. F. Chierichetti, S. Lattanzi, and A. Panconesi. Rumor spreading in social networks. *Theor. Comput. Sci.*, 412(24):2602–2610, 2011.
6. A. E. F. Clementi, P. Crescenzi, C. Doerr, P. Fraigniaud, M. Isopi, A. Panconesi, F. Pasquale, and R. Silvestri. Rumor spreading in random evolving graphs. In *Proc. 21st ESA*, pages 325–336, 2013.

7. A. E. F. Clementi, A. Monti, F. Pasquale, and R. Silvestri. Information spreading in stationary markovian evolving graphs. *IEEE Trans. Parallel Distrib. Syst.*, 22(9):1425–1432, 2011.
8. A. E. F. Clementi, R. Silvestri, and L. Trevisan. Information spreading in dynamic graphs. In *Proc. 31st PODC*, pages 37–46, 2012.
9. B. Doerr, M. Fouz, and T. Friedrich. Social networks spread rumors in sublogarithmic time. In *Proc. 43rd STOC*, pages 21–30, 2011.
10. C. Dutta, G. Pandurangan, R. Rajaraman, Z. Sun, and E. Viola. On the complexity of information spreading in dynamic networks. In *Proc. 24th SODA*, pages 717–736, 2013.
11. U. Feige, D. Peleg, P. Raghavan, and E. Upfal. Randomized broadcast in networks. *Random Struct. Algorithms*, 1(4):447–460, 1990.
12. N. Fountoulakis, K. Panagiotou, and T. Sauerwald. Ultra-fast rumor spreading in social networks. In *Proc. 23rd SODA*, pages 1642–1660, 2012.
13. C. Georgiou, S. Gilbert, R. Guerraoui, and D. R. Kowalski. Asynchronous gossip. *J. ACM*, 60(2):11, 2013.
14. G. Giakkoupis. Tight bounds for rumor spreading in graphs of a given conductance. In *Proc. 28th STACS*, pages 57–68, 2011.
15. G. Giakkoupis. Tight bounds for rumor spreading with vertex expansion. In *Proc. 25th SODA*, pages 801–815, 2014.
16. G. Giakkoupis and T. Sauerwald. Rumor spreading and vertex expansion. In *Proc. 23rd SODA*, pages 1623–1641, 2012.
17. B. Haeupler and D. R. Karger. Faster information dissemination in dynamic networks via network coding. In *Proc. 30th PODC*, pages 381–390, 2011.
18. B. Haeupler and F. Kuhn. Lower bounds on information dissemination in dynamic networks. In *Proc. 26th DISC*, pages 166–180, 2012.
19. R. Karp, C. Schindelhauer, S. Shenker, and B. Vöcking. Randomized rumor spreading. In *Proc. 41st FOCS*, pages 565–574, 2000.
20. H. Kesten and V. Sidoravicius. The spread of a rumor or infection in a moving population. *Annals of Probability*, 33:2402–2462, 2005.
21. F. Kuhn, N. Lynch, and R. Oshman. Distributed computation in dynamic networks. In *Proc. 42nd STOC*, pages 513–522, 2010.
22. H. Lam, Z. Liu, M. Mitzenmacher, X. Sun, and Y. Wang. Information dissemination via random walks in d -dimensional space. In *Proc. 23rd SODA*, pages 1612–1622, 2012.
23. 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 WWW*, pages 695–704, 2008.
24. D. Mosk-Aoyama and D. Shah. Fast distributed algorithms for computing separable functions. *IEEE Trans. Inf. Theory*, 54(7):2997–3007, 2008.
25. Y. Peres, A. Sinclair, P. Sousi, and A. Stauffer. Mobile geometric graphs: Detection, coverage and percolation. In *Proc. 22nd SODA*, pages 412–428, 2011.
26. A. Pettarin, A. Pietracaprina, G. Pucci, and E. Upfal. Tight bounds on information dissemination in sparse mobile networks. In *Proc. 30th PODC*, pages 355–362, 2011.
27. B. Pittel. On spreading a rumor. *SIAM J. Applied Math.*, 47(1):213–223, 1987.
28. A. D. Sarma, A. R. Molla, and G. Pandurangan. Fast distributed computation in dynamic networks via random walks. In *Proc. 26th DISC*, pages 136–150, 2012.
29. T. Sauerwald and A. Stauffer. Rumor spreading and vertex expansion on regular graphs. In *Proc. 22nd SODA*, pages 462–475, 2011.