

Towards Constant Time Multi-Call Rumor Spreading on Small-Set Expanders

Emilio Cruciani  

European University of Rome, Italy

Sebastian Forster  

University of Salzburg, Austria

Tijn de Vos  

TU Graz, Austria

Abstract

We study a multi-call variant of the classic **PUSH&PULL** rumor spreading process where nodes can contact k of their neighbors instead of a single one during both **PUSH** and **PULL** operations. We show that rumor spreading can be made faster at the cost of an increased amount of communication between the nodes. As a motivating example, consider the process on a complete graph of n nodes: while the standard **PUSH&PULL** protocol takes $\Theta(\log n)$ rounds, we prove that our k -**PUSH&PULL** variant completes in $\Theta(\log_k n)$ rounds, with high probability.

We generalize this result in an expansion-sensitive way, as has been done for the classic **PUSH&PULL** protocol for different notions of expansion, e.g., conductance and vertex expansion. We consider small-set vertex expanders, graphs in which every sufficiently small subset of nodes has a large neighborhood, ensuring strong local connectivity. In particular, when the expansion parameter satisfies $\phi > 1$, these graphs have a diameter of $o(\log n)$, as opposed to other standard notions of expansion. Since the graph's diameter is a lower bound on the number of rounds required for rumor spreading, this makes small-set expanders particularly well-suited for fast information dissemination. We prove that k -**PUSH&PULL** takes $O(\log_\phi n \cdot \log_k n)$ rounds in these expanders, with high probability. We complement this with a simple lower bound of $\Omega(\log_\phi n + \log_k n)$ rounds.

2012 ACM Subject Classification Theory of computation $\rightarrow$ Distributed algorithms; Theory of computation $\rightarrow$ Graph algorithms analysis; Mathematics of computing $\rightarrow$ Probabilistic algorithms

Keywords and phrases small set expansion, vertex expansion, rumor spreading, multi-call rumor spreading, push&pull protocol

Digital Object Identifier 10.4230/LIPIcs.DISC.2025.26

Related Version *Full Version*: <https://arxiv.org/abs/2508.18017>

Funding This project has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No 947702) and is supported by the Austrian Science Fund (FWF): P 32863-N and P 36280-N.

1 Introduction

Rumor spreading, which is loosely inspired by biological and social phenomena, is one of the most well-studied stochastic processes on graphs with a rich history in refined analyses [30, 48, 27, 38, 44, 28, 13, 11, 21, 52, 20, 22, 29, 34, 6, 32, 23, 25, 33, 41, 1, 46, 17, 31, 24]. In the classic *random phone call model* [19], an arbitrary node of the graph initially knows a piece of information, the rumor, which spreads across the graph in synchronous rounds until eventually all nodes are informed. In each round, every node calls one of its neighbors uniformly at random: in the **PUSH** protocol each informed node informs the node it calls, in the **PULL** protocol each uninformed node is informed by the node it called, and in the **PUSH&PULL** protocol both of these information exchanges happen simultaneously.



© Emilio Cruciani, Sebastian Forster, and Tijn de Vos;
licensed under Creative Commons License CC-BY 4.0

39th International Symposium on Distributed Computing (DISC 2025).

Editor: Dariusz R. Kowalski; Article No. 26; pp. 26:1–26:25

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

It is well known that PUSH, PULL, and PUSH&PULL on a *complete graph* with n nodes require $\Theta(\log n)$ rounds to complete rumor spreading with high probability¹ [30, 38]. Rumor spreading has been also extensively studied on *expander graphs*, with known tight bounds that relate the completion time of the process with different notions of expansion. In the following, we discuss only PUSH&PULL, due to strong lower bounds that exist for both PUSH and PULL in isolation [11, 52]. In particular, the protocol with high probability requires $\Theta(\varphi^{-1} \log n)$ rounds for graphs with *conductance* $\varphi \in [0, 1]$ [11], and $\Theta(\phi^{-1} \log n \log \Delta)$ rounds for graphs with maximum degree Δ and *vertex-expansion* $\phi \in [0, 1]$ [34, 32]. It is also known that PUSH&PULL requires $\Omega(\sqrt{n})$ rounds with high probability in graphs with constant *edge-expansion* [13]. The study of rumor spreading protocols on expanders, subsequently, has been used as a building block for designing distributed information dissemination algorithms [9, 10, 8, 35], some of which are based on expander decompositions [54].

Faster rumor spreading can be achieved on specific topologies or by modifying the process. On power-law degree distributed graphs, the process completes in $O(\log \log n)$ rounds if the exponent of the power-law β is between 2 and 3, while it needs $\Omega(\log n)$ rounds if β is greater than 3 [29]. A variant of the process, where nodes can use direct addressing, i.e., are able to contact neighbors whose address was learned before, takes $\Theta(\log \log n)$ rounds [36, 6].

With the goal of making rumor spreading even faster, in this paper we study a *multi-call variant* of the classical protocol. We call this the k -PUSH&PULL protocol, where in each round every informed node samples k neighbors and sends the rumor to them, while every uninformed node samples k neighbors and requests the rumor from them. The sampling happens uniformly at random and with replacement in both PUSH and PULL operations.

This process naturally interpolates between classical single-call rumor spreading and full broadcasting. It involves random selective communication, as in the former, while nodes simultaneously interact with multiple neighbors, as in the latter. This interpolation intuitively results in a trade-off between the amount of communication and the spreading time of the two processes. Classical rumor spreading requires minimal communication, while broadcasting is optimal in spreading time, matching the graph's diameter in the worst case.

Somewhat similar processes have been studied for the special case $k = 4$ on random regular graphs [7], with a randomized number of simultaneous calls on the complete graph [47, 24], and in the asynchronous setting considering multiple pulls on the complete graph [42, 50]. For the related random-walk process, a multi-walk extension similar in spirit to our multi-call extension for rumor spreading has been studied by Alon et al. [2]. Apart from these works we are not aware of any paper studying the multi-call setting explicitly. While the analogy to a biological or social process might not be given anymore for the multi-call variant, we still believe that increasing the number of simultaneous calls is a quite natural generalization that captures a relevant aspect of information dissemination under bandwidth constraints. Furthermore, note that the classic “single-call” variant allows each node to *receive* multiple calls at the same time.² Therefore, it seems natural to also allow nodes to *initiate* several calls at the same time.

This leads to the central question of our work:

Question: *How much communication is required for sub-logarithmic time rumor spreading?*

¹ We say that a statement holds *with high probability* (w.h.p. for short) if it holds with probability at least $1 - n^{-c}$, where c is a positive constant hidden in the asymptotic notation.

² Restricted variants in which nodes receiving multiple calls can spread the rumor only to a single neighbor [17] cannot achieve performance guarantees comparable to the unrestricted variant in conductance expanders and vertex expanders [31].

1.1 Our Contribution

The starting point of our investigation is the study of multi-call rumor spreading on the *complete graph*. Using standard arguments that we present as a warm-up in Section B, we show that the k -PUSH&PULL protocol requires $\Theta(\log_k n)$ rounds with high probability. This gives an interesting trade-off between the classic single-call bound of $\Theta(\log n)$ and a constant number of rounds for a polynomial number of simultaneous calls.

Next, we extend our question to *expander graphs*. Our goal is to generalize the trade-off that we get for k -PUSH&PULL on the complete graph to incorporate an expansion parameter ϕ , similarly to the works mentioned above for the single-call model. While this is an intriguing combinatorial question in its own right, we believe that the recent use of expanders in graph neural networks [18, 53] adds additional motivation for revisiting the foundations of information dissemination in expanders.

Our focus is on *small-set expanders*, that have initially been studied for the notion of conductance (see e.g., [49, 3]), and later for edge expansion (see e.g., [43]) and vertex expansion (see e.g., [16, 39, 14]). Small-set vertex expansion has also been identified as a crucial property of networks for designing routing protocols [4]. A non-negligible part of our contribution consists of studying properties of *small-set vertex expanders* in a particular regime where the vertex expansion ϕ is larger than 1. In particular, we establish upper and lower bounds on the *diameter* of such graphs (Section C), which, unlike other standard notions, is *sub-logarithmic*.

This property is crucial for achieving sub-logarithmic rumor spreading time, as the graph diameter provides a fundamental lower bound on the number of rounds required. Roughly speaking, our main result (see Section 2 and Theorem 2) implies that, with high probability, a constant number of rounds suffices for multi-call rumor spreading in a wide regime of small-set vertex expanders when both the number of calls k and the expansion parameter ϕ are polynomial in n . We also give a complementary lower bound (see Theorem 3) showing that both parameters need to be polynomial to achieve a constant number of rounds.

1.2 Formal Statement of Our Results

Small-set vertex expanders

Let $G = (V, E)$ be an unweighted graph. For every $S \subseteq V$, we write $\bar{S} := V \setminus S$ for the complement of S . We also define the neighborhood of a set S as $N(S) := \{v \in V : \exists s \in S, \{s, v\} \in E\}$ and its boundary $\partial S := N(S) \cap \bar{S}$. For a single node v , we write $N(v) := N(\{v\})$. Note that $v \notin N(v)$, but $N(S) \cap S$ can be nonempty for larger S . Using this notation, we are ready to define the class of *small-set vertex expanders* (see, e.g., [16, 39, 14]).

► **Definition 1.** Let $\phi \in (0, n)$ and $\alpha \in (0, \frac{1}{2}]$. We say a graph G with n nodes is a (ϕ, α) -(vertex) expander if

$$\min_{\substack{S \subseteq V \text{ s.t.} \\ 0 < |S| \leq \alpha n}} \frac{|\partial S|}{|S|} \geq \phi.$$

The name small-set comes from the fact that the expansion property holds for sets of size at most αn , opposed to the classical definition where $\alpha = \frac{1}{2}$. The regime $\alpha < \frac{1}{2}$ is conceptually different and has implications on the expansion parameter ϕ , allowing $\phi > 1$. When $\alpha = \frac{1}{2}$, in fact, we *cannot* have $\phi > 1$. Indeed, any set S with $|S| = n/2$ nodes satisfies $\frac{|\partial S|}{|S|} \leq \frac{|V \setminus S|}{|S|} = 1$. More generally, there are no (ϕ, α) -expanders for $\alpha > \frac{1}{1+\phi}$ (see Lemma 18).

However, if we take $\alpha = \frac{1}{1+\phi}$, we only get a very restricted class of graphs: such (ϕ, α) -expanders have extremely low diameter and are very dense; they have diameter 2 and minimum degree $\Theta(n)$ (see Lemma 19). On the other hand, if we take α too small, the expansion property becomes local and the graph is no longer necessarily connected. To be precise, for $\alpha \leq \frac{1}{2+2\phi}$, there exist (ϕ, α) -expanders that are disconnected (see Lemma 20). In between these two bounds, we get a guarantee on the diameter that depends on ϕ . For $\frac{1}{2+2\phi} < \alpha \leq \frac{1}{1+\phi}$, we have that (ϕ, α) -expanders have diameter at most $O(\log_\phi n)$ (see Lemma 21). We note that for $\phi = \omega(1)$ the diameter is $o(\log n)$ and that for polynomial ϕ it becomes constant.

Next, we provide two examples of vertex expanders where the bound on the diameter is tight. First of all, we obtain that complete graphs are (ϕ, α) -vertex expanders for $0 < \phi \leq n-1$ and $\alpha \leq \frac{1}{1+\phi}$ (see Lemma 22). Secondly, we show that random graphs are vertex expanders. More precisely, we show that Erdős-Rényi random graphs, where between each pair of nodes there is an edge with probability $p = \frac{3\phi}{n}$, are (ϕ, α) -vertex expanders for $\alpha \leq \frac{1}{1+1.6\phi}$ (see Lemma 23). We note that in this regime of p these graphs have diameter $\Theta(\log_\phi n)$ with high probability [15].

We also note that concurrent work by Hsieh et al. [37] gives an explicit $\Theta(\phi)$ -regular construction for small-set vertex expanders.

Rumor Spreading

For $\frac{1}{2+2\phi} < \alpha \leq \frac{1}{1+\phi}$, we study rumor spreading on (ϕ, α) -vertex expanders through the k -PUSH&PULL protocol.

► **Theorem 2.** *Let $\phi > 1$, $\alpha > \frac{1}{2+2\phi}$, let $G = (V, E)$ be a (ϕ, α) -expander, and let $k > \log^3 n$. Then w.h.p. rumor spreading with the k -PUSH&PULL protocol requires the following number of rounds:*

$$O\left(\left(\log_\phi n + \phi^{-1}\left(\alpha - \frac{1}{2+2\phi}\right)^{-1}\right) \log_k n\right).$$

Before commenting on this round complexity, we give a lower bound that follows immediately from the fact that on complete graphs we need $\Omega(\log_k n)$ rounds (see Lemma 17) and on Erdős-Rényi random graphs (which are (ϕ, α) -expanders by Lemma 23) we need $\Omega(\log_\phi n)$ rounds, since they have diameter $\Omega(\log_\phi n)$ [15].

► **Theorem 3.** *Let $n \geq 1$, $\phi > 1$, $\alpha \leq \frac{1}{1+1.6\phi}$ and $k \geq 2$. Then there exist a (ϕ, α) -expander $G = (V, E)$ on n nodes such that w.h.p. rumor spreading with the k -PUSH&PULL protocol requires $\Omega(\log_\phi n + \log_k n)$ rounds.*

We make the following observations about the round complexity we get in Theorem 2:

- The round complexity goes to infinity as α approaches $\frac{1}{2+2\phi}$, where expanders can be disconnected and hence rumor spreading never finishes. Interestingly, an analogous term appears in the analysis of the classic protocol on random graphs, making the rumor spreading time go to infinity in the non-connected regime [46].
- For $\alpha \geq \frac{1}{2+2\phi} + \Omega\left(\frac{1}{\phi \log_\phi n}\right)$ the number of rounds simplifies to $O(\log_k n \log_\phi n)$.
- If additionally $\phi = n^{\Omega(1)}$, then we obtain $O(\log_k n)$, which is optimal due to Theorem 3.
- If instead $k = n^{\Omega(1)}$, then we obtain $O(\log_\phi n)$, which is optimal due to Theorem 3.

We believe that our work opens several interesting research directions for the community. Closing the gap between our upper and lower bound is left as an open problem. Extending the proof technique of the state of the art analysis [32] to the multi-call setting in our opinion

appears to be non trivial; it is also not clear if this would be sufficient to actually close the gap. Another direction is to explore the use of such expanders in applications beyond rumor spreading. They can serve as a natural alternative to traditionally used low-diameter graphs, with the benefit of having combinatorial expansion properties.

Lower bounds

The fact that small-set vertex expanders allow $\phi > 1$ and hence sub-logarithmic diameter is fundamental to achieve our results. We observe that the classical notions of expansions considered in the literature do not have these properties and, therefore, the multi-call variant cannot speed up the rumor spreading time on these graphs. The most studied notions of expanders for which these properties do not hold are the following:

1. ϕ -conductance expanders, where for $\phi \in (0, 1)$ we have

$$\min_{\substack{S \subseteq V \text{ s.t.} \\ 0 < \text{Vol}(S) \leq m/2}} \frac{|E(S, \bar{S})|}{\text{Vol}(S)} \geq \phi,$$

where $|E(S, \bar{S})|$ is the number of edges crossing the cut $(S, \bar{S})$, and $\text{Vol}(S) := \sum_{u \in S} \deg(u)$ is the volume of S , with $\deg(u)$ being the degree of a node $u \in V$.

2. ϕ -edge expanders, where for $\phi \in (0, n)$ we have

$$\min_{\substack{S \subseteq V \text{ s.t.} \\ 0 < |S| \leq n/2}} \frac{|E(S, \bar{S})|}{|S|} \geq \phi.$$

3. ϕ -vertex expanders, where for $\phi \in (0, 1)$ we have

$$\min_{\substack{S \subseteq V \text{ s.t.} \\ 0 < |S| \leq n/2}} \frac{|\partial S|}{|S|} \geq \phi.$$

Regarding conductance-expanders, the following statement holds.

► **Lemma 4.** *Let $n \geq 1$, $\phi \in (0, 1)$ and $k \geq 1$. There exists a ϕ -conductance expander on n nodes such that k -PUSH&PULL takes $\Theta(\phi^{-1} \log n)$ rounds.*

The proof follows since the same upper bound holds already for $k = 1$, and the lower bound comes from the diameter of ϕ -conductance expanders [11].

Rumor spreading on ϕ -edge expanders on the other hand, is not very appealing in the first place due to the strong lower bound of $\Omega(\sqrt{n})$ rounds for the single-call variant – which even holds in edge expanders with low diameter [13].

Regarding vertex expanders, we have the following statement.

► **Lemma 5.** *Let $n \geq 1$, $\phi \in (0, 1)$ and $k \geq 1$. There exists a ϕ -vertex expander on n nodes such that k -PUSH&PULL takes $\Omega(\phi^{-1} \log n)$ rounds.*

The lemma follows from a lower bound on the diameter of ϕ -vertex expanders [34]. A tight bound of $\Theta(\phi^{-1} \log n \log \Delta)$ is known for the single-call variant when $\phi \leq 1$ [34, 32]. This means that the potential for savings in the multi-call variant is limited to a single $\log \Delta$ -factor compared to the single-call variant; eliminating this factor is, to the best of our judgment, a quite fine-grained endeavor that we deliberately omit from this paper with its more conceptual focus.

Another class of graphs with very good expansion properties is that of *Ramanujan graphs*, whose spectral gap in the matrix representation of the graph is almost as large as possible. In particular, they can achieve sublogarithmic diameter, but only under very restrictive conditions, i.e., when they are $\Theta(n)$ -regular [51]. In contrast, our class of small-set vertex expanders allows sublogarithmic diameter while being significantly sparser.

1.3 Technical Challenges

The proof we provide for rumor spreading in vertex expanders with $\phi > 1$ is based on the ideas for the case of expansion at most 1, by Giakkoupis and Sauerwald [34]. There are three complications we need to overcome. Let I_t denote the informed nodes in round t . We will show that in any round either I_t grows significantly or the boundary ∂I_t grows significantly. While this growth is a constant factor for the case $\phi < 1$, leading to a $\log n$ in the round complexity, we need to grow more aggressively. Intuitively, if the expected growth in one round is μ , then the expected growth in a round with k -parallel calls is $k\mu$. However, this brings us to the first issue.

Overlap of Parallel Calls

To some extent this challenge already appears in the PUSH model: the fact that two nodes individually push the rumor to a neighbor does not mean two new nodes are informed; they can push the rumor to the same node. This phenomenon is exacerbated when we call k times. Moreover, it now also impacts pulling: if one node has two successful parallel pulls in a round, then this does not contribute as two nodes pulling. In other words, we cannot simply sum the expected gain from separate pushes and pulls to obtain the expected number of newly informed nodes. Let us consider these probabilities to understand this situation better.

Let $u \in \partial I_t$ be an uninformed node. The probability that u pulls the rumor from I_t in one round of k -PUSH&PULL is

$$1 - \left(1 - \frac{|N(u) \cap I_t|}{\deg(u)}\right)^k = p.$$

Now if $\deg(u) \geq k|N(u) \cap I_t|$, we can use a binomial expansion to lower bound this probability by $\frac{k|N(u) \cap I_t|}{2\deg(u)}$, which is a factor $k/2$ higher than the probability that a single call succeeds. For nodes with lower degree we note that $p \geq 1/2$, so we can still say that they get informed with good probability. However, this does not carry over a factor k compared to the single-call situation. We are able to show a sufficient growth by carefully analyzing the number of nodes with high and low degrees and using different arguments for each case.

In this brief description we have only highlighted the issue. The actual solution is more involved but based on this intuition. For details see Section 3.2 and Section 3.3.

Probabilistic Guarantees

The second issue we need to deal with is the probability of successfully growing I_t or ∂I_t through k -PUSH&PULL. Intuitively, there are two sides to what is going on here: on the negative side, we have only a sub-logarithmic number of rounds, which is often not great to give bounds with high probability. On the positive side, we have significant growth in each step, which is good for tail bounds.

From this second observation, we see that in some cases a Chernoff bound suffices. However, in some situations, we cannot use a simple Chernoff bound since the random variables are neither independent nor negatively correlated. In particular, this happens when we want to show that the boundary ∂I_t grows. Let $u \in \partial I_t$ be a node in the boundary of I_t and $v_1, v_2 \in N(u) \setminus I_t$ be two uninformed neighbors of u . The events “ $v_1 \in \partial I_{t+1}$ ” and “ $v_2 \in \partial I_{t+1}$ ” are not independent nor negatively associated as both may occur as a consequence of the event “ $u \in I_{t+1}$ ”. In some cases we can use the bounded difference inequality, see Section A, that can give tail bounds on *functions* of independent random variables which we design to be equal to the sum of the correlated ones discussed above. See for example Section 3.2 and Section 3.3.

In other cases, the function that measures progress does not satisfy the conditions for the bounded difference inequality. Here we use an additional trick: we do not use the full potential of k parallel calls, but only $k/\log n$ calls. Each batch of $k/\log n$ call succeeds with constant probability, and since we have $\log n$ of these independent batches in parallel, at least one succeeds with high probability. This comes at the drawback that we only grow a factor $k/\log n$. By our assumption that $k > \log^3 n$, we get $k/\log n > k^{2/3}$, which means that $\log_{k/\log n} n = O(\log_k n)$.

No Expansion for Large Sets

The last complication we point out here is that small-set expansion is, by definition, only guaranteed for sets S of size at most αn . Therefore, we cannot use the expansion property to argue that I_t keeps growing until it hits $n/2$ nodes – at which point a standard symmetry argument shows that the process completes within a factor 2 in the round complexity. To handle the case where more than αn nodes are informed, we show that the size of the boundary ∂I_t is at least a constant fraction of n , meaning that if the set of informed nodes I_t keeps growing by a constant fraction of the boundary ∂I_t , the process still completes in a constant number of additional iterations. To be precise, this leads to the factor $\phi^{-1}(\alpha - \frac{1}{2+2\phi})^{-1}$ in our round complexity.

2 Multi-Call Rumor Spreading on Small-Set Vertex Expanders

The goal of this section is to prove the following theorem, which extends the study of multi-call rumor spreading on small-set vertex expanders with expansion $\phi > 1$ (see Definition 1).

► **Theorem 2.** *Let $\phi > 1$, $\alpha > \frac{1}{2+2\phi}$, let $G = (V, E)$ be a (ϕ, α) -expander, and let $k > \log^3 n$. Then w.h.p. rumor spreading with the k -PUSH&PULL protocol requires the following number of rounds:*

$$O\left(\left(\log_\phi n + \phi^{-1}\left(\alpha - \frac{1}{2+2\phi}\right)^{-1}\right) \log_k n\right).$$

The proof of this theorem is based on the ideas for the case of expansion at most 1, by Giakkoupis and Sauerwald [34]. Where applicable, we follow their notation.

Let I_t denote the set of informed nodes in round t . Our intermediate goal is to show that either I_t grows or ∂I_t grows. We do this by partitioning the nodes in the boundary according to their degree, and then analyzing each set separately. More formally, we partition the nodes in the boundary ∂I_t into different sets A_i defined as

$$A_i := \{u \in \partial I_t : d_i \leq \deg(u) < 2d_i\}, \text{ with } d_i := 2^{i-1}.$$

Note that there are at most $\log n$ of such sets. We consider only those A_i 's that are sufficiently big and such that the degree of the nodes in the set is bounded by quantities depending on the size of the set itself or is large with respect to the size of the boundary ∂I_t . Formally, let

$$\mathcal{I} := \left\{ i : |A_i| \geq \frac{|\partial I_t|}{4 \log n} \wedge (d_i \leq 16 |A_i| \vee d_i \geq 2 |\partial I_t|) \right\} \quad (1)$$

be the set of indices of the A_i 's we take into account. We note that by considering such sets only, we consider at least half of the nodes in the boundary ∂I_t . We include a proof for completeness.

► **Lemma 6** ([34]). *It holds that $|\bigcup_{i \in \mathcal{I}} A_i| \geq |\partial I_t|/2$.*

Proof. The worst case with respect to the condition $|A_i| \geq \frac{|\partial I_t|}{4 \log n}$ is that all but one of the A_i 's have size $|A_i| = \lceil |\partial I_t| / 4 \log n \rceil - 1 < |\partial I_t| / 4 \log n$. The worst case with respect to the condition $d_i \leq 16 |A_i| \vee d_i \geq 2 |\partial I_t|$ is that $|A_i| = \lceil d_i / 16 \rceil - 1 \leq d_i / 16$ and $d_i < 2 |\partial I_t|$, for all $i < \log(2 |\partial I_t|)$. Hence, since the A_i 's are disjoint, we have

$$\left| \bigcup_{i \notin \mathcal{I}} A_i \right| = \sum_{i \notin \mathcal{I}} |A_i| \leq \frac{|\partial I_t|}{4 \log n} \log n + \frac{1}{16} (2 |\partial I_t| + |\partial I_t| + |\partial I_t|/2 + |\partial I_t|/4 + \dots + 2) \leq \frac{|\partial I_t|}{2}$$

which implies the thesis. $\blacktriangleleft$

In Section 3 we show the following lemma, which is the main technical lemma used in the proof of Theorem 2.

► **Lemma 7.** *For each $i \in \mathcal{I}$ and at every round t , we have that with high probability either at least $\frac{|A_i|}{128}$ nodes from A_i are informed within $r = O(\log_k n)$ rounds, or the boundary ∂I_t grows by at least $\frac{k^{1/6}}{64} |\partial I_t|$ within $O(1)$ rounds.*

Lemmas 6 and 7 lead to the following claim.

► **Lemma 8.** *If $|I_t| \leq \frac{n}{2}$, for $r = O(\log_k n)$ we have that with high probability*

$$|I_{t+r} \setminus I_t| \geq \frac{1}{256} |\partial I_t|.$$

Proof. First, we show that in each phase of r' rounds with high probability we are in either of the following cases:

- a) $|I_{t+r'} \setminus I_t| \geq \frac{1}{256} |\partial I_t|$, for $r' = O(\log_k n)$,
- b) $|\partial I_{t+r'} \setminus \partial I_t| \geq \frac{k^{1/6}}{64} |\partial I_t|$, for $r' = O(1)$.

Lemma 7 guarantees that for each $i \in \mathcal{I}$ with high probability either $\frac{|A_i|}{128}$ nodes from A_i are informed within $r = O(\log_k n)$ rounds, or the boundary ∂I_t grows by at least $\frac{k^{1/6}}{64} |\partial I_t|$ within $r = O(1)$ rounds.

The thesis follows by considering two cases. If there is $i \in \mathcal{I}$ such that the boundary grows then the thesis follows immediately. Otherwise, it must hold that for all $i \in \mathcal{I}$ there are at least $\frac{|A_i|}{128}$ nodes in A_i which get informed. Since the A_i 's are disjoint, by Lemma 6 we conclude that

$$|I_{t+r} \setminus I_t| \geq \sum_{i \in \mathcal{I}} \frac{|A_i|}{128} = \frac{1}{128} \left| \bigcup_{i \in \mathcal{I}} A_i \right| \geq \frac{|\partial I_t|}{256}.$$

Next, we notice that Case b can only occur $O(\log_{\frac{k^{1/6}}{64}} n) = O(\log_k n)$ consecutive times before we cover the entire graph. Hence after $r = O(\log_k n) + O(\log_k n) = O(\log_k n)$ rounds, Case a holds. $\blacktriangleleft$

Next, we use this lemma to show that we reach more than half of the graph. We show that in $O(\log_\phi n)$ phases of $O(\log_k n)$ rounds we reach αn nodes. Let $I_{(t)}$ denote the set of informed nodes in phase t .³ We start by showing by induction that until $|I_{(t)}| \geq \alpha n$, we have $|\partial I_{(t)}| \geq (\frac{\phi}{256})^t$.

³ Note that we use I_t for the set of informed nodes in round t .

- **Base case.** By expansion, every node $v \in V$ has $\deg(v) = |\partial\{v\}| \geq \phi|\{v\}| = \phi$. So we have that $|\partial I_{(1)}| \geq \phi$.
 - **Induction step.** Suppose that $|\partial I_{(t)}| \geq (\frac{\phi}{256})^t$. By Lemma 8 have $|I_{(t+1)}| \geq \frac{1}{256}|\partial I_{(t)}|$. So suppose $|I_{(t+1)}| \leq \alpha n$, then we have $|\partial I_{(t+1)}| \geq \phi|I_{(t+1)}|$ by expansion of $I_{(t+1)}$. We obtain $|\partial I_{(t+1)}| \geq \frac{\phi}{256}|\partial I_{(t)}| \geq \frac{\phi}{256}(\frac{\phi}{256})^t \geq (\frac{\phi}{256})^{t+1}$, using the induction hypothesis. So we conclude that in $O(\log_{\phi/256}(\alpha n)) = O(\log_{\phi} n)$ phases of $O(\log_k n)$ rounds we cover αn nodes, using in total $O(\log_{\phi} n \cdot \log_k n)$ rounds.
- Now suppose that $\alpha n < |I_{(t)}| \leq \frac{n}{2}$. Let $S \subseteq I_{(t)}$ be any subset of size exactly $|S| = \alpha n$, so $|I_{(t)} \setminus S| \leq \frac{n}{2} - \alpha n$. By expansion of S , we have $|\partial S| \geq \phi|S| = \phi\alpha n$. We now see that
- $$|\partial I_{(t)}| \geq |\partial S \setminus I_{(t)}| = |\partial S \setminus (I_{(t)} \setminus S)| \geq |\partial S| - |I_{(t)} \setminus S| \geq \phi\alpha n - (\frac{n}{2} - \alpha n) = (\phi + 1)\alpha n - n/2.$$

By Lemma 8 we now have to have that in each phase of $O(\log_k n)$ rounds, $I_{(t)}$ grows by

$$\frac{1}{256}|\partial I_{(t)}| \geq \frac{1}{256}((\phi + 1)\alpha n - n/2).$$

So we cover more than $n/2$ nodes in the following number of phases

$$\frac{n/2}{\frac{1}{256}((\phi + 1)\alpha n - n/2)} = O\left(\frac{1}{(\phi + 1)\alpha - 1/2}\right) = O\left(\frac{1}{\phi(\alpha - \frac{1}{2+2\phi})}\right).$$

As each phase last $O(\log_k n)$ rounds, in total, we see that the rumor spreads to more than $n/2$ nodes in the following number of rounds

$$R = O\left(\left(\log_{\phi} n + \frac{1}{\phi(\alpha - \frac{1}{2+2\phi})}\right) \log_k n\right).$$

Now Lemma 15 states that the number of rounds we need to spread the rumor from I_t to a node $u \in V \setminus I_t$ is the same as the number of rounds we would need to spread the rumor from u to anywhere in I_t . Since we have shown that we reach more than half of the graph in R rounds, and that $|I_t| > n/2$, a rumor started at u reaches I_t in another R rounds. We conclude that the spreading process finishes after $2R$ rounds in total, proving Theorem 2.

3 Proof of Lemma 7

We look at each set A_i for $i \in \mathcal{I}$ separately, and we upper bound the number of rounds until either a large fraction of nodes in A_i gets informed or the size of the boundary increases by a large factor as a result of nodes in A_i being informed. We distinguish three cases informally involving sets of nodes with low, medium, and high degree. We define these cases according to the following conditions:

1. Low degree: $d_i \leq 1024 \cdot 96 \cdot k$.
2. Medium degree: $1024 \cdot 96 \cdot k < d_i \leq 16|A_i|$.
3. High degree: $d_i \geq 2|\partial I_t|$.

Note that the three cases do not cover all nodes in the boundary ∂I_t . However, they do cover all nodes in the $|A_i|$'s for $i \in \mathcal{I}$. This is important since, due to Lemma 6, it implies that we consider at least half of the nodes in ∂I_t .

We analyze the above three cases separately in the three following subsections, that will be further subdivided into a total of 7 cases: in 4 cases I_t grows, while in 3 cases ∂I_t grows, always with high probability. By summarizing all of them we get with high probability that either

$$|I_{t+r} \setminus I_t| \geq \min\left\{\frac{1}{4}, \frac{1}{20}, \frac{1}{128}, \frac{1}{8}\right\} |A_i| = \frac{|A_i|}{128}$$

within $r = O(\max\{1, \log_k n, 1, 1\}) = O(\log_k n)$ rounds or

$$|\partial I_{t+r} \setminus \partial I_t| \geq \min \left\{ \frac{k}{32 \log n}, \frac{\sqrt{k}}{4}, \frac{k}{64} \right\} |A_i| \geq \frac{k^{1/6}}{64} |\partial I_t|$$

within $r = O(1)$ rounds, since by definition $|A_i| \geq \frac{|\partial I_t|}{4 \log n}$ and by assumption $k > \log^3 n$.

3.1 Sets with low degree nodes

Recall that in this case every $u \in A_i$ is such that

$$d_i \leq 1024 \cdot 96 \cdot k.$$

Since the degree of these boundary nodes is low, it is likely that many of them will directly pull the rumor. We prove a more general lemma that will be used in other cases as well and formalizes the fact that nodes that are “well-connected” to I_t are likely to pull the rumor directly. In particular we consider nodes of a generic subset $B \subseteq \partial I_t$ of the boundary of I_t and bound the number of rounds needed by a constant fraction of nodes in B to pull the rumor from I_t , with high probability.

► **Lemma 9.** *Let $B \subseteq \partial I_t$ be such that $|B| = \Omega(\log n)$. Let $0 < q \leq 1$ and suppose that for every node $u \in B$ at least a q -fraction of its neighbors is in I_t . Then with high probability at least $\frac{|B|}{4}$ nodes of B have pulled the rumor from I_t in $\lceil \frac{1}{qk} \rceil$ rounds.*

Proof. We consider two different cases, depending on the value of q . We show for both cases that the probability of pulling from an informed node is at least $1/2$.

Case 1. If $qk > 1$, we show we only need 1 round. The probability that any $u \in B$ pulls from I_t in one round is $1 - \left(1 - \frac{|N(u) \cap I_t|}{\deg(u)}\right)^k \geq 1 - (1 - q)^k$. Since $qk > 1$, then $1 - (1 - q)^k \geq 1/2$.

Case 2. If $qk \leq 1$, then the probability that a node pulls the rumor in $\lceil \frac{1}{qk} \rceil$ rounds is $1 - (1 - q)^{k \cdot \frac{1}{qk}} \geq 1 - e^{-1} \geq 1/2$.

This means that the expected number of nodes that pull in this many rounds is at least $|B|/2$. Since these pulls are independent random variables and $|B| = \Omega(\log n)$, a Chernoff bound gives us that at least $|B|/4$ nodes are informed with high probability. ◀

We are now ready to show that the set A_i grows sufficiently. For each node $u \in A_i$ the fraction of neighbors that u has in I_t is at least $1/\deg(u) \geq 1/(2d_i)$. Then Lemma 9 gives that with high probability a fraction $1/4$ of the nodes in A_i pull the rumor from I_t in at most $r = \lceil 2 \cdot 1024 \cdot 96 \cdot k/k \rceil = O(1)$ rounds. By definition of $\mathcal{I}$ in Equation (1), with high probability it follows that

$$|I_{t+r} \setminus I_t| \geq \frac{1}{4} |A_i|.$$

3.2 Sets with medium degree nodes

Recall that in this case every $u \in A_i$ is such that

$$1024 \cdot 96 \cdot k < d_i \leq 16 |A_i|.$$

Unlike the previous case, nodes with medium degree can either contribute directly to the growth of I_t or to the growth of ∂I_t . To describe these different types of contributions, for each node $v \in V$ let us denote the number of neighbors of v in A_i as

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

Moreover, let us define the set of nodes that are neither informed or in the boundary of the informed nodes as

$$S_t := V \setminus (I_t \cup \partial I_t).$$

We further distinguish three cases, depending on the *volume* of A_i , we define the volume of a set S as $\text{Vol}(S) := \sum_{v \in S} \deg(v)$.

- i) It holds that $\sum_{v \in S_t \text{ s.t. } h_i(v) \leq d_i \log n/k} h_i(v) \geq \frac{1}{2} \text{Vol}(A_i)$.
- ii) It holds that $\sum_{v \in V \text{ s.t. } h_i(v) \geq d_i \log n/96k} h_i(v) \geq \frac{1}{3} \text{Vol}(A_i)$.
- iii) None of the above conditions are met, i.e., it holds that $\sum_{v \in S_t \text{ s.t. } h_i(v) \leq d_i \log n/k} h_i(v) < \frac{1}{2} \text{Vol}(A_i)$ and $\sum_{v \in V \text{ s.t. } h_i(v) \geq d_i \log n/96k} h_i(v) < \frac{1}{3} \text{Vol}(A_i)$.

In the remainder of this subsection we prove that in case i) ∂I_t grows, while in cases ii) and iii) it is I_t to grow directly.

Case i)

It holds that

$$\sum_{\substack{v \in S_t \text{ s.t.} \\ h_i(v) \leq d_i \log n/k}} h_i(v) \geq \frac{1}{2} \text{Vol}(A_i). \quad (2)$$

In this case we give a lower bound on the number of new nodes in the boundary and prove that it holds with high probability. We start by noting that the probability that a fixed node $u \in A_i$ pulls the rumor from I_t in a given round is

$$1 - \left(1 - \frac{|N(u) \cap I_t|}{\deg(u)}\right)^k \geq 1 - \left(1 - \frac{1}{\deg(u)}\right)^k \geq 1 - \left(1 - \frac{1}{2d_i}\right)^k \geq \frac{k}{4d_i} =: p,$$

where we use that $2d_i > k$.

Let us pessimistically assume that such a probability exactly equals p for every node u . For each $u \in A_i$, let X_u denote the 0/1 random variable that is 1 iff u pulls the rumor from I_t in round $t + 1$. Then we have $\mathbb{P}[X_u = 1] = p$. Further, for each $v \in S_t$ such that $h_i(v) \leq d_i \log n/k$, let Y_v denote the 0/1 random variable which is 1 exactly when v has a neighbor u with $X_u = 1$. Finally let $Y = \sum_{v \in V} Y_v$ be the number of new nodes in the boundary only considering the contribution from pull. Our goal is to prove a lower bound on Y .

We see that

$$\mathbb{P}[Y_v = 1] \geq 1 - \left(1 - \frac{k}{4d_i}\right)^{h_i(v)} \geq 1 - \left(1 - \frac{k}{4d_i}\right)^{h_i(v)/\log n} \geq \frac{kh_i(v)}{8d_i \log n},$$

where the last inequality follows from Taylor series expansion, which needs that $\frac{kh_i(v)}{8d_i \log n} \leq 1$, as ensured by the second to last inequality⁴. This gives us that in expectation

$$\begin{aligned} \mathbb{E}[Y] &= \sum_{\substack{v \in S_t \text{ s.t.} \\ h_i(v) \leq d_i \log n/k}} \mathbb{P}[Y_v = 1] \geq \sum_{\substack{v \in S_t \text{ s.t.} \\ h_i(v) \leq d_i \log n/k}} \frac{kh_i(v)}{8d_i \log n} \\ &\geq \frac{k}{16d_i \log n} \text{Vol}(A_i) \geq \frac{k}{16d_i \log n} \cdot d_i |A_i| = \frac{k}{16 \log n} |A_i|. \end{aligned}$$

⁴ Note that without the extra $\log n$ the rightmost term could be bigger than 1, rendering the inequality trivially false.

Next, we need to show that Y is concentrated around its expectation. A simple Chernoff bound does not suffice: the Y_v are not independent, nor negatively correlated. Instead, we use the method of bounded differences (see Theorem 13). In particular, we apply this theorem with $R_i = X_u$, $f(\{X_u\}_{u \in A_i}) = Y$, so $\mu = \mathbb{E}[Y]$ and $b := \max |f(x) - f(x')| \leq 2d_i$. With $\lambda = \frac{k}{32 \log n} |A_i|$ we get

$$\begin{aligned} \mathbb{P} \left[Y < \frac{k}{32 \log n} |A_i| \right] &= \mathbb{P} \left[Y < \frac{k}{16 \log n} |A_i| - \frac{k}{32 \log n} |A_i| \right] \\ &\leq \exp \left(- \frac{k |A_i|}{32^2 (2 + 1/24) \cdot d_i \log^2 n} \right) \leq \exp \left(- \frac{k |A_i|}{2091 \cdot d_i \log^2 n} \right). \end{aligned}$$

Using that $d_i \leq 16|A_i|$, and that $k > \log^3 n$, we get with high probability that

$$|\partial I_{t+1} \setminus \partial I_t| \geq \frac{k}{32 \log n} |A_i|.$$

Case ii)

It holds that

$$\sum_{\substack{v \in V \text{ s.t.} \\ h_i(v) \geq d_i \log n / 96k}} h_i(v) \geq \frac{1}{3} \text{Vol}(A_i). \quad (3)$$

Intuitively, we want to formalize the fact that the number of informed nodes in A_i must grow even if the nodes are not likely to pull directly from I_t . We do this by looking at a subset A'_i of the nodes in A_i that have many connections. More formally we have the following fact, for a proof we refer to [34].

▷ **Claim 10** (Section 3.4 in [34]). Let $\ell = d_i \log n / 96k$. Then there exist sets $A'_i \subseteq A_i$ and $V' \subseteq V$ such that each node in A'_i has at least $d_i/8$ neighbors in V' and each node in V' has at least $\ell/8$ neighbors in A'_i , and the size of A'_i is at least $|A'_i| \geq |A_i|/20$.

We will show that each node of $s \in A'_i$ gets informed in $O(\log_k n)$ rounds with high probability. By Lemma 15, this is the same as a rumor from s reaching I_t after $O(\log_k n)$ rounds. First, we show that a rumor started at $s \in B$ reaches $\ell/1024$ nodes in A'_i in $O(\log_k n)$ rounds w.h.p. Note that $\ell/1024 \geq 1$ since $d_i \geq 1024 \cdot 96k$. Then the probability that one of these nodes pushes the rumor to I_t in the next $c \cdot 2048 \cdot 96$ rounds is at least

$$1 - \left(1 - \frac{1}{2d_i} \right)^{k \frac{\ell}{1024} c \cdot 2048 \cdot 96} = 1 - \left(1 - \frac{1}{2d_i} \right)^{c \cdot 2d_i \log n} \geq 1 - n^{-c},$$

where the equality follows from the definition of ℓ .

To show that a rumor from $s \in A'_i$ reaches $\ell/1024$ nodes in A'_i in $O(\log_k n)$ rounds, we show that in $\log_k n$ phases of $O(1)$ rounds, the number of informed nodes in A'_i grows by a factor k . We do this in two steps. First we show that by push A'_i informs nodes in V' , and then by pull nodes from A'_i pull the rumor from V' .

Let m denote the number of informed nodes of A'_i at the start of the phase. We show that by $O(1)$ push rounds we have at least $\min\{km, \ell/16\}$ informed nodes in V' . Suppose we have less than $\min\{km, \ell/16\}$ informed nodes V'_i in V' after i pushes. Then the probability of a successful push from $u \in A'_i$ is at least

$$\begin{aligned} \frac{|(N(u) \cap V') \setminus V'_i|}{\deg(u)} &\geq \frac{|N(u) \cap V'| - |V'_i|}{\deg(u)} \geq \frac{d_i/8 - \min\{km, \ell/16\}}{2d_i} \\ &\geq \frac{d_i/8 - d_i/16}{2d_i} \geq \frac{1}{32}. \end{aligned}$$

So the expected number of informed nodes in V' after 1 round is at least $\min\{\frac{mk}{32}, \ell/16\}$. Since these pushes are negatively correlated, Chernoff, Theorem 12, gives that we have with high probability that the number of informed nodes in V' after 1 round is at least

$$\min\left\{\frac{mk}{64}, \ell/16\right\}.$$

Now we consider pulling from V' . Suppose m' nodes V'_{inf} in V' are informed, now the expected number of informed nodes in A'_i after 96 pull rounds is at least

$$\begin{aligned} \sum_{u \in A'_i} 1 - \left(1 - \frac{|N(u) \cap V'_{\text{inf}}|}{\deg(u)}\right)^{96k} &= \sum_{u \in A'_i} 1 - \left(1 - \frac{|N(u) \cap V'_{\text{inf}}|}{\deg(u)}\right)^{d_i \log n / \ell} \\ &\geq \sum_{u \in A'_i} 1 - \left(1 - \frac{|N(u) \cap V'_{\text{inf}}|}{2d_i}\right)^{\frac{d_i}{\ell}} \geq \sum_{u \in A'_i} \frac{d_i}{\ell} \frac{|N(u) \cap V'_{\text{inf}}|}{4d_i} \\ &= \frac{1}{4\ell} \sum_{v \in V'_{\text{inf}}} |N(v) \cap A'_i| \geq \frac{1}{4\ell} \sum_{v \in V'_{\text{inf}}} \frac{\ell}{8} = \frac{|V'_{\text{inf}}|}{32}, \end{aligned}$$

where the second inequality uses that $|N(u) \cap V'_{\text{inf}}| \leq |V'_{\text{inf}}| \leq \ell/16$. So we conclude that in expectation, we have at least $\min\{\frac{mk}{32 \cdot 64}, \frac{\ell}{32 \cdot 16}\}$ informed nodes in A'_i . Since all pulls are independent, Chernoff, Theorem 12, gives us that with high probability we have at least

$$\min\left\{\frac{mk}{4096}, \frac{\ell}{1024}\right\}.$$

informed nodes in A'_i .

We conclude that in $r = O(\log_k n)$ rounds we have

$$|I_{t+r} \setminus I_t| \geq |A'_i| \geq \frac{|A_i|}{20},$$

where the last inequality comes from Claim 10.

Case iii)

None of the previous conditions holds, i.e., we have that

$$\sum_{\substack{v \in S_t \text{ s.t.} \\ h_i(v) \leq d_i \log n / k}} h_i(v) < \frac{1}{2} \text{Vol}(A_i) \quad \text{and} \quad \sum_{\substack{v \in V \text{ s.t.} \\ h_i(v) \geq d_i \log n / 96k}} h_i(v) < \frac{1}{3} \text{Vol}(A_i).$$

By using the two above conditions on the volume of A_i , we can see that the number of edges going from A_i to some $v \in V \setminus S_t$ with $h_i(v) < d_i/96k$ is at least a constant fraction of the volume of A_i , namely

$$\begin{aligned} \sum_{\substack{v \in V \setminus S_t \text{ s.t.} \\ h_i(v) < d_i \log n / 96k}} h_i(v) &= \sum_{v \in V} h_i(v) - \sum_{\substack{v \in V \text{ s.t.} \\ h_i(v) \geq d_i \log n / 96k}} h_i(v) - \sum_{\substack{v \in S_t \text{ s.t.} \\ h_i(v) < d_i \log n / 96k}} h_i(v) \\ &> \text{Vol}(A_i) - \sum_{\substack{v \in V \text{ s.t.} \\ h_i(v) \geq d_i \log n / 96k}} h_i(v) - \sum_{\substack{v \in S_t \text{ s.t.} \\ h_i(v) \leq d_i \log n / k}} h_i(v) \\ &\geq \frac{1}{6} \text{Vol}(A_i). \end{aligned}$$

Moreover, the number of such edges going to ∂I_t is at most $\frac{d_i \log n}{96k} |\partial I_t|$, since $h_i(v) < d_i \log n / 96k$. Therefore, using the condition on $|A_i|$ in Equation (1), that $\deg(u) \geq d_i$ for every $u \in A_i$, and that $k > \log^2 n$, we get

$$|E(A_i, I_t)| \geq \frac{1}{6} \text{Vol}(A_i) - \frac{d_i \log n}{96k} |\partial I_t| \geq \frac{d_i |A_i|}{6} - \frac{d_i \log n}{96k} (4 \log n \cdot |A_i|) \geq \frac{d_i |A_i|}{8}. \quad (4)$$

Let $B \subseteq A_i$ be the set of nodes in A_i that have at least $d_i/16$ neighbors in I_t . Then, using that $\deg(u) < 2d_i$ for every $u \in A_i$, we get

$$|E(A_i, I_t)| \leq |B| \cdot 2d_i + (|A_i| - |B|) \frac{d_i}{16}. \quad (5)$$

By combining Equations (4) and (5) it follows that $|B| \geq \frac{|A_i|}{32}$.

Note that each node $u \in B$ in the set has a fraction of at least $\frac{d_i/16}{\deg(u)} \geq \frac{d_i/16}{2d_i} = 1/32$ neighbors in I_t . We can then apply Lemma 9 and get that with high probability at least $|B|/4 \geq |A_i|/128$ nodes in A_i pull the rumor from I_t in at most $\lceil 1/\lceil k/32 \rceil \rceil = 1$ round, since $k > \log^2 n$. Hence with high probability it follows that

$$|I_{t+1} \setminus I_t| \geq \frac{|A_i|}{128}.$$

3.3 Sets with high degree nodes

Recall that in this case every $u \in A_i$ is such that

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

As in the previous case, also here nodes can either contribute directly to the growth of I_t or to that of the boundary ∂I_t . Recall that $S_t := V \setminus (I_t \cup \partial I_t)$. Let us denote the set of nodes in A_i that have more neighbors toward S_t than I_t as

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

We distinguish three cases:

- i) $|B| \geq \frac{1}{2}|A_i|$, and $\deg(u) \geq \sqrt{k}|B|$.
- ii) $|B| \geq \frac{1}{2}|A_i|$, $\deg(u) < \sqrt{k}|B|$.
- iii) $|B| < \frac{1}{2}|A_i|$.

As before in the remainder of this subsection we prove separately for each of these cases that either I_t or ∂I_t grows sufficiently. In particular in cases i) and ii) the boundary grows sufficiently, while in case iii) it is I_t to grow.

Case i)

It holds that

$$|B| \geq \frac{1}{2}|A_i| \quad \text{and} \quad \deg(u) \geq \sqrt{k}|B|.$$

In this case, we start by showing that after $\lceil \frac{|I_t| + |\partial I_t|}{|B| \log n} \rceil$ rounds, we inform at least one node $u \in B$ with high probability.

The probability that a fixed node $v \in I_t$ does not push to any neighbor in B is $1 - \frac{h_i(v)}{\deg(v)} \leq 1 - \frac{h_i(v)}{|I_t| + |\partial I_t|}$. Hence, the probability that no $v \in I_t$ pushes to B within $\frac{|I_t| + |\partial I_t|}{|B| \log n}$ rounds of k -PUSH&PULL is upper bounded by

$$\prod_{v \in I_t} \left(1 - \frac{h_i(v)}{|I_t| + |\partial I_t|} \right)^{k \frac{|I_t| + |\partial I_t|}{|B| \log n}} \leq e^{-k/\log n \cdot \frac{1}{|B|} \sum_{v \in I_t} h_i(v)} \leq e^{-k/\log n}.$$

This uses that $\sum_{v \in I_t} h_i(v) = \sum_{v \in A_i} |N(u) \cap I_t| \geq |A_i| \geq |B|$, since $A_i \subseteq \partial I_t$.

Hence, the probability that push informs at least one node in these rounds is at least $1 - e^{-k/\log n}$, namely with high probability, since $k > \log^2 n$.

Therefore, since $\deg(u) \geq \sqrt{k}|B|$, and by definition of the set B , we get with high probability that

$$|\partial I_{t+r} \setminus \partial I_t| \geq \frac{\sqrt{k}|B|}{2} \geq \frac{\sqrt{k}|A_i|}{4} \geq \frac{\sqrt{k}|\partial I_t|}{16 \log n}$$

for a number of rounds $r = \lceil \frac{|I_t| + |\partial I_t|}{|B| \log n} \rceil = O(1)$ since $|B| \geq \frac{1}{2}|A_i| \geq \frac{|\partial I_t|}{8 \log n}$, where $|A_i| \geq \frac{|\partial I_t|}{4 \log n}$ since we consider $i \in \mathcal{I}$.

Case ii)

It holds that

$$|B| \geq \frac{1}{2}|A_i| \quad \text{and} \quad \deg(u) < \sqrt{k}|B|.$$

We start by noting that the probability that a fixed node $u \in A_i$ pulls the rumor from I_t in a given round is

$$1 - \left(1 - \frac{|N(u) \cap I_t|}{\deg(u)} \right)^k \geq 1 - \left(1 - \frac{1}{\deg(u)} \right)^k \geq 1 - \left(1 - \frac{1}{2d_i} \right)^k \geq \frac{k}{4d_i} =: p,$$

where we use in the last inequality that $2d_i > k$.

Formally, let X_u be the 0/1 random variable that is 1 iff u pulls the rumor. Note that the X_u 's are independent since they come from pull operations. Let us pessimistically assume that such a probability exactly equals p for every node u . For each node $v \in N(B) \cap S_t$ let Y_v be the 0/1 random variable that is 1 iff v has at least a neighbor u with $X_u = 1$.

We look at $v \in N(B) \cap S_t$ and consider the probability that no neighbor of v in B pulls the rumor, which is at least

$$1 - \prod_{u \in B \cap N(v)} \left(1 - \frac{k}{4d_i} \right) \geq 1 - \exp \left(-k/4 \sum_{u \in B \cap N(v)} \frac{1}{d_i} \right) \geq \frac{k}{8} \sum_{u \in B \cap N(v)} \frac{1}{d_i},$$

where the last inequality follows from the assumption on nodes $u \in B$. The expected number of nodes $v \in N(B) \cap S_t$ that join ∂I_t is now at least

$$\sum_{v \in N(B) \cap S_t} \frac{k}{8} \sum_{u \in B \cap N(v)} \frac{1}{d_i} = \frac{k}{8} \sum_{u \in B} \frac{|N(u) \cap S_t|}{d_i} \geq \frac{k}{16}|B|,$$

where the second to last equality follows by definition of B and the fact that $d_i > 2|\partial I_t|$.

We now show that the expected number of new nodes in the boundary is not far from that lower bound. Formally we lower bound $Y := \sum_{v \in N(B) \cap S_t} Y_v$. Hereto, we use the method of bounded differences (Theorem 13) with $R_i = X_u$, $f(\{X_u\}_{u \in B}) = Y$, and $b \leq 2d_i$. Since $\mathbb{E}[Y] \geq \frac{k}{16}|B|$, with $\lambda = \frac{k}{32}|B|$ we get

$$\mathbb{P}\left(Y < \frac{k}{32}|B|\right) \leq \exp\left(-\frac{k|B|}{2091 \cdot d_i}\right).$$

Since $\deg(u) < 2d_i < 2\sqrt{k}|B|$ and $k > \log^2 n$, the previous bound holds with high probability. Therefore with high probability we have

$$|\partial I_{t+1} \setminus \partial I_t| \geq \frac{k}{64}|A_i|.$$

Case iii)

It holds that

$$|B| < \frac{1}{2}|A_i|.$$

It follows directly from the above condition that $|A_i \setminus B| > |A_i|/2$ and by the definition of B that each node $u \in A_i \setminus B$ iff $|N(u) \cap S_t| < k|N(u) \cap I_t|$. Recall that we are in the case of high degree nodes: $|N(u)| \geq 2|\partial I_t|$, equivalently, we have $|\partial I_t| \leq \frac{1}{2}|N(u)|$. Hence $|N(u) \cap \partial I_t| \leq |\partial I_t| \leq \frac{1}{2}|N(u)|$, and $|N(u) \cap I_t| + |N(u) \cap S_t| \geq \frac{1}{2}|N(u)|$, so $|N(u) \cap \partial I_t| \leq |N(u) \cap I_t| + |N(u) \cap S_t|$. Combining this, we obtain

$$\begin{aligned} \frac{|N(u) \cap I_t|}{|N(u)|} &= \frac{|N(u) \cap I_t|}{|N(u) \cap I_t| + |N(u) \cap S_t| + |N(u) \cap \partial I_t|} \\ &\geq \frac{|N(u) \cap I_t|}{2(|N(u) \cap I_t| + |N(u) \cap S_t|)} > \frac{|N(u) \cap I_t|}{2(k+1)|N(u) \cap I_t|} = \frac{1}{2(k+1)}. \end{aligned}$$

We conclude that u has a $\frac{1}{2(k+1)}$ -fraction of all its neighbors in I_t . Then we apply Lemma 9 and get that with high probability at least one fourth of the nodes in $|A_i \setminus B|$ pull the rumor from I_t in at most $\lceil \frac{2(k+1)}{k} \rceil = O(1)$ rounds. Hence we get with high probability that

$$|I_{t+r} \setminus I_t| \geq \frac{|A_i \setminus B|}{4} \geq \frac{|A_i|}{8}$$

for a number of rounds $r = O(1)$.

References

- 1 Hüseyin Acan, Andrea Collevecchio, Abbas Mehrabian, and Nick Wormald. On the push&pull protocol for rumour spreading: [extended abstract]. In *Proceedings of the 2015 ACM Symposium on Principles of Distributed Computing (PODC 2015)*, pages 405–412. ACM, 2015. doi:10.1145/2767386.2767416.
- 2 Noga Alon, Chen Avin, Michal Koucký, Gady Kozma, Zvi Lotker, and Mark R. Tuttle. Many random walks are faster than one. *Combinatorics, Probability & Computing*, 20(4):481–502, 2011. Announced at SPAA'08. doi:10.1017/S0963548311000125.
- 3 Sanjeev Arora, Boaz Barak, and David Steurer. Subexponential algorithms for unique games and related problems. *Journal of the ACM*, 62(5):42:1–42:25, 2015. Announced at FOCS 2010. doi:10.1145/2775105.
- 4 Sanjeev Arora, Frank Thomson Leighton, and Bruce M. Maggs. On-line algorithms for path selection in a nonblocking network. *SIAM J. Comput.*, 25(3):600–625, 1996. Announced at STOC '90. doi:10.1137/S0097539791221499.

- 5 Richard Arratia and Louis Gordon. Tutorial on large deviations for the binomial distribution. *Bulletin of mathematical biology*, 51(1):125–131, 1989.
- 6 Chen Avin and Robert Elsässer. Breaking the $\log n$ barrier on rumor spreading. *Distributed Computing*, 31(6):503–513, 2018. Announced at DISC 2013. doi:10.1007/S00446-017-0312-4.
- 7 Petra Berenbrink, Robert Elsässer, and Tom Friedetzky. Efficient randomised broadcasting in random regular networks with applications in peer-to-peer systems. *Distributed Comput.*, 29(5):317–339, 2016. Announced at PODC 2008. doi:10.1007/S00446-016-0264-0.
- 8 Keren Censor-Hillel, Bernhard Haeupler, Jonathan A. Kelner, and Petar Maymounkov. Rumor spreading with no dependence on conductance. *SIAM Journal on Computing*, 46(1):58–79, 2017. Announced at STOC 2012. doi:10.1137/14099992X.
- 9 Keren Censor-Hillel and Hadas Shachnai. Partial information spreading with application to distributed maximum coverage. In *Proceedings of the 29th Annual ACM Symposium on Principles of Distributed Computing (PODC 2010)*, pages 161–170. ACM, 2010. doi:10.1145/1835698.1835739.
- 10 Keren Censor-Hillel and Hadas Shachnai. Fast information spreading in graphs with large weak conductance. *SIAM Journal on Computing*, 41(6):1451–1465, 2012. Announced at SODA 2011. doi:10.1137/110845380.
- 11 Flavio Chierichetti, George Giakkoupis, Silvio Lattanzi, and Alessandro Panconesi. Rumor spreading and conductance. *Journal of the ACM*, 65(4):17:1–17:21, 2018. doi:10.1145/3173043.
- 12 Flavio Chierichetti, Silvio Lattanzi, and Alessandro Panconesi. Almost tight bounds for rumour spreading with conductance. In *Proceedings of the 42nd ACM Symposium on Theory of Computing (STOC 2010)*, pages 399–408. ACM, 2010. doi:10.1145/1806689.1806745.
- 13 Flavio Chierichetti, Silvio Lattanzi, and Alessandro Panconesi. Rumor spreading in social networks. *Theoretical Computer Science*, 412(24):2602–2610, 2011. Announced at ICALP 2009. doi:10.1016/J.TCS.2010.11.001.
- 14 Eden Chlamtác, Michael Dinitz, and Yury Makarychev. Minimizing the union: Tight approximations for small set bipartite vertex expansion. In *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2017)*, pages 881–899. SIAM, 2017. doi:10.1137/1.9781611974782.56.
- 15 Fan Chung and Linyuan Lu. The diameter of sparse random graphs. *Advances in Applied Mathematics*, 26(4):257–279, 2001. doi:10.1006/AAMA.2001.0720.
- 16 Julia Chuzhoy, Yury Makarychev, Aravindan Vijayaraghavan, and Yuan Zhou. Approximation algorithms and hardness of the k -route cut problem. *ACM Transactions on Algorithms*, 12(1):2:1–2:40, 2016. Announced at SODA 2012. doi:10.1145/2644814.
- 17 Sebastian Daum, Fabian Kuhn, and Yannic Maus. Rumor spreading with bounded in-degree. *Theoretical Computer Science*, 810:43–57, 2020. Announced at SIROCCO 2016. doi:10.1016/J.TCS.2018.05.041.
- 18 Andreea Deac, Marc Lackenby, and Petar Veličković. Expander graph propagation. In *Learning on Graphs Conference*, pages 38–1. PMLR, 2022.
- 19 Alan J. Demers, Daniel H. Greene, Carl Hauser, Wes Irish, John Larson, Scott Shenker, Howard E. Sturgis, Daniel C. Swinehart, and Douglas B. Terry. Epidemic algorithms for replicated database maintenance. *ACM SIGOPS Oper. Syst. Rev.*, 22(1):8–32, 1988. Announced at PODC 1987. doi:10.1145/43921.43922.
- 20 Benjamin Doerr and Mahmoud Fouz. Asymptotically optimal randomized rumor spreading. *Electronic Notes in Discrete Mathematics*, 38:297–302, 2011. Announced at ICALP 2011. doi:10.1016/J.ENDM.2011.09.049.
- 21 Benjamin Doerr, Mahmoud Fouz, and Tobias Friedrich. Social networks spread rumors in sublogarithmic time. *Electronic Notes in Discrete Mathematics*, 38:303–308, 2011. Announced at STOC 2011. doi:10.1016/J.ENDM.2011.09.050.

- 22 Benjamin Doerr, Mahmoud Fouz, and Tobias Friedrich. Asynchronous rumor spreading in preferential attachment graphs. In *Proceedings of the 13th Scandinavian Symposium and Workshops on Algorithm Theory (SWAT 2012)*, volume 7357 of *LNCS*, pages 307–315. Springer, 2012. doi:10.1007/978-3-642-31155-0_27.
- 23 Benjamin Doerr, Tobias Friedrich, and Thomas Sauerwald. Quasirandom rumor spreading. *ACM Transactions on Algorithms*, 11(2):9:1–9:35, 2014. Announced at SODA 2008. doi:10.1145/2650185.
- 24 Benjamin Doerr and Anatolii Kostygin. Randomized rumor spreading revisited. In *Proceedings of the 44th International Colloquium on Automata, Languages, and Programming (ICALP 2017)*, volume 80 of *LIPICs*, pages 138:1–138:14. LZI, 2017. doi:10.4230/LIPICs.ICALP.2017.138.
- 25 Benjamin Doerr and Marvin Künnemann. Tight analysis of randomized rumor spreading in complete graphs. In *Proceedings of the Eleventh Workshop on Analytic Algorithmics and Combinatorics (ANALCO 2014)*, pages 82–91. SIAM, 2014. doi:10.1137/1.9781611973204.8.
- 26 Devdatt P. Dubhashi and Alessandro Panconesi. *Concentration of measure for the analysis of randomized algorithms*. Cambridge University Press, 2009.
- 27 Uriel Feige, David Peleg, Prabhakar Raghavan, and Eli Upfal. Randomized broadcast in networks. *Random Structures & Algorithms*, 1(4):447–460, 1990. Announced at SIGAL 1990. doi:10.1002/RSA.3240010406.
- 28 Nikolaos Fountoulakis and Konstantinos Panagiotou. Rumor spreading on random regular graphs and expanders. *Random Structures & Algorithms*, 43(2):201–220, 2013. Announced at APPROX-RANDOM 2010. doi:10.1002/RSA.20432.
- 29 Nikolaos Fountoulakis, Konstantinos Panagiotou, and Thomas Sauerwald. Ultra-fast rumor spreading in social networks. In *Proceedings of the Twenty-Third Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2012)*, pages 1642–1660. SIAM, 2012. doi:10.1137/1.9781611973099.130.
- 30 Alan M. Frieze and Geoffrey R. Grimmett. The shortest-path problem for graphs with random arc-lengths. *Discrete Applied Mathematics*, 10(1):57–77, 1985. doi:10.1016/0166-218X(85)90059-9.
- 31 Mohsen Ghaffari and Calvin C. Newport. How to discreetly spread a rumor in a crowd. In *Proceedings of the 30th International Symposium on Distributed Computing (DISC 2016)*, volume 9888 of *LNCS*, pages 357–370. Springer, 2016. doi:10.1007/978-3-662-53426-7_26.
- 32 George Giakkoupis. Tight bounds for rumor spreading with vertex expansion. In *Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2014)*, pages 801–815. SIAM, 2014. doi:10.1137/1.9781611973402.59.
- 33 George Giakkoupis, Yasamin Nazari, and Philipp Woelfel. How asynchrony affects rumor spreading time. In *Proceedings of the 2016 ACM Symposium on Principles of Distributed Computing (PODC 2016)*, pages 185–194. ACM, 2016. doi:10.1145/2933057.2933117.
- 34 George Giakkoupis and Thomas Sauerwald. Rumor spreading and vertex expansion. In *Proceedings of the Twenty-Third Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2012)*, pages 1623–1641. SIAM, 2012. doi:10.1137/1.9781611973099.129.
- 35 Bernhard Haeupler. Simple, fast and deterministic gossip and rumor spreading. *Journal of the ACM*, 62(6):47:1–47:18, 2015. Announced at SODA 2013. doi:10.1145/2767126.
- 36 Bernhard Haeupler and Dahlia Malkhi. Optimal gossip with direct addressing. In *Proceedings of the 2014 ACM symposium on Principles of distributed computing*, pages 176–185, 2014. doi:10.1145/2611462.2611489.
- 37 Jun-Ting Hsieh, Alexander Lubotzky, Sidhanth Mohanty, Assaf Reiner, and Rachel Yun Zhang. Explicit lossless vertex expanders. *arXiv preprint arXiv:2504.15087*, 2025. doi:10.48550/arXiv.2504.15087.
- 38 Richard M. Karp, Christian Schindelhauer, Scott Shenker, and Berthold Vöcking. Randomized rumor spreading. In *Proceedings of the 41st Annual Symposium on Foundations of Computer*

- Science (FOCS 2000)*, pages 565–574. IEEE Computer Society, 2000. doi:10.1109/SFCS.2000.892324.
- 39 Anand Louis and Yury Makarychev. Approximation algorithms for hypergraph small-set expansion and small-set vertex expansion. *Theory of Computing*, 12(1):1–25, 2016. Announced at APPROX-RANDOM 2014. doi:10.4086/TOC.2016.V012A017.
 - 40 Colin McDiarmid. Concentration. In *Probabilistic methods for algorithmic discrete mathematics*, pages 195–248. Springer, 1998.
 - 41 Abbas Mehrabian and Ali Pourmiri. Randomized rumor spreading in poorly connected small-world networks. *Random Structures & Algorithms*, 49(1):185–208, 2016. Announced at DISC 2014. doi:10.1002/RSA.20624.
 - 42 Yves Mocquard, Bruno Sericola, and Emmanuelle Anceaume. Analysis of rumor spreading with 2-pull or 3-pull operations. In *Proceedings of the 20th IEEE International Symposium on Network Computing and Applications (NCA 2021)*, pages 1–8. IEEE, 2021. doi:10.1109/NCA53618.2021.9685439.
 - 43 Guy Moshkovitz and Asaf Shapira. Decomposing a graph into expanding subgraphs. *Random Structures & Algorithms*, 52(1):158–178, 2018. Announced at SODA 2015. doi:10.1002/RSA.20727.
 - 44 Damon Mosk-Aoyama and Devavrat Shah. Fast distributed algorithms for computing separable functions. *IEEE Transactions on Information Theory*, 54(7):2997–3007, 2008. Announced at PODC 2006. doi:10.1109/TIT.2008.924648.
 - 45 Charlotte Out, Nicolás Rivera, Thomas Sauerwald, and John Sylvester. Rumors with changing credibility. In Venkatesan Guruswami, editor, *15th Innovations in Theoretical Computer Science Conference, ITCS 2024, January 30 to February 2, 2024, Berkeley, CA, USA*, volume 287 of *LIPICs*, pages 86:1–86:23. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2024. doi:10.4230/LIPICs.ITCS.2024.86.
 - 46 Konstantinos Panagiotou, Xavier Pérez-Giménez, Thomas Sauerwald, and He Sun. Randomized rumour spreading: The effect of the network topology. *Combinatorics, Probability & Computing*, 24(2):457–479, 2015. doi:10.1017/S0963548314000194.
 - 47 Konstantinos Panagiotou, Ali Pourmiri, and Thomas Sauerwald. Faster rumor spreading with multiple calls. *The Electronic Journal of Combinatorics*, 22(1):1, 2015. Announced at ISAAC 2013. doi:10.37236/4314.
 - 48 Boris Pittel. On spreading a rumor. *SIAM Journal on Applied Mathematics*, 47(1):213–223, 1987. doi:10.1137/0147013.
 - 49 Prasad Raghavendra and David Steurer. Graph expansion and the unique games conjecture. In *Proceedings of the 42nd ACM Symposium on Theory of Computing (STOC 2010)*, pages 755–764. ACM, 2010. doi:10.1145/1806689.1806792.
 - 50 Frédérique Robin, Bruno Sericola, Emmanuelle Anceaume, and Yves Mocquard. Stochastic analysis of rumor spreading with multiple pull operations. *Methodology and Computing in Applied Probability*, 24(3):2195–2211, 2022.
 - 51 Naser T Sardari. Diameter of ramanujan graphs and random cayley graphs. *Combinatorica*, 39(2):427–446, 2019. doi:10.1007/S00493-017-3605-0.
 - 52 Thomas Sauerwald and Alexandre Stauffer. Rumor spreading and vertex expansion on regular graphs. In *Proceedings of the Twenty-Second Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2011)*, pages 462–475. SIAM, 2011. doi:10.1137/1.9781611973082.37.
 - 53 Hamed Shirzad, Ameya Velingker, Balaji Venkatachalam, Danica J. Sutherland, and Ali Kemal Sinop. Exphormer: Sparse transformers for graphs. In *Proceedings of the International Conference on Machine Learning (ICML 2023)*, volume 202 of *Proceedings of Machine Learning Research*, pages 31613–31632. PMLR, 2023. arXiv:2303.06147.
 - 54 Daniel A. Spielman and Shang-Hua Teng. Nearly-linear time algorithms for graph partitioning, graph sparsification, and solving linear systems. In *Proceedings of the 36th Annual ACM Symposium on Theory of Computing (STOC 2004)*, pages 81–90. ACM, 2004. doi:10.1145/1007352.1007372.

A Concentration Bounds

► **Definition 11** (Negatively Associated Random Variables). *The random variables $X_1, \dots, X_n$ are said to be negatively associated if for all disjoint subsets $I, J \subseteq \{1, \dots, n\}$ and all non-decreasing functions f and g it holds that*

$$\mathbb{E}[f(X_i, i \in I) \cdot g(X_j, j \in J)] \leq \mathbb{E}[f(X_i, i \in I)] \cdot \mathbb{E}[g(X_j, j \in J)].$$

► **Theorem 12** (Chernoff Bounds [26]). *Let $X_1, \dots, X_n$ be independent 0/1 random variables. Let $X := \sum_{i=1}^n X_i$ and let $\mu_L \leq \mathbb{E}[X] \leq \mu_H$. It holds that*

$$\mathbb{P}[X < (1 - \delta)\mu_L] \leq e^{-\frac{\delta^2}{2}\mu_L}, \quad \forall \delta \in (0, 1);$$

$$\mathbb{P}[X > (1 + \delta)\mu_H] \leq e^{-\frac{\delta^2}{2+\delta}\mu_H}, \quad \forall \delta > 0.$$

The bounds also hold if $X_1, \dots, X_n$ are negatively associated.

► **Theorem 13** (Bounded Differences Inequality [40]). *Let $R_1, \dots, R_n$ be independent 0/1 random variables with $\mathbb{P}[R_i = 1] \leq p \leq 1/2$. Let f be a bounded real function defined on $\{0, 1\}^n$. Define $\mu := \mathbb{E}[f(R_1, \dots, R_n)]$, and $b := \max |f(x) - f(x')|$, where the maximum is over all $x, x' \in \{0, 1\}^n$ that differ only in one position. Then for any $\lambda > 0$*

$$\mathbb{P}[f(R_1, \dots, R_n) \leq \mu - \lambda] \leq \exp\left(-\frac{\lambda^2}{2pnb^2 + 2b\lambda/3}\right).$$

► **Theorem 14** ([5]). *Let $B(n, p)$ denote a random variable following the binomial distribution with n trials and success probability p . Then, for $p < a < 1$, we have that*

$$\mathbb{P}[B(n, p) \geq an] \leq \exp\left(-n \left(a \log\left(\frac{a}{p}\right) + (1 - a) \log\left(\frac{1 - a}{1 - p}\right)\right)\right).$$

B Warm-Up: Multi-Call Rumor Spreading on Complete Graphs

This section serves as a warm-up for our main result; we investigate what impact parallel calls have on rumor spreading on complete graphs. We provide a full characterization with matching upper and lower bounds. For $k = 1$, we know that rumor spreading takes $\Theta(\log n)$ rounds [30, 38], so we only consider $k \geq 2$. The analysis in [24] yields precise results that include the leading constant (in fact tight apart from an additive number of rounds) for the case $k = O(1)$. However, it does not extend to super-constant values of k , where it leads to an upper bound of $O(k \cdot \log_k n)$, which is asymptotically worse than ours.

The proofs in the section are generalizations of standard arguments for rumor spreading on cliques. Independent work by Out, Rivera, Sauerwald, and Sylvester [45] considers rumor spreading with a time-dependent credibility function. Some of their proofs resemble the results in this section.

In the proof for the upper bound, we use the following standard lemma, which is a trivial generalization of the version with $k = 1$ (see, e.g., [12, Lemma 3], [34, Lemma 3.3], or [8, Lemma 2.1]).

► **Lemma 15.** *For $S, T \subseteq V$, let $T_{k\text{-PUSH}\&\text{PULL}}(S, T)$ be the number of rounds for $k\text{-PUSH}\&\text{PULL}$ until a rumor that is initially known to all nodes in S to spread to at least one node of $v \in T$. Let $V_1, V_2 \subseteq V$. Then the random variables $T_{k\text{-PUSH}\&\text{PULL}}(V_1, V_2)$ and $T_{k\text{-PUSH}\&\text{PULL}}(V_2, V_1)$ have the same distribution.*

► **Lemma 16.** *Rumor spreading in a complete graph with the k -PUSH&PULL protocol for every $k \geq 2$ requires $O(\log_k n)$ rounds with high probability.*

Proof. An upper bound of $O(\log n)$ has been known for a long time for $k = 1$ [30, 38]. Since parallel spreading can only make the process faster, we have this upper-bound for any k . In particular, for any constant k we have that the k -PUSH&PULL protocol requires $O(\log n) = O(\log_k n)$ rounds with high probability. In the following we assume $k \geq 17$.

Let I_t denote the informed nodes at round t . First we show using only k -PUSH that $|I_t| \geq (k/4)^t$ until $k|I_t| > n - 1$. Assume $k|I_t| \leq n - 1$. Each node receives a message from at least one node in I_t with probability $1 - (1 - \frac{1}{n-1})^{k|I_t|} = \frac{k|I_t|}{n-1} - \frac{1}{2}k|I_t|(k|I_t| - 1)\frac{1}{(n-1)^2} + \dots \geq \frac{k|I_t|}{2(n-1)}$ where we use binomial expansion. Now we see that

$$\mathbb{E}[|I_{t+1} \setminus I_t|] \geq (n - |I_t|) \frac{k|I_t|}{2(n-1)} \geq \left(n - \frac{n-1}{k}\right) \frac{k|I_t|}{2(n-1)} \geq \frac{k-1}{2}|I_t|.$$

$|I_{t+1} \setminus I_t|$ can be written as a sum of 0/1 random variables. Note that they are not independent, but they are negatively associated since the probability that u receives a push can only decrease under the assumption that v receives a push. Hence, as discussed in the appendix (see Theorem 12), Chernoff still gives

$$\mathbb{P}[|I_{t+1} \setminus I_t| < \frac{k-1}{4}|I_t|] \leq e^{-(k-1)|I_t|/16}.$$

For $k \geq 16(c+1)\log n$, this is bounded by $n^{-(c+1)}$, which gives the required probability. For smaller k , we need to analyze the process more carefully.

We consider $\beta \cdot 2\log_k n$ rounds, for some constant $\beta = \beta(c) \geq 2$ to be decided later. By a Chernoff bound, the probability that any such round does *not* have the required growth is bounded by $e^{-(k-1)/16}$. We compute the probability that more than $(\beta - 1) \cdot 2\log_k n$ rounds fail in reaching the required growth. In fact, if at most $(\beta - 1) \cdot 2\log_k n$ rounds fail, we have at least $2\log_k n \geq \frac{\log(n/k)}{\log(k/4)}$ successes which is enough to reach $|I_t| > (n - 1)/k$ nodes.

To count the number of failures, we see that this is dominated by a binomial random variable $B(N, p)$ with $N = \beta \cdot 2\log_k n$ trials and probability of success $p = e^{-(k-1)/16}$. We bound the probability that $B(N, p) \geq (\beta - 1) \cdot 2\log_k n = \frac{\beta-1}{\beta}N$. By using standard tail bounds on the binomial distribution (see Theorem 14) we get that for $p < \frac{\beta-1}{\beta}$ we have

$$\mathbb{P}\left[B(N, p) \geq \frac{\beta-1}{\beta}N\right] \leq \exp\left(-N\left(\frac{\beta-1}{\beta}\log\left(\frac{\frac{\beta-1}{\beta}}{p}\right) + \left(1 - \frac{\beta-1}{\beta}\right)\log\left(\frac{1 - \frac{\beta-1}{\beta}}{1-p}\right)\right)\right).$$

Note that we have $p < \frac{\beta-1}{\beta}$, since $p = e^{-(k-1)/16} \leq e^{-1} < 1/2 \leq \frac{\beta-1}{\beta}$, where the last inequality follows from the assumption that $\beta \geq 2$.

Now we simplify this bound

$$\begin{aligned} \mathbb{P}\left[B(N, p) \geq \frac{\beta-1}{\beta}N\right] &\leq \exp\left(-N\left(\frac{\beta-1}{\beta}\log\left(\frac{\frac{\beta-1}{\beta}}{p}\right) + \left(1 - \frac{\beta-1}{\beta}\right)\log\left(\frac{1 - \frac{\beta-1}{\beta}}{1-p}\right)\right)\right) \\ &= \exp\left(-\beta 2\log_k(n)\left(\frac{\beta-1}{\beta}\log\left(\frac{\frac{\beta-1}{\beta}}{e^{-(k-1)/16}}\right) + \frac{1}{\beta}\log\left(\frac{\frac{1}{\beta}}{1 - e^{-(k-1)/16}}\right)\right)\right) \\ &\leq \exp\left(-2\log_k(n)\left((\beta-1)\left(\log\left(\frac{\beta-1}{\beta}\right) + (k-1)/16\right) - \log\beta\right)\right). \end{aligned}$$

We add the constraint that $\beta \geq 3$, so that $\log\left(\frac{\beta-1}{\beta}\right) > -1/2 \geq -(k-1)/32$, hence we have

$$\begin{aligned} \mathbb{P}\left[B(N, p) \geq \frac{\beta-1}{\beta}N\right] &\leq \exp(-2\log_k(n) \cdot (\beta-1)(k-1)/32 - \log \beta) \\ &= \exp\left(-2\log_k(n) \cdot (\beta-1) \left((k-1)/32 - \frac{\log \beta}{\beta-1}\right)\right). \end{aligned}$$

We add the further constraint that $\beta \geq 11$, so that $\frac{\log \beta}{\beta-1} < 1/4 \leq (k-1)/64$. This gives

$$\begin{aligned} \mathbb{P}\left[B(N, p) \geq \frac{\beta-1}{\beta}N\right] &\leq \exp(-2\log_k(n) \cdot (\beta-1)(k-1)/64) \\ &= \exp\left(-\log(n) \cdot \frac{\beta-1}{32} \cdot \frac{k-1}{\log k}\right) \leq n^{-(c+1)}, \end{aligned}$$

where the last inequality holds for $\beta \geq 32c + 33$. In other words, we have shown that with high probability after at most $O(\log_k n)$ rounds we have that $k|I_t| > n - 1$.

Next, we show that $|I_{t+1}| > n/2$ w.h.p., so assume $|I_t| \leq n/2$. Let $v \in V \setminus I_t$. Then the probability that v pulls from I_t is

$$1 - \left(1 - \frac{|I_t|}{n-1}\right)^k \geq 1 - \left(1 - \frac{1}{k}\right)^k \geq 1 - 1/e.$$

So the expected number of nodes that pull is at least $(n - |I_t|)(1 - 1/e)$. In this case, $|I_{t+1} \setminus I_t|$ is a standard binomial random variable (i.e., sum of independent 0/1 random variables). Hence, using Chernoff we see that

$$\begin{aligned} \mathbb{P}[|I_{t+1} \setminus I_t| \leq \frac{n-|I_t|}{2}] &= \mathbb{P}[|I_{t+1} \setminus I_t| \leq (1 - \frac{1+1/e}{2})(n - |I_t|)(1 - 1/e)] \\ &\leq e^{-\left(\frac{1+1/e}{2}\right)^2 (n-|I_t|)(1-1/e)/2} \leq e^{-n/14}, \end{aligned}$$

using that $n - |I_t| \geq n/2$. Now we have that $|I_{t+1}| > |I_t| + \frac{n-|I_t|}{2} \geq n/2$.

By a standard symmetry argument, see Lemma 15, we see that we use at most $2O(\log_k n) = O(\log_k n)$ rounds with probability $1 - n^{-c}$. $\blacktriangleleft$

Next, we show the matching lower bound.

► **Lemma 17.** *Rumor spreading in a complete graph with the k -PUSH&PULL protocol for every $k \geq 2$ requires $\Omega(\log_k n)$ rounds with high probability.*

Proof. We show that it takes $\Omega(\log_k n)$ rounds to get n/k nodes informed. First, let us assume that we are at a round t_0 such that $54(c+1)\log n \leq |I_{t_0}| \leq 216(c+1)k\log n$ for some positive constant c . We will later show that this round exists w.h.p. With such an assumption, we show by induction that $|I_{t_0+t}| \leq (3k)^t |I_{t_0}|$, until $|I_t| \geq n/k$.⁵ Clearly k -PUSH can inform at most $k|I_t|$ nodes in one round. So it remains to show that the new informed nodes due to k -PULL in one round are at most $2k|I_t|$ w.h.p.

Let $v \in V \setminus I_t$. Then the probability that v pulls from I_t equals

$$1 - \left(1 - \frac{|I_t|}{n-1}\right)^k \leq \frac{3}{2}k \frac{|I_t|}{n-1}.$$

⁵ Note that this requires that $216(c+1)k\log n < n/k \iff k < \sqrt{n/(216(c+1)\log n)}$, which we can assume since $1 = \Omega(\log_k n)$ is a trivial lower bound for such k .

So we can bound the expected gain from pull $|I_{t+1}^{(\text{pull})}|$ by

$$\mathbb{E}[|I_{t+1}^{(\text{pull})}|] \leq (n - |I_t|) \cdot \frac{3}{2} k \frac{|I_t|}{n-1} \leq \frac{3}{2} k |I_t|.$$

Note that $|I_{t+1}^{(\text{pull})}|$ is a binomial random variable, hence we can use Chernoff to bound the probability that this is bigger than $2k|I_t|$ and get

$$\mathbb{P}[|I_{t+1}^{(\text{pull})}| > 2k|I_t|] = \mathbb{P}\left[|I_{t+1}^{(\text{pull})}| > \left(1 + \frac{1}{3}\right) \frac{3}{2} k |I_t|\right] \leq e^{-\mathbb{E}[|I_{t+1}^{(\text{pull})}|]/27}.$$

To see that this probability is indeed low, we lower bound the expected increase $|I_{t+1}^{(\text{pull})}|$ as well. The probability that any $v \in V \setminus I_t$ pulls from I_t equals

$$1 - \left(1 - \frac{|I_t|}{n-1}\right)^k \geq k \frac{|I_t|}{2(n-1)},$$

since we are assuming $|I_t| \leq n/k$. So we see

$$\mathbb{E}[|I_{t+1}^{(\text{pull})}|] \geq (n - |I_t|) k \frac{|I_t|}{2(n-1)} \geq \frac{n - n/k}{2(n-1)} k |I_t| = \frac{n(k-1)}{2(n-1)} |I_t| \geq \frac{k-1}{2} |I_t| \geq \frac{|I_t|}{2}.$$

Since we have by assumption that $|I_t| \geq |I_{t_0}| \geq 54(c+1) \log n$ the result holds at each step with probability at least $1 - n^{-(c+1)}$. If $|I_{t_0+t}| \leq (3k)^t |I_{t_0}|$, then we need

$$\frac{\log\left(\frac{n/k}{|I_{t_0}|}\right)}{\log(3k)} = \Omega(\log_k n)$$

rounds to inform at least n/k nodes.

It remains to show that until we inform $54(c+1) \log n$ nodes, we cannot suddenly overshoot $216(c+1) k \log n$ nodes. We note that $I_{t+1} = I_t \cup I_{t+1}^{(\text{push})} \cup I_{t+1}^{(\text{pull})}$. It is clear that k -PUSH can inform at most $k \cdot 54(c+1) \log n$ nodes. So it remains to show that it is very unlikely that the process reaches more than $216(c+1) k \log n$ nodes through k -PULL. We do this by applying a Chernoff bound in a different way. More precisely, assume that $|I_t| \leq 54(c+1) \log n$, then we show that the probability that $|I_{t+1}^{(\text{pull})}| > 108(c+1) k \log n$ is small:

$$\begin{aligned} \mathbb{P}[|I_{t+1}^{(\text{pull})}| > 108(c+1) k \log n] &= \mathbb{P}\left[|I_{t+1}^{(\text{pull})}| > \left(1 + \left(\frac{1}{2} + \frac{3}{2}\right) \frac{(54(c+1) k \log n - 1)}{\mathbb{E}[|I_{t+1}^{(\text{pull})}|]}\right) \mathbb{E}[|I_{t+1}^{(\text{pull})}|]\right] \\ &\leq \mathbb{P}\left[|I_{t+1}^{(\text{pull})}| > \left(1 + \frac{27(c+1) k \log n}{\mathbb{E}[|I_{t+1}^{(\text{pull})}|]}\right) \mathbb{E}[|I_{t+1}^{(\text{pull})}|]\right] \leq \exp\left(-\frac{\left(\frac{27(c+1) k \log n}{\mathbb{E}[|I_{t+1}^{(\text{pull})}|]}\right)^2}{2 + \frac{27(c+1) k \log n}{\mathbb{E}[|I_{t+1}^{(\text{pull})}|]}} \cdot \mathbb{E}[|I_{t+1}^{(\text{pull})}|]\right) \\ &\leq \exp(-(27/7)(c+1) k \log n) \leq n^{-(c+1)}. \end{aligned}$$

C On Small-Set Vertex Expanders with Expansion Larger than 1

As previously discussed, when the parameter $\alpha < 1/2$ we talk about small-set expanders. These graphs are substantially different from the classical notion of expanders, where $\alpha = 1/2$. In fact, if $\alpha = 1/2$ it holds that the vertex-expansion of the graph is bounded, having $\phi \leq 1$. The following lemma formalizes this fact.

► **Lemma 18.** *Let $\phi > 0$. There are no (ϕ, α) -expanders for $\alpha > \frac{1}{1+\phi}$.*

Proof. Let S be any set of size $\frac{n}{1+\phi} < |S| \leq \alpha n$. We demand that $\phi(S) \geq \phi$, in other words:

$$|S \cup N(S)| = |S| + |\partial S| \geq |S| + \phi|S| > (1 + \phi) \frac{n}{1+\phi} = n,$$

which clearly is not possible. $\blacktriangleleft$

Now the follow-up question is: why do we not simply take $\alpha = \frac{1}{1+\phi}$? The answer is that this is a very restricted graph class as we show next. Recall that the neighborhood of S is defined as $N(S) := \{v \in V : \exists s \in S, \{s, v\} \in E\}$. We now define the *inclusive neighborhood* of a set S as $N[S] := S \cup N(S)$. We define the i -th neighborhood as $N^i[v] := \{u \in V : d(u, v) \leq i\} = N[N^{i-1}[v]]$ where $d(u, v)$ is the shortest-path distance between u and v .

► **Lemma 19.** *Let $\phi \geq 1$. If G is a $(\phi, \frac{1}{1+\phi})$ -expander, then G has diameter at most 2 and minimum degree $\frac{\phi}{1+\phi}n$.*

Proof. Let D denote the diameter of G . Let $u, v \in V$ be a pair of nodes with distance D . Then we have that $u \in N^D[v] \setminus N^{D-1}[v]$. That means that $N^{D-1}[v] \neq V$, so $|N^{D-2}[v]| < \frac{n}{1+\phi}$, and $|N^D[v] \setminus N^{D-1}[v]| < \frac{n}{1+\phi}$. That means that

$$|N^{D-1}[v] \setminus N^{D-2}[v]| = n - |N^{D-2}[v]| - |N^D[v] \setminus N^{D-1}[v]| > n - 2 \frac{n}{1+\phi} = (\phi - 1) \frac{n}{1+\phi}.$$

But that means that the nodes at distance $D - 1$, i.e., $N^{D-1}[v] \setminus N^{D-2}[v]$, have edges to the entire graph, so also to v . Hence $N^{D-1}[v] \setminus N^{D-2}[v] = N[v] \setminus v$, and thus $D = 2$.

For the minimum degree we now note that

$$\deg(v) = |N(v)| = n - |N^2[v] \setminus N[v]| - 1 \geq n - \frac{n}{1+\phi} = \frac{\phi}{1+\phi}n,$$

since $|N^2[v] \setminus N[v]| < \frac{n}{1+\phi}$, so $|N^2[v] \setminus N[v]| + 1 \leq \frac{n}{1+\phi}$. $\blacktriangleleft$

So to make the definition less restrictive, we allow for α to take more values. However, for small values of α , the expansion property becomes local, which does not help us for rumor spreading.

► **Lemma 20.** *Let $\alpha \leq \frac{1}{2+2\phi}$. For any even $n \in \mathbb{N}$, there exists a (ϕ, α) -expander G that is not connected.*

Proof. Let $G = K_{n/2} + K_{n/2}$ be the sum of two disjoint complete graphs, which is clearly not connected. We show that G is a (ϕ, α) -expander. Let $S \subset V$ with $|S| \leq \alpha n$. We have

$$|\partial S| \geq \frac{n}{2} - |S| \geq \frac{n}{2} - \alpha n \geq \frac{n}{2} - \frac{n}{2+2\phi} = \frac{n}{2} \left(1 - \frac{1}{1+\phi}\right) = \frac{\phi n}{2+2\phi} \geq \phi \alpha n \geq \phi |S|. \quad \blacktriangleleft$$

We conclude that for rumor spreading we can focus on the regime $\frac{1}{2+2\phi} < \alpha \leq \frac{1}{1+\phi}$. We show that in this regime a (ϕ, α) -expander has small diameter.

► **Lemma 21.** *If G is a (ϕ, α) -expander for $\alpha > \frac{1}{2+2\phi}$, then it has diameter at most $O(\log_\phi n)$.*

Proof. Let $v \in V$ be an arbitrary node, we show that $|N^{\lceil \log_\phi n \rceil}[v]| > n/2$, hence for any pair of nodes $u, v \in V$, we have $N^{\lceil \log_\phi n \rceil}[u] \cap N^{\lceil \log_\phi n \rceil}[v] \neq \emptyset$, so there is a path from u to v of length at most $2\lceil \log_\phi n \rceil$.

By expansion we have that $|N^i[v]| \geq \phi^i$, if $|N^{i-1}[v]| \leq \alpha n$. Now if $|N^{\lceil \log_\phi n \rceil}[v]| \geq \phi^{\lceil \log_\phi n \rceil} = n$, we are done. So suppose not, i.e., suppose $|N^{\lceil \log_\phi n \rceil}[v]| < \alpha n$. For every $S \subset N^{\lceil \log_\phi n \rceil}[v]$ with $|S| = \alpha n$ it holds that $|S \cup \partial S| \geq (1 + \phi) \cdot \alpha n$ since G is a (ϕ, α) -expander. Hence $|N^{\lceil \log_\phi n \rceil}[v]| \geq |S \cup \partial S| \geq (1 + \phi) \cdot \alpha n$. Since we assume $\alpha > \frac{1}{2+2\phi}$ we have that $(1 + \phi)\alpha n > \frac{1+\phi}{2+2\phi}n = \frac{n}{2}$, which completes the proof. $\blacktriangleleft$

D Examples of Small-Set Vertex Expanders

A trivial example of (ϕ, α) -expanders are complete graphs.

► **Lemma 22.** *The complete graph on n nodes is a (ϕ, α) -vertex expander for every $0 < \phi \leq n - 1$ and $\alpha \leq \frac{1}{1+\phi}$.*

However, we can show also the existence of significantly sparser graphs that are (ϕ, α) -expanders. We consider Erdős-Rényi random graphs $G(n, p)$ on n nodes, where there is an edge between any pair of nodes with probability p . We also know that w.h.p. the graphs $G(n, \phi/n)$ have diameter $\Theta(\log_\phi n)$ for $\phi = \omega(1)$ [15].

► **Lemma 23.** *Let $\alpha < \frac{1}{1+\phi/(1-e^{-1})} \approx \frac{1}{1+1.58\phi}$. Let a be such that $\alpha = \frac{1}{1+a\phi}$ and $\phi \geq \Theta\left(\left(\frac{1}{a(1-e^{-1})-1}\right)^2 \log n\right)$. Erdős-Rényi random graphs $G(n, p)$ with $p \geq \frac{3\phi}{n}$ are (ϕ, α) -vertex expanders with high probability.*