

Faster rumor spreading with multiple calls

Konstantinos Panagiotou

Ludwig-Maximilians-University
Munich, Germany

kpanagio@math.lmu.de

Ali Pourmiri

Max Planck Institute for Informatics
Saarbrücken, Germany

pourmiri@mpi-inf.mpg.de

Thomas Sauerwald

University of Cambridge
Cambridge, United Kingdom

thomas.sauerwald@cl.cam.ac.uk

Submitted: Apr 25, 2014; Accepted: Dec 16, 2014; Published: Feb 9, 2015

Mathematics Subject Classifications: 68W20; 68M14; 05C90

Abstract

We consider the random phone call model introduced by Demers et al., which is a well-studied model for information dissemination in networks. One basic protocol in this model is the so-called **Push** protocol that proceeds in synchronous rounds. Starting with a single node which knows of a rumor, every informed node calls in each round a random neighbor and informs it of the rumor. The **Push-Pull** protocol works similarly, but additionally every uninformed node calls a random neighbor and may learn the rumor from it.

It is well-known that both protocols need $\Theta(\log n)$ rounds to spread a rumor on a complete network with n nodes. Here we are interested in how much the spread can be speeded up by enabling nodes to make more than one call in each round. We propose a new model where the number of calls of a node is chosen independently according to a probability distribution R . We provide both lower and upper bounds on the rumor spreading time depending on statistical properties of R such as the mean or the variance (if they exist). In particular, if R follows a power law distribution with exponent $\beta \in (2, 3)$, we show that the **Push-Pull** protocol spreads a rumor in $\Theta(\log \log n)$ rounds. Moreover, when $\beta = 3$, the **Push-Pull** protocol spreads a rumor in $\Theta(\frac{\log n}{\log \log n})$ rounds.

Contents

1	Introduction	3
2	Notations and Preliminaries	5
3	Some Useful Facts for Power Law Distributions	6
4	Push Protocol	7
5	Push Protocol with Bounded Mean and Bounded Variance	9
5.1	The Preliminary Phase	11
5.2	The Middle Phase	11
5.3	The Final Phase	15
6	Push Protocol with Bounded Mean	19
7	Lower Bound for Push-Pull	19
8	Push-Pull Protocol with Bounded Mean and Bounded Variance	21
9	Push-Pull Protocol with Power Law Distribution $2 < \beta < 3$	21
9.1	Proof of the Upper Bound	21
9.2	Proof of the Lower Bound	25
10	Push-Pull Protocol with Power Law Distribution $\beta = 3$	28
10.1	Proof of Lower Bound	28
10.2	Proof of Upper Bound	29
11	Generating a New C_u^t in Each Round (Theorem 1.6)	32

1 Introduction

Randomized rumor spreading is an important primitive for spreading information in networks. The goal is to spread a piece of information, the so-called rumor, from an arbitrary node to all the other nodes. Randomized rumor spreading protocols are based on the simple idea that every node picks a random neighbor and these two nodes are able to exchange information in that round. This paradigm ensures that the protocol is local, scalable, and robust against network failures (cf. [13, 15]). Therefore, these protocols have been successfully applied in other contexts such as replicated databases [8], failure detection [30], resource discovery [23], load balancing [3], data aggregation [25], and analysis of the spread of computer viruses [2].

The most basic variant of randomized rumor spreading is the **Push** protocol. At the beginning, there is a single node who knows of some rumor. Then in each of the following rounds every *informed* node calls a random neighbor chosen independently and uniformly at random and informs it of the rumor. The **Pull** protocol is symmetric, here every *uninformed* node calls a random neighbor chosen independently and uniformly at random, and if that neighbor happens to be informed the node becomes informed. The **Push-Pull** protocol is simply the combination of both protocols. Most studies in randomized rumor spreading concern the *runtime*, which is the number of rounds required until the rumor reaches all other nodes, and the *communication overhead*, which is the total number of information exchanges, produced by these protocols (see e.g. [24]).

In one of the first papers in this area, Frieze and Grimmett [19] proved that if the underlying graph is a complete graph with n nodes, then the runtime of **Push** is $\log_2 n + \log n \pm o(\log n)$ with high probability¹, where $\log n$ denotes the natural logarithm of n . This result was later strengthened by Pittel [29]. For the standard **Push-Pull** protocol, Karp et al. [24] proved a runtime bound of $\log_3 n + \mathcal{O}(\log \log n)$. In order to overcome the large number of $\Theta(n \log n)$ calls, Karp et al. also presented an extension of the **Push-Pull** protocol together with a termination mechanism that spreads a rumor in $\mathcal{O}(\log n)$ rounds using only $\mathcal{O}(n \log \log n)$ messages. Doerr and Fouz [9] proposed a new protocol using only **Push** calls that achieves a runtime of $(1 + o(1)) \log_2 n$ using only $\mathcal{O}(n \cdot f(n))$ calls (and messages), where $f(n)$ is an arbitrarily slowly growing function.

Besides the complete graph, randomized rumor spreading protocols have been shown to be efficient also on other topologies. In particular, their runtime is at most logarithmic in n for topologies ranging from basic networks, such as random graphs [15, 14, 16] and hypercubes [15], random regular graphs [1, 17], graphs with constant conductance [27, 6, 20], constant weak conductance [4] or constant vertex expansion [22, 21], to more complex structures including preferential attachment graphs modeling social networks [5]. In particular, recent studies establishing a sub-logarithmic runtime on certain social network models [10, 11, 18] raise the question whether it is possible to achieve a sub-logarithmic runtime also on the complete graph. In addition to analyses on static graphs, there are also studies on mobile geometric graphs, e.g., [7, 28].

¹with probability $1 - o(1)$ as $n \rightarrow \infty$. For simplicity, we sometimes omit the “with high probability” in the introduction.

Since all aforementioned protocols require $\Theta(\log n)$ rounds to spread the rumor on a complete graph, we equip nodes with the possibility of calling more than one node in each round. Specifically, we assume that the *power* of a node u , denoted by C_u , is a random variable, which has the same distribution as a random variable R with support on the positive integers and which is independent of u . In order to keep the overall communication cost small, we focus on distributions R satisfying $\sum_{u \in V} C_u = \mathcal{O}(n)$ with high probability – in particular, R has bounded mean. Our aim is to understand the impact of the distribution of R on the runtime of randomized rumor spreading. In particular, we seek for conditions on R which are necessary (and/or sufficient) for a sublogarithmic runtime.

Our first result concerns the **Push** protocol for the case where R has bounded mean and bounded variance, which is the most basic setting. Let T_{total} be the first round in which all nodes are informed.

Theorem 1.1. *Consider the Push protocol and assume that R is a distribution with $\mathbf{E}[R] = \mathcal{O}(1)$ and $\mathbf{Var}[R] = \mathcal{O}(1)$. Then $|T_{total} - (\log_{1+\mathbf{E}[R]} n + \log_{e^{\mathbf{E}[R]}} n)| = o(\log n)$.*

Note that by putting $R \equiv 1$, we retain the classic result by Frieze and Grimmett. Our next result addresses the case where we drop the assumption on the variance, and it provides a lower bound of $\Omega(\log n)$ on the number of rounds. Although this result is less precise than Theorem 1.1, it demonstrates that it is necessary to consider the **Push-Pull** protocol in order to achieve a sub-logarithmic runtime.

Theorem 1.2. *Assume that R is any distribution with $\mathbf{E}[R] = \mathcal{O}(1)$. Then with probability $1 - o(1)$, the Push protocol needs at least $\Omega(\log n)$ rounds to inform all nodes.*

We point out that the lower bound in Theorem 1.2 is tight up to constant factors, as the results in [19, 29] for the standard **Push** protocol imply an upper bound of $\mathcal{O}(\log n)$ rounds. We now consider the **Push-Pull** protocol and extend the lower bound of $\Omega(\log n)$ from Theorem 1.1.

Theorem 1.3. *Assume that R is any distribution with $\mathbf{E}[R] = \mathcal{O}(1)$ and $\mathbf{Var}[R] = \mathcal{O}(1)$. Then for any constant $\epsilon > 0$, with probability $1 - \epsilon$ the Push-Pull protocol needs at least $\Omega(\log n)$ rounds to inform all nodes.*

Theorem 1.3 establishes that an unbounded variance is necessary to break the $\Omega(\log n)$ lower bound. An important distribution with bounded mean but unbounded variance is the *power law distribution* with exponent $\beta \leq 3$, i.e., there are constants $0 < c_1 \leq c_2$ such that $c_1 z^{1-\beta} \leq \mathbf{Pr}[R \geq z] \leq c_2 z^{1-\beta}$ for any $z \geq 1$, and $\mathbf{Pr}[R \geq 1] = 1$. We are especially interested in power law distributions, because they are scale invariant and have been observed in a variety of settings. Our main result below shows that this natural distribution achieves a sublogarithmic runtime.

Theorem 1.4. *Assume that R is a power law distribution with $2 < \beta < 3$. Then the Push-Pull protocol informs all nodes in $\Theta(\log \log n)$ rounds with probability $1 - o(1)$.*

Notice that if R is a power law distribution with $\beta > 3$, then Theorem 1.3 applies because the variance of R is bounded. Hence our results reveal a dichotomy in terms of the exponent β : if $2 < \beta < 3$, then the **Push-Pull** protocol finishes in $\mathcal{O}(\log \log n)$ rounds, whereas for $\beta > 3$ the **Push-Pull** protocol finishes in $\Theta(\log n)$ rounds². While a very similar dichotomy was shown in [18] for random graphs with a power law degree distribution, our result here concerns the spread of the rumor from one to *all* nodes.

In the case $\beta = 3$ we show that the runtime is close to the one in the $\beta > 3$ case.

Theorem 1.5. *Assume that R is a power law distribution with $\beta = 3$. Then the **Push-Pull** protocol informs all nodes in $\Theta\left(\frac{\log n}{\log \log n}\right)$ rounds with probability $1 - o(1)$.*

Finally, we also argue that it is necessary that the C_u 's are chosen once and for all at the beginning, and they are not updated in each round. Indeed, suppose we generate in the t th round a new variable C_u^t , which is the number of calls made by u in that round. Then we prove the following lower bound.

Theorem 1.6. *Assume that R is any distribution with $\mathbf{E}[R] = \mathcal{O}(1)$. Then with probability $1 - o(1)$, the **Push-Pull** protocol needs $\Omega(\log n)$ rounds to inform all nodes.*

2 Notations and Preliminaries

We introduce some notation that will be used throughout the paper without further reference. In our setting, the **Push**, **Pull** and **Push-Pull** protocols proceed like the classic ones except that in each round, every (un)informed node u calls C_u node(s) chosen independently and uniformly at random and sends (requests) the rumor. For any of these protocols, we let \mathcal{I}_t be the set of informed nodes at the end of round t and \mathcal{U}_t the set of uninformed nodes. We write $\mathcal{V} = \mathcal{I}_t \cup \mathcal{U}_t$ for the vertex set of the graph, and we assume $|\mathcal{V}| = n$. The size of \mathcal{I}_t and \mathcal{U}_t is denoted by I_t and U_t . We indicate the set of newly informed nodes in round $t + 1$ by \mathcal{N}_t and its size is N_t . Let S_t be the number of **Push** calls in round $t + 1$, so $S_t = \sum_{u \in \mathcal{I}_t} C_u \geq I_t$. Let us define $\mathcal{N}_t^{\text{Pull}}$ and $\mathcal{N}_t^{\text{Push}}$ to be the set of newly informed nodes by **Pull** and **Push** calls in round $t + 1$, respectively. The size of $\mathcal{N}_t^{\text{Pull}}$ and $\mathcal{N}_t^{\text{Push}}$ are denoted by N_t^{Pull} and N_t^{Push} . The size of every set divided by n will be denoted by the corresponding small letter, so i_t , n_t and s_t are used to denote I_t/n , N_t/n , and S_t/n , respectively. Further, let

$$\mathcal{L}(z) := \{u \in \mathcal{V} : C_u \geq z\} \quad \text{and set} \quad L(z) = \mathcal{L}(z).$$

Moreover, let $\Delta = \max_{u \in \mathcal{V}} C_u$.

We will use extensively the following two concentration inequalities. The first one is a Chernoff-type bound.

²We do not consider the case $\beta \leq 2$, since then there exists at least one node with degree $\Omega(n)$ and the rumor is spread in constant time. Additionally, $\mathbf{E}[R]$ is no longer bounded.

Theorem 2.1 ([12]). *Suppose that $X_1, X_2, \dots, X_n \in \{0, 1\}$ are independent and identically distributed random variables and let $X := \sum_{i=1}^n X_i$. Then for any $\lambda \geq 0$*

$$\Pr [|X - \mathbf{E}[X]| \geq \lambda] \leq 2 \cdot e^{-\frac{\lambda^2}{2(\mathbf{E}[X] + \lambda/3)}}.$$

In particular,

$$\Pr [|X - \mathbf{E}[X]| \geq \mathbf{E}[X]/2] \leq 2 \cdot e^{-\frac{\mathbf{E}[X]^2}{8(\mathbf{E}[X] + \mathbf{E}[X]/6)}} < 2 \cdot e^{-\frac{\mathbf{E}[X]}{10}}.$$

The next inequality is known as the Bounded Difference inequality.

Theorem 2.2 ([26]). *Suppose that X_1, X_2, \dots, X_n are independent random variables and every X_i , $1 \leq i \leq n$, takes a value from a finite set A_i . Let $f : \prod_{1 \leq i \leq n} A_i \rightarrow \mathbb{R}$ be a real-valued function so that there exist c_1, c_2, \dots, c_n with*

$$\sup_{x_1, x_2, \dots, x_n, x'_i} |f(x_1, x_2, \dots, x_i, \dots, x_n) - f(x_1, x_2, \dots, x'_i, \dots, x_n)| \leq c_i, \text{ for every } 1 \leq i \leq n.$$

Then, for every $\lambda > 0$,

$$\Pr [|f(X_1, X_2, \dots, X_n) - \mathbf{E}[f(X_1, X_2, \dots, X_n)]| \geq \lambda] \leq 2 \cdot e^{-\frac{\lambda^2}{2 \sum_{i=1}^n c_i^2}}.$$

3 Some Useful Facts for Power Law Distributions

Let R be a power law probability distribution with exponent β , i.e., there are constants $0 < c_1 < c_2$ so that for every integer $z \geq 1$,

$$c_1 \cdot z^{1-\beta} \leq \Pr [C_u \geq z] \leq c_2 \cdot z^{1-\beta},$$

and $\Pr [C_u \geq 1] = 1$. In this section we collect some basic properties of R .

Fact 3.1. *If R is a power law distribution with $\beta > 3$, then $\mathbf{Var}[R] = \mathcal{O}(1)$.*

Proof. Since $\beta > 3$

$$\mathbf{Var}[R] \leq \mathbf{E}[R^2] = \sum_{z \geq 1} \Pr [R^2 \geq z] \leq 1 + \sum_{z \geq 2} \sqrt{c_2 \cdot z^{1-\beta}} < \infty.$$

□

Fact 3.2. *Let $\beta > 2$. Let C_u , $u \in \mathcal{V}$ be independent, power-law distributed random variables with exponent β . Then, with probability $1 - o(\frac{1}{\log n})$,*

$$\Delta := \max_{u \in \mathcal{V}} C_u \leq n^{\frac{1}{\beta-1}} \cdot \log n.$$

Proof. By definition,

$$\Pr \left[C_u \geq n^{\frac{1}{\beta-1}} \log n \right] \leq c_2 \cdot n^{-1} \cdot \log^{1-\beta}(n).$$

Applying the union bound yields the claim. \square

Recall that $\mathcal{L}(z) := \{u \in \mathcal{V}, C_u \geq z\}$ and $L(z) := |\mathcal{L}(z)|$.

Proposition 3.3. *Let $\beta > 2$. Let $C_u, u \in \mathcal{V}$ be independent, power-law distributed random variables with exponent β . Then, for every $z = \mathcal{O}(n^{\frac{1}{\beta-1}} / \log n)$, with probability $1 - o(1/n)$*

$$\frac{n \cdot c_1 \cdot z^{1-\beta}}{2} \leq L(z) \leq \frac{3 \cdot n \cdot c_2 \cdot z^{1-\beta}}{2}.$$

Proof. For $u \in \mathcal{V}$ let I_u be the indicator random variable for the event $C_u \geq z$. Since the C_u 's are independent and identically distributed, so are the I_u 's. By linearity of expectation

$$n \cdot c_1 \cdot z^{1-\beta} \leq \mathbf{E}[L(z)] \leq n \cdot c_2 \cdot z^{1-\beta}.$$

Applying Theorem 2.1 to the random variable $X := \sum_{u \in \mathcal{V}} I_u$ yields that

$$\Pr [|L(z) - \mathbf{E}[L(z)]| > \mathbf{E}[L(z)]/2] < 2 \cdot e^{-\frac{\mathbf{E}[L(z)]}{10}} \leq 2 \cdot e^{-\frac{n \cdot c_1 \cdot z^{1-\beta}}{10}}$$

Since $z = \mathcal{O}(n^{\frac{1}{\beta-1}} / \log n)$, the claim follows. \square

4 Push Protocol

In this section we will show two general lemmas for the Push protocol that are valid for any R with support on the positive integers. They will be used when analyzing the Push and the Push-Pull protocols.

Lemma 4.1. *Consider the Push protocol and suppose that $S_t \leq \log^c n$, where $c > 0$ is an arbitrary constant. Then with probability at least $1 - \mathcal{O}(n^{-1} \log^{2c} n)$ we have*

$$I_{t+1} = I_t + S_t.$$

Proof. Recall that S_t is the number of Push calls in round $t + 1$. By applying the union bound, the probability that an informed node receives a call in round $t + 1$ is bounded by $\frac{S_t I_t}{n}$. So, with probability at least $1 - \frac{S_t I_t}{n}$, none of the calls are sent to a node in \mathcal{I}_t . Conditioning on this event, consider all calls one by one in an arbitrary order; then the probability that the i -th call informs a different node from the previous $i - 1$ calls is $1 - \frac{i-1}{U_t}$. Therefore the conditional probability that S_t calls inform S_t different nodes is at least

$$\prod_{i=1}^{S_t-1} \left(1 - \frac{i}{U_t} \right) > \left(1 - \frac{S_t-1}{U_t} \right)^{S_t} \geq 1 - \frac{S_t^2}{U_t}.$$

So the probability that S_t calls inform S_t different uninformed nodes is at least

$$\left(1 - \frac{S_t I_t}{n}\right) \cdot \left(1 - \frac{S_t^2}{U_t}\right) = 1 - \mathcal{O}\left(\frac{S_t^2}{n}\right),$$

where the above equality holds because $I_t \leq S_t \leq \log^c n$ and $U_t = n(1 - o(1))$. The claim follows. □

Lemma 4.2. *Consider the Push protocol. Then with probability at least $1 - o(\frac{1}{\log n})$*

$$s_t - 2s_t^2 - 2\sqrt{\frac{s_t \log \log n}{n}} \leq n_t \leq s_t.$$

Proof. Since N_t is always bounded by S_t , $n_t \leq s_t$. To see the lower bound, let for $v \in \mathcal{U}_t$ Z_v be the indicator random variable for the event $v \in \mathcal{I}_{t+1}$. Then $N_t = \sum_{v \in \mathcal{U}_t} Z_v$. Since the Z_v 's are identically distributed random variables,

$$\mathbf{E}[N_t] = U_t \cdot \mathbf{Pr}[Z_v = 1].$$

Let $X_i \in \mathcal{V}$, $1 \leq i \leq N = S_t$, denote the target of the i -th call. Define

$$f(X_1, X_2, \dots, X_N) := N_t$$

to be the function counting the number of newly informed nodes in round $t + 1$. Then $\mathbf{E}[f(X_1, X_2, \dots, X_N)] = \mathbf{E}[N_t]$. For each change in just one coordinate of f , the following statement holds:

$$\sup_{x_1, x_2, \dots, x_i, x'_i \in \mathcal{V}} |f(x_1, x_2, \dots, x_i, \dots, x_N) - f(x_1, x_2, \dots, x'_i, \dots, x_N)| \leq 1.$$

Therefore by applying Theorem 2.2, we obtain

$$\mathbf{Pr}\left[|N_t - \mathbf{E}[N_t]| \geq \sqrt{4 \cdot S_t \cdot \log \log n}\right] \leq 2 \cdot e^{-4S_t \log \log n / 2S_t} = o(1/\log n).$$

So with probability $1 - o(1/\log n)$ we have

$$N_t > \mathbf{E}[N_t] - 2\sqrt{S_t \log \log n} = U_t \cdot \mathbf{Pr}[Z_v = 1] - 2\sqrt{S_t \log \log n}. \tag{1}$$

Now we estimate $\mathbf{Pr}[Z_v = 1]$. By the definition of Push

$$\mathbf{Pr}[Z_v = 1] = 1 - \prod_{u \in \mathcal{I}_t} \left(1 - \frac{1}{n}\right)^{C_u}.$$

Using that $1 - x \leq e^{-x} \leq 1 - x + x^2$ for any $x \geq 0$

$$\mathbf{Pr}[Z_v = 1] \geq 1 - e^{-\sum_{u \in \mathcal{I}_t} C_u/n} = 1 - e^{-s_t} \geq s_t - s_t^2.$$

We now plug the value obtained by the above formula into (1) and normalize it. So we obtain

$$\begin{aligned} n_t &\geq (1 - i_t) \cdot (s_t - s_t^2) - 2\sqrt{\frac{s_t \log \log n}{n}} \\ &= s_t - s_t^2 - i_t \cdot (s_t - s_t^2) - 2\sqrt{\frac{s_t \log \log n}{n}} \geq s_t - 2s_t^2 - 2\sqrt{\frac{s_t \log \log n}{n}}, \end{aligned}$$

where the last inequality comes from the fact that $i_t \leq s_t$. □

Corollary 4.3. *Consider the Push protocol. Then with probability at least $1 - o(1/\log n)$ for any round t in which $S_t \leq n/8$ we have that $I_{t+1} \geq I_t + S_t/2$.*

Proof. If $1 \leq S_t \leq \log n$, then Lemma 4.1 yields that with probability $1 - o(1/\log n)$ we have $N_t = S_t$. If $\log n \leq S_t \leq n/8$, then $2s_t^2 \leq s_t/4$ and $2\sqrt{s_t n^{-1} \log \log n} \leq s_t/4$. Thus, Lemma 4.2 guarantees that with probability at least $1 - o(1/\log n)$

$$n_t \geq s_t - 2s_t^2 - 2\sqrt{\frac{s_t \log \log n}{n}} \geq \frac{s_t}{2}.$$

□

Corollary 4.4. *Consider the Push protocol. For any round t and positive integer $k = \mathcal{O}(\log n)$ in which $S_{t+k} = o(n)$, with probability $1 - o(\frac{k}{\log n})$*

$$I_{t+k} \geq I_t \cdot (3/2)^k.$$

Proof. By assumption we have for every $1 \leq i \leq k$ that $S_{t+i} = o(n)$. Applying Corollary 4.3 shows that with probability $1 - o(\frac{1}{\log n})$

$$I_{t+i} \geq I_{t+i-1} + \frac{S_{t+i-1}}{2} \geq I_{t+i-1} \cdot \frac{3}{2}.$$

Using an inductive argument and the union bound for k implies the statement. □

5 Push Protocol with Bounded Mean and Bounded Variance

This section is devoted to the proof of Theorem 1.1. Recall that $T_{total} := \min\{t \mid I_t = n\}$, i.e., the first round in which all nodes are informed. We claim that if $\mathbf{E}[R] = \mathcal{O}(1)$ and $\mathbf{Var}[R] = \mathcal{O}(1)$, then $|T_{total} - (\log_{1+\mathbf{E}[R]} n + \log_{e^{\mathbf{E}[R]}} n)| = o(\log n)$. To prove this result, we study the protocol in three consecutive phases. In the following we give a brief overview of the proof.

- **The Preliminary Phase.** This phase starts with just one informed node and ends when $I_t \geq \log^5 n$. Similar to the proof of the birthday paradox we show that in each round every Push call informs a different uninformed node and thus the number of informed nodes increases by $S_t \geq I_t$. Hence after $\mathcal{O}(\log \log n)$ rounds there are at least $\log^5 n$ informed nodes. Further, since $\mathbf{E}[R] = \mathcal{O}(1)$, after $\mathcal{O}(\log \log n)$ rounds we also have $S_t \leq \log^{\mathcal{O}(1)} n$ for all these rounds.

- **The Middle Phase.** This phase starts when $\log^5 n \leq I_t \leq S_t \leq \log^{\mathcal{O}(1)} n$ and ends when $I_t \geq n/\log \log n$. First we show that the number of **Push** calls S_t increases by a factor of approximately $1 + \mathbf{E}[R]$ as long as the number of informed nodes is $o(n)$. Then we prove that the number of newly informed nodes in round $t + 1$ is roughly the same as S_t . Therefore an inductive argument shows that it takes $\log_{1+\mathbf{E}[R]} n \pm o(\log n)$ rounds to reach $n/\log \log n$ informed nodes.
- **The Final Phase.** This phase starts when $I_t \geq \frac{n}{\log \log n}$ and ends when all nodes are informed with high probability. In this phase, we first prove that after $o(\log n)$ rounds the number of uninformed nodes decreases to $n/\log^5 n$. Then we show the probability that an arbitrary uninformed node remains uninformed is $e^{-\mathbf{E}[R] \pm o(1/\log n)}$. Finally, an inductive argument establishes that it takes $\log_{e^{\mathbf{E}[R]}} n \pm o(\log n)$ rounds until every node is informed.

In the following we present the detailed proofs for these phases. Before that we show the following proposition.

Proposition 5.1. *Let $\epsilon > 0$ and let R be a random variable with support on the positive integers such that $\mathbf{E}[R] = \mathcal{O}(1)$ and $\mathbf{Var}[R] = \mathcal{O}(1)$. Let $\delta \geq 0$ be such that $U_t = n^{1-\delta}$, for some round t . Then, with probability $1 - o(1/\log n)$,*

$$\sum_{u \in \mathcal{U}_t} C_u = \mathcal{O}(n^{1-\delta/2} \cdot \log^{1+\epsilon} n).$$

Proof. For $k \in \mathbb{N}$ let us define a random variable

$$W_k := \sum_{u \in \mathcal{V}} C_u \cdot \mathbb{1}(C_u \geq k).$$

By linearity of expectation

$$\begin{aligned} \mathbf{E}[W_k] &= \sum_{u \in \mathcal{V}} \mathbf{E}[C_u \mathbb{1}(C_u \geq k)] \\ &= n \cdot \mathbf{E}[C_u \mathbb{1}(C_u \geq k)] = n \cdot \sum_{l \geq k} l \cdot \mathbf{Pr}[C_u = l] \leq \frac{n}{k} \cdot \sum_{l \geq k} l^2 \cdot \mathbf{Pr}[C_u = l]. \end{aligned}$$

Since C_u has bounded variance the last sum is in $\mathcal{O}(1)$. Thus, $\mathbf{E}[W_k] = \mathcal{O}(n/k)$. Markov's inequality implies that with probability $1 - o(1/\log n)$, $W_k = \mathcal{O}(n \cdot \log^{1+\epsilon} n \cdot k^{-1})$. If we set $k = n^{\delta/2}$, then

$$\sum_{u \in \mathcal{U}_t} C_u = \sum_{\{u \in \mathcal{U}_t : C_u \geq k\}} C_u + \sum_{\{u \in \mathcal{U}_t : C_u < k\}} C_u \leq W_k + \mathcal{O}(n^{1-\delta} \cdot k) = \mathcal{O}(n^{1-\delta/2} \cdot \log^{1+\epsilon} n).$$

□

5.1 The Preliminary Phase

This phase starts with one informed node and ends when $I_t \geq \log^5 n$ and $S_t \leq \log^{\mathcal{O}(1)} n$. Let T_0 be the first round in which the number of informed nodes exceeds $\log^5 n$.

Lemma 5.2. *For any $t = \mathcal{O}(\log \log n)$, with probability at least $1 - \log^{-3} n$, we have $S_t = \log^{\mathcal{O}(1)} n$.*

Proof. We will bound the expected number of calls in each round t as follows:

$$\mathbf{E}[S_t | S_{t-1}] = S_{t-1} + \mathbf{E}\left[\sum_{u \in \mathcal{N}_{t-1}} C_u \mid S_{t-1}\right] = S_{t-1} + N_{t-1} \cdot \mathbf{E}[R] \leq S_{t-1} \cdot (1 + \mathbf{E}[R]),$$

where the last inequality comes from the fact that $N_{t-1} \leq S_{t-1}$. Since the origin of the rumor is chosen before determining the C_u 's we have $\mathbf{E}[S_0] = \mathbf{E}[R]$. Applying the law of total expectation yields

$$\mathbf{E}[S_t] = \mathbf{E}[\dots \mathbf{E}[\mathbf{E}[S_t | S_{t-1}] | S_{t-2}] \dots | S_0] \leq (1 + \mathbf{E}[R])^t \mathbf{E}[S_0] = (1 + \mathbf{E}[R])^t \mathbf{E}[R].$$

By using Markov's inequality we have that

$$\Pr[S_t \geq (1 + \mathbf{E}[R])^t \cdot \mathbf{E}[R] \cdot \log^3 n] \leq \log^{-3} n.$$

and the claim follows for any $t = \mathcal{O}(\log \log n)$. □

Corollary 5.3. *With probability $1 - o(1)$ we have $T_0 = \mathcal{O}(\log \log n)$.*

Proof. Lemma 5.2 asserts that with probability at least $1 - \mathcal{O}(\log^{-3} n)$, $S_t = \log^{\mathcal{O}(1)} n$ for any $t = \mathcal{O}(\log \log n)$. Conditioning on this event, Lemma 4.1 guarantees that with probability $1 - (n^{-1} \log^{\mathcal{O}(1)} n)$, for any $t = \mathcal{O}(\log \log n)$,

$$I_{t+1} = I_t + S_t \geq 2I_t,$$

where the inequality comes from the fact that $S_t \geq I_t$. So, with probability at least

$$\left(1 - \frac{1}{\log^3 n}\right) \left(1 - \mathcal{O}(\log \log n) \cdot \frac{\log^{\mathcal{O}(1)} n}{n}\right) = 1 - o(1),$$

there exists a round $T_0 = \mathcal{O}(\log \log n)$ such that $I_{T_0} \geq \log^5 n$ and $S_{T_0} \leq \log^{\mathcal{O}(1)} n$. □

5.2 The Middle Phase

The phase starts when $\log^5 n \leq I_t \leq S_t \leq \log^{\mathcal{O}(1)} n$ and ends when $I_t \geq n/\log \log n$. Let T_1 be the first round so that $I_{T_1} \geq n/\log \log n$. The main result of this subsection is that $|T_1 - \log_{1+\mathbf{E}[R]} n| = o(\log n)$ with high probability.

Lemma 5.4. *Suppose for a round t we have $s_t = \Omega(n^{-1} \cdot \log^5 n)$ and $s_t = o(1)$. Then for any $k = \mathcal{O}(\log n)$ with $(1 + \mathbf{E}[R])^k s_t = o(1)$, with probability $1 - o(k/\log n)$,*

$$\text{for all } 1 \leq i \leq k, \quad (1 + \mathbf{E}[R])^i \cdot s_t \cdot (1 - o(1)) \leq s_{t+i} \leq (1 + \mathbf{E}[R])^i \cdot s_t \cdot (1 + o(1)). \quad (2)$$

Proof. Consider the random variable $\sum_{u \in \mathcal{N}_t} C_u$. Since \mathcal{N}_t is fixed and the random variables C_u , $u \in \mathcal{N}_t$ are independent we obtain that $\mathbf{E}[\sum_{u \in \mathcal{N}_t} C_u] = N_t \cdot \mathbf{E}[R]$. Moreover,

$$\mathbf{Var} \left[\sum_{u \in \mathcal{N}_t} C_u \right] = N_t \cdot \mathbf{Var}[R].$$

Chebychev's inequality implies that

$$\Pr \left[\left| \sum_{u \in \mathcal{N}_t} C_u - N_t \mathbf{E}[R] \right| \geq \sqrt{N_t \log^2 n} \right] \leq \frac{N_t \mathbf{Var}[R]}{N_t \log^2 n} = o\left(\frac{1}{\log n}\right).$$

Since $S_{t+1} = S_t + \sum_{u \in \mathcal{N}_t} C_u$, it follows that with probability $1 - o(1/\log n)$,

$$S_t + N_t \cdot \mathbf{E}[R] - \sqrt{N_t \log^2 n} \leq S_{t+1} \leq S_t + N_t \cdot \mathbf{E}[R] + \sqrt{N_t \log^2 n}. \quad (3)$$

Using the above formula and the fact that $N_t \leq S_t$ we have

$$S_{t+1} \leq S_t + S_t \cdot \mathbf{E}[R] + \sqrt{S_t \log^2 n} \leq S_t \cdot \left(1 + \mathbf{E}[R] + \sqrt{\frac{\log^2 n}{S_t}} \right).$$

Since S_t is a non-decreasing function in t and $\log^5 n \leq I_t \leq S_t$, with probability $1 - o(1/\log n)$

$$s_{t+1} \leq s_t \cdot (1 + \mathbf{E}[R]) \left(1 + \sqrt{\frac{\log^2 n}{(1 + \mathbf{E}[R])^2 \log^5 n}} \right) < s_t \cdot (1 + \mathbf{E}[R]) \left(1 + \frac{1}{\log^{\frac{3}{2}} n} \right).$$

An inductive argument and the union bound for all k events that violate the above inequality shows that for any $k = \mathcal{O}(\log n)$ with probability $1 - o(k/\log n)$,

$$\text{for all } 1 \leq i \leq k, \quad s_{t+i} \leq s_t \cdot (1 + \mathbf{E}[R])^i (1 + o(1)). \quad (4)$$

In order to prove the left hand side of (2), we use Lemma 4.2 which states with probability $1 - o(1/\log n)$,

$$n_t \geq s_t - 2s_t^2 - 2\sqrt{\frac{s_t \log \log n}{n}}.$$

Using the lower bound in (3) and the above formula implies that with probability $1 - o(1/\log n)$,

$$\begin{aligned}
s_{t+1} &\geq s_t + n_t \cdot \mathbf{E}[R] - \sqrt{\frac{n_t \log^2 n}{n}} \\
&\geq s_t + s_t \cdot \mathbf{E}[R] - 2s_t^2 \cdot \mathbf{E}[R] - 2\sqrt{\frac{s_t \log \log n}{n}} \cdot \mathbf{E}[R] - \sqrt{\frac{s_t \cdot \log^2 n}{n}} \\
&\geq (1 + \mathbf{E}[R])s_t - 2\mathbf{E}[R]s_t^2 - 2\sqrt{\frac{s_t \log^2 n}{n}} \\
&\geq (1 + \mathbf{E}[R])s_t - F(s_t),
\end{aligned}$$

where $F(s_t) = 2\mathbf{E}[R]s_t^2 + 2\sqrt{n^{-1} \cdot s_t \cdot \log^2 n}$. An inductive argument and the union bound for all k events that violate the above inequality show that for any integer k for which $(1 + \mathbf{E}[R])^k \cdot s_t = o(1)$ with probability $1 - o(k/\log n)$,

$$\text{for all } 1 \leq i \leq k, \quad s_{t+i} \geq (1 + \mathbf{E}[R])^i s_t - \sum_{j=0}^{i-1} (1 + \mathbf{E}[R])^j F(s_{t+i-j}). \quad (5)$$

Inequality (4) yields that with probability $1 - o(k/\log n)$,

$$\text{for all } 1 \leq i \leq k = \mathcal{O}(\log n), \quad s_{t+i} \leq a \cdot s_t \cdot (1 + \mathbf{E}[R])^i,$$

where $a := 1 + o(1)$. $F(s_t)$ is a non-decreasing function in s_t and hence for any $k = \mathcal{O}(\log n)$ and $1 \leq j \leq k$,

$$\begin{aligned}
F(s_{t+i-j}) &\leq F(a \cdot (1 + \mathbf{E}[R])^{i-j} s_t) \\
&\leq 2\mathbf{E}[R] (1 + \mathbf{E}[R])^{2(i-j)} (a \cdot s_t)^2 + 2(1 + \mathbf{E}[R])^{\frac{i-j}{2}} \sqrt{\frac{a \cdot s_t \log^2 n}{n}}.
\end{aligned}$$

Hence by combining the above inequality and (5), we conclude that for any integer k , where $(1 + \mathbf{E}[R])^k s_t = o(1)$ and $k = \mathcal{O}(\log n)$ with probability $1 - o(k/\log n)$, for all $1 \leq i \leq k$

$$\begin{aligned}
&s_{t+i} \\
&\geq (1 + \mathbf{E}[R])^i s_t - 2\mathbf{E}[R] \sum_{j=0}^{i-1} (1 + \mathbf{E}[R])^{2i-j} (c \cdot s_t)^2 - 2 \sum_{j=0}^{i-1} (1 + \mathbf{E}[R])^{\frac{i+j}{2}} \sqrt{\frac{c \cdot s_t \log^2 n}{n}} \\
&\geq (1 + \mathbf{E}[R])^i s_t - d_1 \cdot (1 + \mathbf{E}[R])^{2i} s_t^2 - d_2 \cdot (1 + \mathbf{E}[R])^i \cdot \sqrt{\frac{s_t \log^2 n}{n}} \\
&= (1 + \mathbf{E}[R])^i s_t \cdot \left(1 - d_1 \cdot (1 + \mathbf{E}[R])^i s_t - d_2 \cdot \sqrt{\frac{\log^2 n}{s_t n}} \right),
\end{aligned}$$

where d_1 and d_2 are constants which do not depend on i . Since $(1 + \mathbf{E}[R])^k s_t = o(1)$ and $s_t = \Omega(\frac{\log^5 n}{n})$, for any $1 \leq i \leq k$,

$$s_{t+i} \geq (1 + \mathbf{E}[R])^i \cdot s_t \cdot (1 - o(1)).$$

□

Lemma 5.5. *Suppose that $\frac{\log^5 n}{n} \leq i_t \leq s_t \leq \frac{\log^{\mathcal{O}(1)} n}{n}$. Then for any $k = \mathcal{O}(\log n)$ with $(1 + \mathbf{E}[R])^k s_t = o(1)$, with probability $1 - o(1)$,*

$$i_t + f_2 \cdot (1 + \mathbf{E}[R])^k \cdot s_t \cdot (1 - o(1)) \leq i_{t+k} \leq i_t + f_1 \cdot (1 + \mathbf{E}[R])^k \cdot s_t \cdot (1 + o(1)),$$

where $f_1 > 0$ and $f_2 > 0$ are constants.

Proof. It is easy to see that

$$i_{t+k} = i_t + \sum_{i=0}^{k-1} n_{t+i} \leq i_t + \sum_{i=0}^{k-1} s_{t+i}.$$

Applying Lemma 5.4 implies that for any integer k for which $(1 + \mathbf{E}[R])^k \cdot s_t = o(1)$, with probability $1 - o(\frac{k}{\log n})$ the following upper bound holds:

$$\begin{aligned} i_{t+k} &\leq i_t + \sum_{i=0}^{k-1} s_{t+i} \leq i_t + s_t \cdot (1 + o(1)) \cdot \sum_{i=0}^{k-1} (1 + \mathbf{E}[R])^i \\ &= i_t + f_1 \cdot (1 + \mathbf{E}[R])^k \cdot s_t \cdot (1 + o(1)), \end{aligned}$$

where $f_1 > 0$ is a constant. On the other hand, Lemma 4.2 yields that with probability $1 - o(\frac{1}{\log n})$,

$$n_t \geq s_t - 2s_t^2 - 2\sqrt{\frac{s_t \log \log n}{n}}.$$

Another application of Lemma 5.4 shows that with probability $1 - o(\frac{k}{\log n})$, for all integers $1 \leq i \leq k$ in which $(1 + \mathbf{E}[R])^k s_t = o(1)$ and $s_t \geq \frac{\log^5 n}{n}$,

$$(1 + \mathbf{E}[R])^i \cdot s_t \cdot (1 - o(1)) \leq s_{t+i} \leq (1 + \mathbf{E}[R])^i \cdot s_t \cdot (1 + o(1)).$$

Using these two inequalities, as long as $(1 + \mathbf{E}[R])^k s_t = o(1)$, we have with probability

$$1 - o\left(\frac{k}{\log n}\right),$$

$$\begin{aligned} i_{t+k} &= i_t + \sum_{i=0}^{k-1} n_{t+i} \\ &\geq i_t + \sum_{i=0}^{k-1} s_{t+i} - \sum_{i=0}^{k-1} \left\{ 2s_{t+i}^2 + 2\sqrt{\frac{s_{t+i} \log \log n}{n}} \right\} \\ &\geq i_t + (1 - o(1)) \sum_{i=0}^{k-1} (1 + \mathbf{E}[R])^i s_t \\ &\quad - (2 + o(1)) \sum_{i=0}^{k-1} \left\{ (1 + \mathbf{E}[R])^{2i} s_t^2 + (1 + \mathbf{E}[R])^{i/2} \sqrt{\frac{s_t \log \log n}{n}} \right\} \\ &\geq i_t + f_2 \cdot (1 + \mathbf{E}[R])^k \cdot s_t - d \cdot \left((1 + \mathbf{E}[R])^{2k} s_t^2 + (1 + \mathbf{E}[R])^{k/2} \sqrt{\frac{s_t \log \log n}{n}} \right) \\ &\geq i_t + f_2 \cdot (1 + \mathbf{E}[R])^k \cdot s_t \cdot \left(1 - f \cdot (1 + \mathbf{E}[R])^k \cdot s_t - d \cdot (1 + \mathbf{E}[R])^{-k/2} \sqrt{\frac{\log \log n}{s_t n}} \right), \end{aligned} \tag{6}$$

where $f_2 > 0$ and $d > 0$ are constants. Since $\frac{\log^5 n}{n} \leq i_t \leq s_t$, we obtain that

$$i_{t+k} \geq i_t + f_2 \cdot (1 + \mathbf{E}[R])^k \cdot s_t \cdot (1 - o(1)). \tag{7}$$

By combining equations (7) and (6) we infer that with probability $1 - o\left(\frac{k}{\log n}\right)$,

$$i_t + f_2 \cdot (1 + \mathbf{E}[R])^k \cdot s_t \cdot (1 - o(1)) \leq i_{t+k} \leq i_t + f_1 \cdot (1 + \mathbf{E}[R])^k \cdot s_t \cdot (1 + o(1)).$$

□

Corollary 5.6. *With probability $1 - o(1)$ we have $|T_1 - \log_{1+\mathbf{E}[R]} n| = o(\log n)$.*

Proof. Applying Corollary 5.3 shows that with probability $1 - o(1)$, $T_0 = \mathcal{O}(\log \log n)$, where T_0 is the first round in which $\frac{\log^5 n}{n} \leq i_{T_0} \leq s_{T_0} \leq \frac{\log^{\mathcal{O}(1)} n}{n}$. Now we can apply Lemma 5.5 and set $k = \log_{1+\mathbf{E}[R]} n - o(\log n)$ such that with probability at least $1 - o(1)$ we have $\frac{1}{\log \log n} \leq i_{T_0+k} \leq \frac{A}{\log \log n}$, where $A > 1$ is a constant. Then we conclude that with probability $1 - o(1)$, $|T_1 - \log_{1+\mathbf{E}[R]} n| = o(\log n)$. □

5.3 The Final Phase

This phase starts with at least $\frac{n}{\log \log n}$ informed nodes and ends when all nodes get informed. Let T_1 be the first round in which $I_{T_1} \geq \frac{n}{\log \log n}$ and let T_2 be the first round in which all nodes are informed with probability $1 - o(1)$. We will show that with probability $1 - o(1)$, $|(T_2 - T_1) - \log_{e^{\mathbf{E}[R]}} n| = o(\log n)$.

Lemma 5.7. *With probability $1 - o(1)$,*

$$|(T_2 - T_1) - \log_{e^{\mathbf{E}[R]}} n| = o(\log n).$$

Proof. We define the indicator random variable Z_v for every $v \in \mathcal{U}_t$ and any round $t \geq T_1$:

$$Z_v = \begin{cases} 1 & \text{if } v \text{ does not get informed in round } t+1, \\ 0 & \text{otherwise.} \end{cases}$$

Thus,

$$\mathbf{E}[U_{t+1} | U_t] = \mathbf{E}\left[\sum_{v \in \mathcal{U}_t} Z_v\right] = U_t \cdot \mathbf{Pr}[Z_v = 1],$$

where for simplicity we omit the conditioning of U_{t+1} on U_t when dealing with the Z_v 's. Using the fact that $1 - \frac{1}{n} = e^{-\frac{1}{n} - \mathcal{O}(\frac{1}{n^2})}$, we can approximate the value $\mathbf{Pr}[Z_v = 1]$ as follows,

$$\begin{aligned} \mathbf{Pr}[Z_v = 1] &= \prod_{u \in \mathcal{I}_t} \left(1 - \frac{1}{n}\right)^{C_u} = \prod_{u \in \mathcal{I}_t} e^{-\frac{C_u}{n} - \mathcal{O}(\frac{C_u}{n^2})} \\ &= e^{-\sum_{u \in \mathcal{I}_t} (\frac{C_u}{n} + \mathcal{O}(\frac{C_u}{n^2}))} = e^{-s_t - \mathcal{O}(\frac{s_t}{n})}. \end{aligned}$$

Since $\frac{s_t}{n} = \mathcal{O}(\frac{1}{n})$ for any round and $e^{-\mathcal{O}(\frac{1}{n})} = 1 - \mathcal{O}(\frac{1}{n})$,

$$\mathbf{E}[U_{t+1} | U_t] = U_t e^{-s_t} \cdot e^{-\mathcal{O}(\frac{1}{n})} = U_t e^{-s_t} - \mathcal{O}\left(\frac{U_t}{n}\right). \quad (8)$$

Since for every $u, v \in \mathcal{U}_t$,

$$\mathbf{Pr}[Z_u = 1 \cap Z_v = 1] = \mathbf{Pr}[Z_u = 1 | Z_v = 1] \cdot \mathbf{Pr}[Z_v = 1] \leq \mathbf{Pr}[Z_v = 1] \cdot \mathbf{Pr}[Z_u = 1],$$

we have that

$$\mathbf{E}[Z_u \cdot Z_v] \leq \mathbf{E}[Z_u] \cdot \mathbf{E}[Z_v].$$

Therefore,

$$\begin{aligned} \mathbf{Var}\left[\sum_{v \in \mathcal{U}_t} Z_v\right] &= \sum_{v \in \mathcal{U}_t} \mathbf{E}[Z_v^2] + \sum_{u \neq v} (\mathbf{E}[Z_u \cdot Z_v] - \mathbf{E}[Z_u] \cdot \mathbf{E}[Z_v]) \\ &\leq \sum_{v \in \mathcal{U}_t} \mathbf{E}[Z_v^2] = U_t \cdot \mathbf{Pr}[Z_v = 1] = \mathbf{E}[U_{t+1} | U_t] \leq U_t. \end{aligned}$$

Applying Chebychev's inequality implies that with probability $1 - o(\frac{1}{\log n})$,

$$\left|U_{t+1} - \mathbf{E}[U_{t+1} | U_t]\right| \leq \sqrt{U_t \log^2 n}. \quad (9)$$

Combining inequalities (8) and (9) yields that with probability $1 - o(\frac{1}{\log n})$,

$$|U_{t+1} - U_t e^{-s_t}| \leq \sqrt{U_t \log^2 n} + \mathcal{O}\left(\frac{U_t}{n}\right) \leq 2\sqrt{U_t \log^2 n}. \quad (10)$$

According to the value of U_t , we consider two cases.

- Suppose that $U_t \geq \frac{n}{\log^5 n}$. Note that $s_t \geq i_t \geq \frac{1}{\log \log n}$ by the assumption of the lemma. Since s_t is a non-decreasing value in t and $U_t < n$ the recursive formula (10) implies that with probability $1 - o(\frac{1}{\log n})$,

$$U_{t+1} \leq U_t \cdot e^{\frac{-1}{\log \log n}} + 2\sqrt{n \log^2 n}.$$

Using an inductive argument shows that with probability $1 - o(\frac{k}{\log n})$,

$$U_{t+k} \leq U_t \cdot e^{\frac{-k}{\log \log n}} + \sum_{i=0}^{k-1} e^{\frac{-i}{\log \log n}} \cdot \left(2\sqrt{n \log^2 n}\right).$$

Hence after at most $k_0 = 6 \log \log^2 n$ rounds with probability $1 - o(1)$ the number of uninformed nodes decreases to $\frac{n}{\log^6 n} + \mathcal{O}(\sqrt{n \log^2 n})$, where $c > 0$ is a constant.

- Suppose that $U_t \leq \frac{n}{\log^5 n}$. If we set $n^\delta = \log^5 n$, then applying Proposition 5.1 implies that for any t for which $U_t = \mathcal{O}(\frac{n}{\log^5 n})$ with probability $1 - o(\frac{1}{\log n})$,

$$\sum_{u \in \mathcal{U}_t} C_u = o\left(\frac{n}{\log n}\right). \quad (11)$$

On the other hand, using Chebychev's inequality yields that with probability $1 - o(\frac{1}{\log n})$,

$$\left| \sum_{u \in \mathcal{V}} C_u - n \cdot \mathbf{E}[R] \right| \leq \sqrt{n \cdot \log^2 n}.$$

Combining the above equality and equality (11) results into an approximation for s_t which is not best possible but it suffices for our purpose. We know that

$$s_t = \sum_{u \in \mathcal{V}} C_u - \sum_{u \in \mathcal{U}_t} C_u.$$

So,

$$\mathbf{E}[R] - \sqrt{\frac{\log^2 n}{n}} - o\left(\frac{1}{\log n}\right) \leq s_t \leq \mathbf{E}[R] + \sqrt{\frac{\log^2 n}{n}}.$$

Therefore, s_t can be replaced by $\mathbf{E}[R] \pm o(\frac{1}{\log n})$ with probability $1 - o(\frac{1}{\log n})$. Inequality (10) implies that

$$\alpha \cdot U_t - 2\sqrt{U_t \log^2 n} \leq U_{t+1} \leq \alpha \cdot U_t + 2\sqrt{U_t \log^2 n}, \quad (12)$$

where $\alpha = e^{-\mathbf{E}[R] \pm o(1/\log n)}$. So as long as $U_t \geq \log^5 n$ with probability $1 - o(\frac{1}{\log n})$,

$$\begin{aligned} U_{t+1} &\leq \alpha \cdot U_t + 2\sqrt{U_t \log^2 n} = \alpha \cdot U_t \left(1 + 2\sqrt{\frac{\log^2 n}{\alpha^2 U_t}} \right) \\ &\leq \alpha \cdot U_t \left(1 + 2\sqrt{\frac{\log^2 n}{\alpha^2 \log^5 n}} \right) \leq \alpha \cdot U_t \left(1 + \frac{2}{\alpha \log^{\frac{3}{2}} n} \right). \end{aligned}$$

Now for any k for which $U_t e^{-k \mathbf{E}[R]} \geq \log^5 n$, with probability $1 - o(\frac{k}{\log n})$,

$$U_{t+k} \leq \alpha^k \cdot U_t \cdot \left(1 + \frac{2}{\alpha \log^{\frac{3}{2}} n} \right)^k = \alpha^k \cdot U_t \cdot (1 + o(1)). \quad (13)$$

In order to lower bound U_{t+k} we apply the lower bound (12) inductively. So we have that with probability $1 - o(\frac{k}{\log n})$,

$$U_{t+k} \geq \alpha^k \cdot U_t - \sum_{i=0}^{k-1} 2 \cdot \alpha^i \cdot \sqrt{U_{t+k-i-1} \log^2 n}.$$

Applying inequality (13) yields that with probability $1 - o(\frac{k}{\log n})$,

$$\sqrt{U_{t+k-i} \log^2 n} \leq \alpha^{\frac{k-i}{2}} \cdot \sqrt{U_t (1 + o(1)) \log^2 n}.$$

Thus,

$$\begin{aligned} U_{t+k} &\geq \alpha^k \cdot U_t - (1 + o(1)) \sum_{i=0}^{k-1} \alpha^{\frac{k-i-1}{2}} \cdot \sqrt{U_t \log^2 n} \\ &\geq \alpha^k \cdot U_t - c \cdot \alpha^{\frac{k}{2}} \cdot \sqrt{U_t \log^2 n}, \end{aligned} \quad (14)$$

where $c > 0$ is a constant and the last inequality holds because $\sum_{i=0}^{k-1} \alpha^{\frac{k-i-1}{2}} = \mathcal{O}(\alpha^{\frac{k}{2}})$. Combining the inequalities (13) and (14) yields for any k satisfying

$$U_t e^{-k \mathbf{E}[R]} \geq \log^5 n$$

with probability $1 - o(\frac{k}{\log n})$,

$$\alpha^k \cdot U_t(1 - o(1)) \leq U_{t+k} \leq \alpha^k \cdot U_t(1 + o(1)).$$

Hence by taking $k = \log_{e^{\mathbf{E}[R]}} n - o(\log n)$, with probability $1 - o(1)$, the number of uninformed nodes after $T_1 + k_0 + k$ rounds decreases to $\log^5 n$, so we have at most $\log^5 n$ uninformed nodes. Using the fact that for every $x \geq 0$, $1 - x \leq e^{-x}$, the probability that a node does not get informed after k_1 additional rounds is bounded from above by

$$\prod_{u \in \mathcal{I}_t} \left(1 - \frac{1}{n}\right)^{C_u \cdot k_1} \leq e^{-k_1 \sum_{u \in \mathcal{I}_t} C_u}.$$

We already know that $s_t = \mathbf{E}[R] \pm o\left(\frac{1}{\log n}\right)$ and s_t is a non-decreasing value in t so

$$\sum_{u \in \mathcal{I}_t} C_u = s_t > \frac{\mathbf{E}[R]}{2}.$$

Thus the union bound implies that the probability that every node in \mathcal{U}_t does not get informed is bounded by $\log^5 n \cdot e^{-\frac{k_1 \mathbf{E}[R]}{2}}$. By choosing $k_1 = \Theta(\log \log n)$ we conclude that with probability $1 - o(1)$ all nodes get informed. So we have with probability at least $1 - o(1)$ that $T_2 \leq T_1 + k_0 + k + k_1$, and $k_0 + k + k_1 = \log_{e^{\mathbf{E}[R]}} n + o(\log n)$. □

6 Push Protocol with Bounded Mean

This section is devoted to the proof of Theorem 1.2.

Proof. In the Push protocol, in round $t + 1$, at most S_t randomly chosen uninformed nodes are informed. This implies that $\mathbf{E}[S_{t+1} | S_t]$ increases by at most $\mathbf{E}[R] \cdot S_t$. Since the origin of the rumor is chosen without knowing C_u , $\mathbf{E}[S_0] = \mathbf{E}[R]$. Using the law of total expectation yields that

$$\mathbf{E}[S_t] = \mathbf{E}[\dots \mathbf{E}[\mathbf{E}[S_t | S_{t-1}] | S_{t-2}] \dots | S_0] \leq (1 + \mathbf{E}[R])^t \cdot \mathbf{E}[R].$$

By applying Markov's inequality, we conclude that

$$\Pr[I_t \geq n] \leq \Pr[S_t \geq n] \leq \frac{(1 + \mathbf{E}[R])^t \cdot \mathbf{E}[R]}{n}.$$

Hence $\Omega(\log n)$ rounds are necessary to inform all nodes with probability $1 - o(1)$. □

7 Lower Bound for Push-Pull

Before we present our results about the Push-Pull protocol we show the following general lemma. Recall that $S_0 = C_u$, where u is the single node that is aware of the rumor at the beginning of the protocol.

Lemma 7.1. Consider the Push-Pull protocol and $\{C_u : u \in \mathcal{V}\}$ be a sequence of positive integers. Then with probability $1 - o(1)$, the Push-Pull protocol needs at least

$$\Omega\left(\frac{\log n - \log S_0}{\log \sum_{u \in \mathcal{V}} C_u^2/n}\right)$$

rounds to inform all nodes.

Proof. We know the probability that an uninformed node u gets informed by Pull in round $t + 1$ is bounded by $I_t \cdot C_u/n$. Therefore using this bound we have

$$\begin{aligned} & \sum_{u \in \mathcal{U}_t} \mathbf{E}[C_u \mathbb{1}(u \text{ gets informed by Pull}) \mid S_t] \\ &= \sum_{u \in \mathcal{U}_t} C_u \cdot \Pr[u \text{ gets informed by Pull in round } t + 1] \\ &\leq \sum_{u \in \mathcal{U}_t} C_u \cdot \frac{I_t \cdot C_u}{n} \leq I_t \cdot \sum_{u \in \mathcal{V}} \frac{C_u^2}{n} \end{aligned}$$

On the other hand the probability that a node $u \in \mathcal{U}_t$ gets informed by Push in round $t + 1$ is at most S_t/n . So we get that

$$\begin{aligned} & \sum_{u \in \mathcal{U}_t} \mathbf{E}[C_u \mathbb{1}(u \text{ gets informed by Push}) \mid S_t] \\ &= \sum_{u \in \mathcal{U}_t} C_u \cdot \Pr[u \text{ gets informed by Push in round } t + 1] \\ &\leq \sum_{u \in \mathcal{U}_t} C_u \cdot \frac{S_t}{n} \leq S_t \cdot \sum_{u \in \mathcal{V}} \frac{C_u^2}{n}, \end{aligned}$$

where the last inequality follows by $C_u \leq C_u^2$. Combining the above inequalities implies that

$$\mathbf{E}[S_{t+1} \mid S_t] \leq S_t + (S_t + I_t) \cdot \left(\sum_{u \in \mathcal{V}} C_u^2/n\right) \leq \left(1 + 2 \cdot \sum_{u \in \mathcal{V}} C_u^2/n\right) \cdot S_t,$$

Applying the law of total expectation yields that

$$\mathbf{E}[S_t] = \mathbf{E}[\dots \mathbf{E}[\mathbf{E}[S_t \mid S_{t-1}] \mid S_{t-2}] \dots \mid S_0] \leq \left(1 + 2 \cdot \sum_{u \in \mathcal{V}} C_u^2/n\right)^t \cdot S_0.$$

Using Markov's inequality implies that

$$\Pr[I_t = n] \leq \Pr[S_t > n/2] \leq \frac{\mathbf{E}[S_t]}{n/2} \leq \frac{\left(1 + 2 \cdot \sum_{u \in \mathcal{V}} C_u^2/n\right)^t \cdot S_0}{n/2},$$

and the claim follows. □

8 Push-Pull Protocol with Bounded Mean and Bounded Variance

This section is devoted to the proof of Theorem 1.3.

Proof. $\{C_u : u \in \mathcal{V}\}$ be a sequence of positive integers each of which is generated independently according to some distribution R with $\mathbf{E}[R] = \mathcal{O}(1)$ and $\mathbf{Var}[R] = \mathcal{O}(1)$. We call $\{C_u : u \in \mathcal{V}\}$ a *good* sequence if $\sum_{u \in \mathcal{V}} C_u^2 = \mathcal{O}(n)$ and $S_0 = \mathcal{O}(1)$. Since the origin of the rumor is chosen without knowing C_u , $\mathbf{E}[S_0] = \mathbf{E}[R]$. Applying Markov's inequality implies that for any constant $\epsilon > 0$ with probability at least $1 - \epsilon/2$, $S_0 = \mathcal{O}(1)$. Since R is a probability distribution with bounded variance, $\sum_{u \in \mathcal{V}} \mathbf{E}[C_u^2] = \mathcal{O}(n)$. Another application of Markov's inequality implies that with probability $1 - \epsilon/2$, $\sum_{u \in \mathcal{V}} C_u^2 = \mathcal{O}(n)$. Therefore using a union bound for failure probability of two mentioned events implies that for fixed $\epsilon > 0$ with probability at least $1 - \epsilon$, $\{C_u : u \in \mathcal{V}\}$ is a good sequence. Conditioning on the event that $\{C_u : u \in \mathcal{V}\}$ is a good sequence, using Lemma 7.1 implies that with probability at least $1 - o(1)$ the Push-Pull protocol needs $\Omega(\log n)$ rounds to inform n nodes and the result follows. \square

9 Push-Pull Protocol with Power Law Distribution $2 < \beta < 3$

In this section we analyze the Push-Pull protocol where R is a power law distribution with $2 < \beta < 3$ and show that it only takes $\Theta(\log \log n)$ rounds to inform all with probability $1 - o(1)$. To prove the upper bound of $\mathcal{O}(\log \log n)$, we study the protocol in three consecutive phases and show that each phase takes only $\mathcal{O}(\log \log n)$ rounds. After that we show the lower bound $\Omega(\log \log n)$.

9.1 Proof of the Upper Bound

The Preliminary Phase.

This phase starts with just one informed node and ends when $I_t \geq n^{\frac{1}{\beta-1}} / (2 \cdot \log n)$. Let T_1 be the first round where $I_{T_1} \geq n^{\frac{1}{\beta-1}} / (2 \log n)$. We will show that $T_1 = \mathcal{O}(\log \log n)$. First we claim that $\mathcal{O}(\log \log n)$ rounds are sufficient to have $\log^{\mathcal{O}(1)} n$ informed nodes. Then we will show that in round $t + 1$ with probability $1 - e^{-\Omega(\log n)}$ there exists a node u with $C_u \geq I_t^{1+\gamma}$, $\gamma := \frac{3-\beta}{2(\beta-2)} > 0$, which pulls the rumor and consequently $S_{t+1} \geq I_t^{1+\gamma}$. Then considering only Push calls it follows that with probability $1 - o(\frac{1}{\log n})$,

$$I_{t+2} = I_{t+1} + N_{t+1} \geq I_{t+1} + S_{t+1}(1 - o(1)) > \frac{1}{2} I_t^{1+\gamma}.$$

So in every two rounds, I_t is increased by a factor of $\frac{1}{2} I_t^\gamma$ and hence after $\mathcal{O}(\log \log n)$ rounds the phase ends. For a complete proof see the following lemma.

Lemma 9.1. *With probability $1 - o(1)$, $T_1 = \mathcal{O}(\log \log n)$.*

Proof. At first we only consider Push calls and apply Lemma 4.1 which states that as long as $S_t \leq \log^{\frac{2}{3-\beta}} n$, with probability $1 - \mathcal{O}(\frac{\log^{\frac{4}{3-\beta}} n}{n})$,

$$I_{t+1} = I_t + S_t \geq 2I_t.$$

Thus as long as $S_t \leq \log^{\frac{2}{3-\beta}} n$, in each round the number of informed nodes is at least doubled. So we conclude that with probability $1 - o(1)$, $\mathcal{O}(\log \log n)$ rounds are sufficient to inform $\log^{\frac{2}{3-\beta}} n$ nodes. Let T_0 be the first round when $I_{T_0} \geq \log^{\frac{2}{3-\beta}} n$. Let us define the constant $\gamma := \frac{3-\beta}{2(\beta-2)} > 0$. Let T be the first round such that

$$I_{T-1}^{(1+\gamma)} \leq n^{\frac{1}{\beta-1}} / \log n < I_T^{(1+\gamma)}.$$

Now for any $T_0 \leq t \leq T$, we can apply Proposition 3.3 and conclude that with probability $1 - o(\frac{1}{n})$,

$$\sum_{u \in \mathcal{L}(I_t^{1+\gamma})} C_u \geq L(I_t^{1+\gamma}) \cdot I_t^{1+\gamma} \geq \frac{n \cdot c_1 \cdot I_t^{(1+\gamma)(2-\beta)}}{2}. \quad (15)$$

So,

$$\frac{I_t}{n} \sum_{u \in \mathcal{L}(I_t^{1+\gamma})} C_u \geq \frac{c_1 \cdot I_t^{1+(1+\gamma)(2-\beta)}}{2} = \frac{c_1 \cdot I_t^{3-\beta+\gamma(2-\beta)}}{2}.$$

We will bound the probability that none of $u \in \mathcal{L}(I_t^{1+\gamma})$ gets informed by Pull calls in round $t + 1$ as follows,

$$\prod_{u \in \mathcal{L}(I_t^{1+\gamma})} \left(1 - \frac{I_t}{n}\right)^{C_u} = \left(1 - \frac{I_t}{n}\right)^{\sum_{u \in \mathcal{L}(I_t^{1+\gamma})} C_u} \leq e^{-c_1 \cdot I_t^{3-\beta+\gamma(2-\beta)}} = e^{-c_1 \cdot I_t^{\frac{3-\beta}{2}}}.$$

Since for any $t \geq T_0$, $I_t \geq \log^{\frac{2}{3-\beta}} n$, we have that with probability at least $1 - n^{-c_1}$, at least one node in $\mathcal{L}(I_t^{1+\gamma})$ gets informed by Pull in round $t + 1$. Hence we have that

$$S_{t+1} \geq I_t^{1+\gamma}.$$

Let us now consider the Push calls in round $t + 2$. By applying Lemma 4.1 we know that as long as $S_{t+1} = o(n)$ with probability $1 - o(\frac{1}{\log n})$,

$$S_{t+1}(1 - o(1)) \leq N_{t+1}.$$

Thus,

$$I_{t+2} \geq I_{t+1} + S_{t+1}(1 - o(1)) > \frac{I_t^{1+\gamma}}{2}.$$

An inductive argument shows that for any integer $k \geq 1$ as long as $I_{T_0+2k-2}^{1+\gamma} \leq n^{\frac{1}{\beta-1}} / \log n$, with probability $1 - o(\frac{k}{\log n})$

$$I_{T_0+2k} > \left(\frac{1}{2}\right)^{\sum_{i=0}^{k-1} (1+\gamma)^i} I_{T_0}^{(1+\gamma)^k} = \left(\frac{I_{T_0}}{2^\gamma}\right)^{(1+\gamma)^k} \cdot 2^{1/\gamma} > \left(\frac{\log^{\frac{2}{3-\beta}} n}{C'}\right)^{(1+\gamma)^k},$$

where $C' = 2^\gamma = \mathcal{O}(1)$. So we conclude that after T_0+2k rounds, where $k = o(\log_{1+\gamma} \log n)$, there are two cases: either $I_{T_0+2k} \geq n^{\frac{1}{\beta-1}} / (2 \log n)$ which means $T_1 \leq T_0+2k = \mathcal{O}(\log \log n)$ and we are done, or

$$I_{T_0+2k} < n^{\frac{1}{\beta-1}} / (2 \log n) < n^{\frac{1}{\beta-1}} / \log n < I_{T_0+2k}^{1+\gamma}.$$

In the latter case, we change the value γ to γ' which satisfies $I_{T_0+2k}^{1+\gamma'} = n^{\frac{1}{\beta-1}} / \log n$ and a similar argument shows that

$$I_{T_0+2k+2} \geq n^{\frac{1}{\beta-1}} / (2 \log n).$$

□

The Middle Phase.

This phase starts with at least $n^{\frac{1}{\beta-1}} / (2 \log n)$ informed nodes and ends when $I_t \geq \frac{n}{\log n}$. Let T_2 be the first round in which $\frac{n}{\log n}$ nodes are informed. We will show that $T_2 - T_1 = \mathcal{O}(\log \log n)$. In contrast to the Preliminary Phase where we focus only on an informed node with maximal C_u , we now consider the number of informed nodes u with a C_u above a certain threshold Z_{t+1} which is inversely proportional to I_t .

Lemma 9.2. *Suppose that $I_t \geq n^{\frac{1}{\beta-1}} / (2 \log n)$ for some round t . Then with probability $1 - o(\frac{1}{n})$,*

$$|\mathcal{L}(Z_{t+1}) \cap \mathcal{I}_{t+1}| \geq \frac{1}{4} L(Z_{t+1}),$$

where $Z_{t+1} := \frac{n \log \log n}{I_t}$.

Proof. We consider two cases. If at least $\frac{1}{4}$ of the nodes in $\mathcal{L}(Z_{t+1})$ are already informed (before round $t+1$), then the statement of the lemma is true. Otherwise $|\mathcal{L}(Z_{t+1}) \cap \mathcal{U}_{t+1}| > \frac{3}{4} L(Z_{t+1})$. In the latter case, we define

$$\mathcal{L}'(Z_{t+1}) := \mathcal{L}(Z_{t+1}) \cap \mathcal{U}_{t+1}.$$

Let X_u be an indicator random variable for every $u \in \mathcal{L}'(Z_{t+1})$ so that

$$X_u := \begin{cases} 1 & \text{if } u \text{ gets informed by Pull in round } t+1, \\ 0 & \text{otherwise.} \end{cases}$$

Then we define a random variable X to be $X := \sum_{u \in \mathcal{L}'(Z_{t+1})} X_u$. Since for every $u \in \mathcal{L}'(Z_{t+1})$, $C_u \geq Z_{t+1} = \frac{n \log \log n}{I_t}$, it follows that

$$\Pr[X_u = 1] = 1 - \left(1 - \frac{I_t}{n}\right)^{C_u} \geq 1 - \left(1 - \frac{I_t}{n}\right)^{Z_{t+1}} = 1 - e^{-\Omega(\log \log n)} = 1 - o(1).$$

Thus $\Pr[X_u = 1] > \frac{3}{4}$ and $\mathbf{E}[X] = \sum_{u \in \mathcal{L}'(Z_{t+1})} \Pr[X_u = 1] > \frac{3}{4} |\mathcal{L}'(Z_{t+1})|$. Since

$$|\mathcal{L}'(Z_{t+1})| = |\mathcal{L}(Z_{t+1}) \cap \mathcal{U}_{t+1}| > \frac{3}{4} L(Z_{t+1}),$$

it follows that

$$\mathbf{E}[X] \geq \frac{9}{16} L(Z_{t+1}).$$

We know that $I_t \geq n^{\frac{1}{\beta-1}} / (2 \log n)$ and also I_t is a non-decreasing function in t , so

$$Z_{t+1} = \frac{n \log \log n}{I_t} \leq 2 \cdot n^{\frac{\beta-2}{\beta-1}} \log n \log \log n < n^{\frac{1}{\beta-1}} / \log n,$$

where the last inequality holds because $\beta < 3$. Now we can apply Proposition 3.3 (see appendix) to infer that with probability $1 - o(\frac{1}{n})$,

$$L(Z_{t+1}) \geq \frac{n \cdot c_1 \cdot Z_{t+1}^{1-\beta}}{2} \geq \frac{c_1 \cdot \log^{\beta-1} n}{2}.$$

Therefore,

$$\mathbf{E}[X] \geq \frac{9 \cdot c_1 \cdot \log^{\beta-1} n}{32}.$$

Then applying Theorem 2.1 results into

$$\Pr\left[X < \frac{\mathbf{E}[X]}{2}\right] \leq \Pr\left[|X - \mathbf{E}[X]| \geq \frac{\mathbf{E}[X]}{2}\right] < 2e^{-\frac{\mathbf{E}[X]}{10}} \leq 2e^{-\Omega(\log^{\beta-1} n)}. \quad (16)$$

So with probability $1 - o(\frac{1}{n})$, we have that

$$|\mathcal{L}(Z_{t+1}) \cap \mathcal{I}_{t+1}| \geq X \geq \frac{\mathbf{E}[X]}{2} > \frac{3|\mathcal{L}'(Z_{t+1})|}{8} \geq \frac{1}{4} L(Z_{t+1}),$$

where the last inequality holds because $|\mathcal{L}'(Z_{t+1})| > \frac{3}{4} L(Z_{t+1})$. □

Lemma 9.3. *With probability $1 - o(1)$, $T_2 - T_1 = \mathcal{O}(\log \log n)$.*

Proof. Since $I_t \geq n^{\frac{1}{\beta-1}} / (2 \log n)$, $Z_{t+1} = \frac{n \log \log n}{I_t} < n^{\frac{1}{\beta-1}} / \log n$, using Proposition 9.2 results into a lower bound for $|\mathcal{L}(Z_{t+1}) \cap \mathcal{I}_{t+1}|$. So with probability $1 - o(\frac{1}{n})$,

$$S_{t+1} = \sum_{u \in \mathcal{I}_{t+1}} C_u \geq |\mathcal{L}(Z_{t+1}) \cap \mathcal{I}_{t+1}| \cdot Z_{t+1} \geq \frac{1}{4} L(Z_{t+1}) \cdot Z_{t+1}.$$

By applying Proposition 3.3, we conclude that with probability $1 - o(\frac{1}{n})$, $L(Z_{t+1}) \geq \frac{n \cdot c_1 \cdot Z_{t+1}^{1-\beta}}{2}$. Therefore, with probability $1 - o(\frac{1}{n})$,

$$S_{t+1} \geq \frac{n \cdot c_1 \cdot Z_{t+1}^{2-\beta}}{8}.$$

As long as $S_{t+1} = o(n)$, we can apply Lemma 4.2 for the Push protocol to round $t + 2$ implying that with probability $1 - o(\frac{1}{\log n})$,

$$I_{t+2} = I_{t+1} + N_t \geq I_{t+1} + S_{t+1}(1 - o(1)).$$

Thus,

$$I_{t+2} > \frac{S_{t+1}}{2} \geq \frac{c_1}{16} n \cdot Z_{t+1}^{2-\beta} = \frac{c_1}{16} \cdot n^{3-\beta} \cdot \log \log^{2-\beta} n \cdot I_t^{\beta-2}.$$

By an inductive argument, we obtain that for any integer $k \geq 1$ with $S_{t+k} = o(n)$, it holds with probability $1 - o(\frac{k}{\log n})$,

$$I_{t+2k} > \left(\frac{c}{16} n^{3-\beta} \cdot \log \log^{2-\beta} n \right)^{\sum_{i=0}^{k-1} (\beta-2)^i} I_t^{(\beta-2)^k} = \left(\frac{c}{16} n^{3-\beta} \cdot \log \log^{2-\beta} n \right)^{\frac{1-(\beta-2)^k}{3-\beta}} I_t^{(\beta-2)^k}.$$

Therefore there exists $k = \mathcal{O}(\log_{\frac{1}{\beta-2}} \log n)$ such that

$$\begin{aligned} I_{t+2k} &\geq \left(\frac{c}{16} n^{3-\beta} \cdot \log \log^{2-\beta} n \right)^{\frac{1-\mathcal{O}(1/\log n)}{3-\beta}} I_t^{1/\log n} \\ &= \Omega \left(n^{1-\mathcal{O}(1/\log n)} \left(\frac{c}{16} \cdot \log \log^{2-\beta} n \right)^{\frac{1-\mathcal{O}(1/\log n)}{3-\beta}} \right) = \Omega \left(\frac{n}{\log \log^\delta n} \right), \end{aligned}$$

where $\delta = \frac{\beta-2}{3-\beta}(1 - \mathcal{O}(1/\log n)) > 0$. Hence $T_2 \leq T_1 + 2k = T_1 + \mathcal{O}(\log \log n)$ with probability $1 - o(1)$. \square

The Final Phase.

This phase starts with at least $\frac{n}{\log n}$ informed nodes. Since the runtime of our Push-Pull protocol is stochastically smaller than the runtime of the standard Push-Pull protocol (i.e. $C_u = 1$ for every $u \in V$), we simply use the result by Karp et. al in [24, Theorem 2.1] for the standard Push-Pull protocol which states that once $I_t \geq \frac{n}{\log n}$, additional $\mathcal{O}(\log \log n)$ rounds are with probability $1 - o(1)$ sufficient to inform all n nodes. \square

9.2 Proof of the Lower Bound

Since increasing the number of informed nodes can only decrease the runtime of the protocol, we may assume that at the beginning there are $\log^b n$ random informed nodes, where $b := \max\{4, 2 + \frac{3(3-\beta)}{\beta-2}\}$. Applying Markov's inequality to the random variable S_0

implies that with probability $1 - o(\frac{1}{\log n})$, $\log^b n \leq S_0 \leq \log^{2+b} n$. In the following we lower bound the number of rounds to reach $n^{\frac{1}{\log \log n}}$ informed nodes. We do this by keeping track of the largest value of C_u among all informed nodes and show that this value does not exceed $I_t^{\frac{1}{\beta-2}} \log^{\frac{3}{\beta-2}} n$ with high probability.

By Fact 3.2, with probability $1 - o(\frac{1}{\log n})$ we have $\max_{u \in \mathcal{V}} C_u \leq n^{\frac{1}{\beta-1}} \log n$. Let i^* be the smallest positive integer so that $2^{i^*} \geq n^{\frac{1}{\beta-1}} / \log n$. Then $i^* < \log n$. Let us define the set $\mathcal{M}_i := \{u \in \mathcal{V} : 2^{i-1} \leq C_u < 2^i\}$ for $1 \leq i \leq i^* - 1$ and $\mathcal{M}_{i^*} := \{u \in \mathcal{V} : 2^{i^*-1} \leq C_u \leq n^{\frac{1}{\beta-1}} \log n\}$. We denote the size of \mathcal{M}_i with M_i . By definition, for any $1 \leq i \leq i^*$, $M_i \leq L(2^{i-1})$. Applying Proposition 3.3 implies that with probability $1 - o(\frac{1}{n})$ for any $1 \leq i \leq i^*$ we have $M_i \leq \frac{3}{2} \cdot c_2 \cdot n \cdot 2^{(i-1)(1-\beta)}$. Let us define the indicator random variable Z_u^i for every $u \in \mathcal{U}_t \cap \mathcal{M}_i$ as follows:

$$Z_u^i := \begin{cases} 1 & \text{if } u \text{ gets informed by Pull in round } t+1, \\ 0 & \text{otherwise.} \end{cases}$$

Hence, $\Pr[Z_u^i = 1] \leq C_u \cdot \frac{I_t}{n} \leq \frac{I_t \cdot 2^i}{n}$. Let P_i be the probability that at least one node in $\mathcal{U}_t \cap \mathcal{M}_i$ gets informed by Pull in round $t + 1$. Then, for any $1 \leq i \leq i^* - 1$,

$$P_i \leq \sum_{u \in \mathcal{U}_t \cap \mathcal{M}_i} \Pr[Z_u^i = 1] \leq M_i \cdot \frac{I_t}{n} \cdot 2^i \leq 3 \cdot c_2 \cdot I_t \cdot 2^{(i-1)(2-\beta)}.$$

Since $2^{i^*} \geq n^{\frac{1}{\beta-1}} / \log n$ and $C_u \leq n^{\frac{1}{\beta-1}} \log n$ with probability $1 - o(\frac{1}{\log n})$,

$$\begin{aligned} P_{i^*} &\leq \sum_{u \in \mathcal{U}_t \cap \mathcal{M}_{i^*}} \Pr[Z_u^{i^*} = 1] \leq \frac{3}{2} \cdot c_2 \cdot n \cdot 2^{(i^*-1)(1-\beta)} \cdot \frac{I_t}{n} \cdot n^{\frac{1}{\beta-1}} \log n \\ &\leq 6 \cdot c_2 \cdot I_t \cdot n^{\frac{2-\beta}{\beta-1}} \log^{\beta-1} n. \end{aligned}$$

So as long as $I_t \leq n^{\frac{1}{\log \log n}}$, $P_{i^*} = o(\frac{1}{\log^3 n})$. We define $\Delta_t := S_t^{\frac{1}{\beta-2}} \log^{\frac{3}{\beta-2}} n$. Let i_t be the smallest integer so that $2^{i_t} \geq \Delta_t$. Then for any $i_t \leq i \leq i^*$ we have,

$$P_i \leq 3 \cdot c_2 \cdot 2^{\beta-2} \cdot I_t \cdot 2^{i(2-\beta)} \leq 6 \cdot c_2 \cdot I_t \cdot \Delta_t^{2-\beta} \leq 6 \cdot c_2 \cdot \log^{-3} n.$$

Let E_t be the event that no node with $C_u \geq \Delta_t$ gets informed by Pull in round $t + 1$. Then we have

$$\Pr[E_t] \geq 1 - \sum_{i=i_t}^{i^*} P_i \geq 1 - o\left(\frac{1}{\log n}\right). \quad (17)$$

Let us define $S_{t+1}^{(1)} := \sum_{u \in \mathcal{N}_t^{\text{Pull}}} C_u$. Conditioning on the event E_t we obtain that

$$\begin{aligned} \mathbf{E} \left[S_{t+1}^{(1)} \mid S_t \right] &\leq \sum_{i=1}^{i_t} \sum_{u \in \mathcal{U}_t \cap \mathcal{M}_i} C_u \cdot \frac{\Pr[Z_u^i = 1]}{\Pr[E_t]} \\ &\leq (1 + o(1)) \cdot \sum_{i=1}^{i_t} 2^i \cdot M_i \cdot \frac{I_t}{n} \cdot 2^i \leq (1 + o(1)) \cdot \frac{S_t}{n} \cdot \sum_{i=1}^{i_t} 2^i \cdot M_i \cdot 2^i. \end{aligned}$$

By definition of i_t , we have $2^{i_t} \leq 2 \cdot \Delta_t$ and also we have $M_i \leq L(2^{i-1}) \leq \frac{3}{2} \cdot c_2 \cdot 2^{(i-1)(1-\beta)} \cdot n$. Hence the last sum is bounded by

$$\begin{aligned} (1 + o(1)) \cdot \sum_{i=1}^{i_t} 2^{2i} \cdot I_t \cdot 2^{(i-1)(1-\beta)} &\leq 24 \cdot c_2 \cdot I_t \cdot 2^{i_t(3-\beta)} \leq 24 \cdot c_2 \cdot I_t \cdot (2 \cdot \Delta_t)^{3-\beta} \\ &\leq 24 \cdot c_2 \cdot S_t^{1+\frac{3-\beta}{\beta-2}} \log^{\frac{3(3-\beta)}{\beta-2}} n. \end{aligned}$$

Conditioning on the event E_t and applying Markov's inequality imply that with probability $1 - o(\frac{1}{\log n})$,

$$S_{t+1}^{(1)} \leq \log^2 n \cdot \mathbf{E} \left[S_{t+1}^{(1)} \mid S_t \right] \leq 24 \cdot c_2 \cdot S_t^{1+\frac{3-\beta}{\beta-2}} \log^{2+\frac{3(3-\beta)}{\beta-2}} n. \quad (18)$$

Let us define the indicator random variable Y_u for every $u \in \mathcal{U}_t$ as follows:

$$Y_u := \begin{cases} 1 & \text{if } u \text{ gets informed by Push in round } t+1, \\ 0 & \text{otherwise.} \end{cases}$$

Then we have $\Pr[Y_u = 1] \leq S_t/n$. Let A denote the event that $\sum_{u \in \mathcal{V}} C_u \leq n \cdot \log^2 n$. Since $\mathbf{E}[R] = \mathcal{O}(1)$, applying Markov's inequality implies that $\Pr[A] \geq 1 - o(1/\log n)$. Let us define $S_{t+1}^{(2)} := \sum_{u \in \mathcal{N}_t^{\text{Push}}} C_u$. Conditioning on the event A we have

$$\mathbf{E} \left[S_{t+1}^{(2)} \mid S_t \right] = \sum_{u \in \mathcal{U}_t} C_u \cdot \frac{\Pr[Y_u = 1]}{\Pr[A]} \leq (1 + o(1)) \cdot \sum_{u \in \mathcal{V}} C_u \cdot \frac{S_t}{n} \leq (1 + o(1)) \cdot S_t \cdot \log^2 n.$$

Conditioning on A and applying Markov's inequality implies that with probability $1 - o(1/\log n)$,

$$S_{t+1}^{(2)} \leq \log^2 n \cdot \mathbf{E} \left[S_{t+1}^{(2)} \mid S_t \right] \leq S_t \cdot \log^4 n \quad (19)$$

Combining inequalities (18) and (19) implies that with probability $1 - o(1/\log n)$ for every $0 \leq t \leq \log \log n$

$$\begin{aligned} S_{t+1} &\leq S_t + S_{t+1}^{(1)} + S_{t+1}^{(2)} \leq S_t + 24 \cdot c_2 \cdot S_t^{1+\frac{3-\beta}{\beta-2}} \log^{2+\frac{3(3-\beta)}{\beta-2}} n + S_t \cdot \log^4 n \\ &\leq S_t + 24 \cdot c_2 \cdot S_t^{b+1} + S_t^2 \leq S_t^{b+2}, \end{aligned}$$

where the last inequality holds because $b = \max\{4, 2 + \frac{3(3-\beta)}{\beta-2}\}$ and $\log^b n \leq I_t \leq S_t$. We know that with probability $1 - o(1/\log n)$ we have $S_0 \leq \log^{b+2} n$. An inductive argument shows that for every $1 \leq t \leq \log \log n$ with probability $1 - o(1)$, $S_t \leq S_0^{(b+2)^t} \leq \log^{(b+2)^{t+1}}$. If we set $T := \frac{1}{2} \cdot \log_{b+2} \log n$, then with probability $1 - o(1)$ we have $S_T < n^{1/\log \log n}$. Thus $T = \Omega(\log \log n)$ rounds are necessary to inform all nodes with probability $1 - o(1)$. This finishes the proof of the lower bound of $\Omega(\log \log n)$. \square

10 Push-Pull Protocol with Power Law Distribution $\beta = 3$

In this section we analyse the Push-Pull protocol where R is a power law distribution with $\beta = 3$ and show that the Push-Pull protocol takes $\Theta\left(\frac{\log n}{\log \log n}\right)$ rounds to inform all n nodes. Throughout this section we assume that the power law distribution with $\beta = 3$ has an additional property in which for every positive integer z

$$\Pr[R = z] \geq c \cdot z^{-3}, \quad (20)$$

where $c > 0$ is fixed. Let us define $\mathcal{L}^*(z) = \{u : C_u = z\}$ and $L^*(z) = |\mathcal{L}^*(z)|$. Also we define $\mathcal{I}_t(z) = \mathcal{I}_t \cap \mathcal{L}^*(z)$ and $\mathcal{N}_t(z) = \mathcal{N}_t \cap \mathcal{L}^*(z)$, whose sizes are denoted by $I_t(z)$ and $N_t(z)$ respectively. $N_t^{\text{Push}}(z)$ and $N_t^{\text{Pull}}(z)$ are denoted the size of the newly informed nodes with $C_u = z$ by Push and Pull transmissions respectively. In the following we show a useful fact about the $L^*(z)$.

Fact 10.1. *Let R be a power law distribution with $\beta = 3$. Then for every $z = \mathcal{O}(n^{1/4})$, with probability $1 - o(\frac{1}{n})$ we have that*

$$\frac{n \cdot \Pr[R = z]}{2} \leq L^*(z) \leq \frac{3 \cdot n \cdot \Pr[R = z]}{2}.$$

Proof. We know that $\mathbf{E}[L^*(z)] = n \cdot \Pr[R = z]$. By using the inequality (20) we have that for any $z = \mathcal{O}(n^{1/4})$, $\Pr[R = z] = \Omega(n^{-3/4})$. Then we have that $\mathbf{E}[L^*(z)] = \Omega(n^{2/5})$ and using a Chernoff bound 2.1 yields that with probability $1 - o(\frac{1}{n})$ the inequality in the statement holds. \square

10.1 Proof of Lower Bound

Theorem 10.2. *With probability $1 - o(1)$, the Push-Pull needs at least $\Omega\left(\frac{\log n}{\log \log n}\right)$ rounds to inform all n nodes.*

Proof. Let $\{C_u : u \in \mathcal{V}\}$ be a sequence of positive integers where every C_u is generated independently according to a power law distribution with $\beta = 3$. We call a sequence $\{C_u : u \in \mathcal{V}\}$ is good if it fulfills three conditions:

1. For every $u \in \mathcal{V}$, $C_u < n$.
2. $S_0 = \mathcal{O}(\log n)$.
3. $\sum_{u \in \mathcal{V}} \frac{C_u^2}{n} = \mathcal{O}(\log^2 n)$.

In the following we show that with probability $1 - o(1)$ every sequence $\{C_u, u \in \mathcal{V}\}$ is good. By definition of power law distribution for $\beta = 3$ we have that

$$\Pr[C_u \leq n] > 1 - \frac{c_1}{n^2} = 1 - o(1).$$

We know that $\mathbf{E}[R] = \mathcal{O}(1)$, so Markov's inequality implies that with probability $1 - \mathcal{O}(\frac{1}{\log n})$, $S_0 = \mathcal{O}(\log n)$. Conditioning on the event that for every $u \in \mathcal{V}$, $C_u < n$ we get

$$\mathbf{E}[C_u^2 | C_u \leq n] \leq \frac{\sum_{z=1}^n \mathbf{Pr}[R^2 \geq z]}{\mathbf{Pr}[C_u \leq n]} \leq (1 + o(1)) \cdot c_1 \sum_{z=1}^n \frac{1}{z} = (1 + o(1)) \cdot c_1 \cdot \log n.$$

So applying Markov's inequality yields that with probability $1 - \mathcal{O}(\frac{1}{\log n})$,

$$\sum_{u \in \mathcal{V}} \frac{C_u^2}{n} = \mathcal{O}(\log^2 n).$$

Therefore we have that with probability $1 - o(1)$, the sequence $\{C_u : u \in \mathcal{V}\}$ is good. Conditioning on this event and then applying Lemma 7.1 shows that with probability $1 - o(1)$ the Push-Pull needs at least

$$\Omega\left(\frac{\log n - \log S_0}{\log \sum_{u \in \mathcal{V}} C_u^2/n}\right) = \Omega\left(\frac{\log n}{\log \log n}\right)$$

rounds to inform n nodes. □

10.2 Proof of Upper Bound

Before we present a proof for the upper bound we show following two lemmas.

Lemma 10.3. *Suppose that $S_t \leq \frac{n}{\log^6 n}$ and $z \leq \min\{\frac{n}{I_t \cdot \log^6 n}, \mathcal{O}(n^{\frac{1}{4}})\}$. Then with probability $1 - o(\frac{1}{\log n})$, for any round $t = \mathcal{O}(\log n)$ we have that*

$$|\mathcal{U}_t(z) \cap \mathcal{L}^*(z)| \geq \frac{L^*(z)}{2} \geq \frac{n \cdot \mathbf{Pr}[R = z]}{4}.$$

Proof. By considering the Push call we have that the size of newly informed nodes is bounded by S_t . Since they are chosen randomly, we have that

$$\mathbf{E}[N_t^{\text{Push}}(z) | S_t] \leq S_t \cdot \mathbf{Pr}[R = z]. \tag{21}$$

On the other hand we have that

$$\begin{aligned} \mathbf{E}[N_t^{\text{Pull}}(z) | I_t] &\leq \sum_{u \in \mathcal{U}_t \cap \mathcal{L}^*(z)} \mathbf{Pr}[u \text{ gets informed by Pull in round } t + 1] \\ &\leq L^*(z) \cdot \mathbf{Pr}[u \text{ gets informed by Pull in round } t + 1], \quad \text{since } |\mathcal{U}_t \cap \mathcal{L}^*(z)| \leq L^*(z). \\ &= L^*(z) \cdot \left(1 - \left(1 - \frac{I_t}{n}\right)^z\right) \leq L^*(z) \cdot \frac{2 \cdot I_t \cdot z}{n} \end{aligned}$$

where the last inequality holds because we assume that $\frac{I_t}{n} \leq \frac{S_t}{n} < \frac{1}{2}$ and for any $0 \leq a \leq \frac{\log 2}{2}$, $e^{-2a} \leq 1 - a \leq e^{-a}$. Applying Fact 10.1 shows that for any $z = \mathcal{O}(n^{\frac{1}{4}})$ with probability $1 - o(\frac{1}{n})$ we have

$$\frac{n \cdot \mathbf{Pr}[R = z]}{2} \leq L^*(z) \leq \frac{3 \cdot n \cdot \mathbf{Pr}[R = z]}{2}. \tag{22}$$

Thus,

$$\mathbf{E} [N_t^{\text{Pull}}(z)|S_t] \leq 3 \cdot I_t \cdot z \cdot \mathbf{Pr} [R = z]. \quad (23)$$

Combining (21) and (23) implies that

$$\mathbf{E} [N_t(z)|S_t, I_t] \leq S_t \cdot \mathbf{Pr} [R = z] + 3 \cdot I_t \cdot z \cdot \mathbf{Pr} [R = z]$$

We know that $I_{t+1}(z) = I_0(z) + \sum_{i=1}^t N_i(z)$. Using the linearity of expectation we have that

$$\begin{aligned} \mathbf{E} [I_{t+1}(z)|S_i, I_i, 0 \leq i \leq t] &= I_0(z) + \sum_{i=0}^t \mathbf{E} [N_i(z)|S_i, I_i] \\ &\leq I_0(z) + \mathbf{Pr} [R = z] \cdot \sum_{i=0}^t (S_i + 3 \cdot I_i \cdot z) \\ &\leq 1 + \mathbf{Pr} [R = z] \cdot (t + 1) \cdot (S_t + z \cdot 3 \cdot I_t), \end{aligned}$$

where the last inequality comes from the fact that S_i and I_i are non-decreasing function in t . By assumption $z \leq \min\{\frac{n}{I_t \cdot \log^6 n}, \mathcal{O}(n^{\frac{1}{4}})\}$ and $S_t \leq \frac{n}{\log^6 n}$, for any round $t = \mathcal{O}(\log n)$ we have that

$$\mathbf{E} [I_{t+1}(z)|S_i, I_i, 1 \leq i \leq t] \leq 2 \cdot (t + 1) \cdot (S_t + 3 \cdot I_t \cdot z) \cdot \mathbf{Pr} [R = z] \leq \frac{n \cdot \mathbf{Pr} [R = z]}{\log^4 n}.$$

Applying Markov's inequality shows that with probability $1 - o(\frac{1}{\log n})$ for any round $t = \mathcal{O}(\log n)$,

$$I_{t+1}(z) \leq \log^2 n \cdot \mathbf{E} [I_{t+1}(z)|S_i, I_i, 0 \leq i \leq t] \leq \frac{n \cdot \mathbf{Pr} [R = z]}{\log^2 n} \leq \frac{L^*(z)}{2},$$

where the last inequality follows from inequality (22). Therefore we infer that with probability $1 - o(\frac{1}{\log n})$, $|\mathcal{U}_t(z) \cap \mathcal{L}^*(z)| \geq \frac{L^*(z)}{2}$. \square

Lemma 10.4. *Suppose that $I_t = e^{\Omega(\frac{\log n}{\log \log n})}$ and $S_t \leq \frac{n}{\log^6 n}$. Then with probability $1 - o(1)$, the Push-Pull protocol needs $\mathcal{O}\left(\frac{\log n}{\log \log n}\right)$ rounds to inform at least $e^{\log n - \frac{\log n}{\log \log n}}$ nodes.*

Proof. Let X_u be an indicator random variable for every $u \in \mathcal{U}_t(z) \cap \mathcal{L}^*(z)$ so that

$$X_u := \begin{cases} 1 & \text{if } u \text{ gets informed by Pull in round } t + 1, \\ 0 & \text{otherwise.} \end{cases}$$

Then we define the random variable $X_t(z) := \sum_{u \in \mathcal{U}_t(z) \cap \mathcal{L}^*(z)} X_u$. Let us define $z_t = \min\{I_t^{1/4}, \frac{n}{I_t \cdot \log^6 n}\}$. Using the approximation $e^{-2 \cdot a} \leq 1 - a \leq e^{-a}$, $0 \leq a \leq 1/2$, results that for any $z \leq z_t$ we have

$$\mathbf{Pr} [X_u = 1] = 1 - \left(1 - \frac{I_t}{n}\right)^z \geq 1 - e^{-\frac{I_t \cdot z}{n}} \geq \frac{I_t \cdot z}{2 \cdot n},$$

Applying Lemma 10.3 shows that with probability $1 - o(\frac{1}{\log n})$ for any $z \leq z_t$ and any round $t = \mathcal{O}(\log n)$,

$$\mathbf{E}[X_t(z)] = \sum_{u \in \mathcal{U}_t(z) \cap \mathcal{L}^*(z)} \Pr[X_u = 1] > \frac{L^*(z) \cdot I_t \cdot z}{4 \cdot n} \geq \frac{I_t \cdot z \cdot \Pr[R = z]}{8} \geq \frac{c \cdot I_t}{I_t^{\frac{3}{4}}}, \quad (24)$$

where the last inequality holds because $\Pr[R = z] \geq \frac{c}{z^3}$. Since $I_t = e^{\Omega(\frac{\log n}{\log \log n})}$ and X_u 's are independent, applying a Chernoff bound 2.1 implies that with probability $1 - o(\frac{1}{n})$,

$$X_t(z) \geq \frac{\mathbf{E}[X_t(z)]}{2}.$$

Using the above inequality and inequality (24) shows that that with probability $1 - o(\frac{1}{\log n})$ there exists a constant C so that

$$S_{t+1} \geq \sum_{z=1}^{z_t} X_t(z) \cdot z \geq \frac{I_t}{16} \sum_{z=1}^{z_t} z^2 \cdot \Pr[R = z] \geq \frac{c \cdot I_t}{16} \sum_{z=1}^{z_t} \frac{1}{z} = I_t \cdot C \cdot \log z_t.$$

For any positive integer k in which $I_{t+k} \in [e^{\Omega(\frac{\log n}{\log \log n})}, e^{\log n - \frac{\log n}{\log \log n}}]$, we have that

$$e^{\Omega(\frac{\log n}{\log \log n})} \leq z_t.$$

Hence from the above inequality we conclude that here exists a constant C_1 so that

$$S_{t+1} \geq C_1 \cdot I_t \cdot \frac{\log n}{\log \log n} \geq C_1 \cdot I_t \cdot \sqrt{\log n}.$$

Considering only Push transmission for $S_t = o(n)$ and applying Lemma 4.2 implies that with probability $1 - o(\frac{1}{\log n})$

$$I_{t+2} \geq \frac{S_{t+1}}{2} \geq \frac{C_1 \cdot I_t \cdot \sqrt{\log n}}{2}$$

An inductive argument shows that for any integer k as long as $S_{t+2k} = \frac{n}{\log^6 n}$ with probability $1 - o(1)$,

$$I_{t+2k} \geq I_t \cdot \left(\frac{C_1 \cdot \sqrt{\log n}}{2} \right)^k.$$

Thus there is a $k = \mathcal{O}\left(\frac{\log n}{\log \log n}\right)$ so that after $t + 2k$ rounds there are at least $e^{\log n - \frac{\log n}{\log \log n}}$ informed nodes. □

Corollary 10.5. *let R be a power law distribution with $\beta = 3$. Then with probability $1 - o(1)$, the Push-Pull protocol informs all n nodes in $\mathcal{O}\left(\frac{\log n}{\log \log n}\right)$ rounds.*

Proof. Applying Corollary 4.4 results that as long as $S_t = o(n)$ with probability $1 - o(1)$, for any round $t = \mathcal{O}(\log n)$,

$$I_t \geq \left(\frac{3}{2}\right)^t \cdot I_0.$$

So after $\mathcal{O}\left(\frac{\log n}{\log \log n}\right)$ rounds there are at least $e^{\Omega\left(\frac{\log n}{\log \log n}\right)}$ informed nodes. Now we apply Lemma 10.4 and conclude that after $\mathcal{O}\left(\frac{\log n}{\log \log n}\right)$ rounds we have at least $e^{\log n - \frac{\log n}{\log \log n}}$ informed nodes. Another application of Corollary 4.4 implies that after $\mathcal{O}\left(\frac{\log n}{\log \log n}\right)$ rounds we have at least $\frac{n}{\log \log n}$ informed nodes. Since we have enough number of informed nodes using the result by Karp et. al in [24, Theorem 2.1] for the standard Push-Pull protocol shows that once $I_t \geq \frac{n}{\log n}$, with probability $1 - o(1)$ additional $\mathcal{O}(\log \log n)$ rounds are sufficient to inform all n nodes. \square

11 Generating a New C_u^t in Each Round (Theorem 1.6)

In this section we analysis the Push-Pull protocol for a new model. In this model according to some distribution R , at the beginning of each round t , every node u generates a random natural number $C_u^t \geq 1$ independent of all other nodes. Then in round t , the Push-Pull protocol disseminates the information according to $\{C_u^t: u \in \mathcal{V}\}$, i.e., node u calls C_u^t random nodes. In the following we show that if we have $\mathbf{E}[R] = \mathcal{O}(1)$. Then with probability $1 - o(1)$, the Push-Pull protocol needs $\Omega(\log n)$ rounds to inform all nodes.

Proof. The probability that a node $u \in \mathcal{U}_t$ gets informed by Pull is as follows:

$$\begin{aligned} & \Pr [u \text{ gets informed by Pull in round } t + 1] \\ &= \sum_{x=1}^{\infty} \Pr [u \text{ gets informed by Pull in round } t + 1 \mid R_u^{t+1} = x] \cdot \Pr [R_u^{t+1} = x] \\ &= \sum_{x=1}^{\lfloor \frac{n}{2I_t} \rfloor} \left(1 - \left(1 - \frac{I_t}{n}\right)^x\right) \cdot \Pr [R_u^{t+1} = x] + \sum_{x=\lfloor \frac{n}{2I_t} \rfloor + 1}^{\infty} \left(1 - \left(1 - \frac{I_t}{n}\right)^x\right) \cdot \Pr [R_u^{t+1} = x] \\ &\leq \frac{I_t}{n} \sum_{x=1}^{\lfloor \frac{n}{2I_t} \rfloor} x \cdot \Pr [R_u^{t+1} = x] + \sum_{x=\lfloor \frac{n}{2I_t} \rfloor + 1}^{\infty} \Pr [R_u^{t+1} = x] \quad \left(\text{since } 1 - \left(1 - \frac{I_t}{n}\right)^x \leq \frac{I_t \cdot x}{n}\right) \\ &\leq \frac{I_t}{n} \cdot \mathbf{E}[R] + \Pr \left[R_u^{t+1} > \left\lfloor \frac{n}{2I_t} \right\rfloor \right] \leq \frac{I_t}{n} \cdot \mathbf{E}[R] + \frac{2I_t}{n} \cdot \mathbf{E}[R], \end{aligned}$$

where the last inequality is due to Markov's inequality. Recall that N_t^{Pull} and N_t^{Push} are the number of newly informed nodes by Pull and Push calls in round $t + 1$ respectively.

Hence,

$$\begin{aligned} \mathbf{E} [N_t^{\text{Pull}} | I_t] &= \sum_{u \in \mathcal{U}_t} \Pr [u \text{ gets informed by Pull in round } t + 1] \\ &\leq \frac{U_t \cdot I_t \cdot 3 \cdot \mathbf{E} [R]}{n} < 3 \cdot I_t \cdot \mathbf{E} [R]. \end{aligned}$$

Recall that S_t is the number of Push calls by informed nodes in round $t + 1$. Therefore, $N_t^{\text{Push}} \leq S_t$ and

$$\mathbf{E} [N_t^{\text{Push}} | I_t] \leq \mathbf{E} [S_t | I_t] = \sum_{u \in I_t} \mathbf{E} [C_u^{t+1}] = I_t \cdot \mathbf{E} [R].$$

Hence,

$$\mathbf{E} [I_{t+1} | I_t] \leq I_t + \mathbf{E} [N_t^{\text{Pull}} | I_t] + \mathbf{E} [N_t^{\text{Push}} | I_t] \leq (1 + 4 \cdot \mathbf{E} [R]) \cdot I_t.$$

By using the law of total expectation, we conclude that $\mathbf{E} [I_t] < (1 + 4 \cdot \mathbf{E} [R])^t$. If we set $T = c \cdot \log n$, where $c > 0$ is a small constant, then

$$\Pr [I_T \geq \sqrt{n}] \leq \frac{\mathbf{E} [I_T]}{\sqrt{n}} \leq \frac{(1 + 4 \cdot \mathbf{E} [R])^T}{\sqrt{n}} = o(1).$$

So with probability $1 - o(1)$, we need at least $c \cdot \log n$ rounds to inform all nodes. \square

References

- [1] P. Berenbrink, R. Elsässer, and T. Friedetzky. Efficient randomised broadcasting in random regular networks with applications in peer-to-peer systems. In *Proc. 27th Symp. Principles of Distributed Computing (PODC)*, pages 155–164, 2008.
- [2] N. Berger, C. Borgs, J. Chayes, and A. Saberi. On the Spread of Viruses on the Internet. In *Proc. 16th Symp. Discrete Algorithms (SODA)*, pages 301–310, 2005.
- [3] S. Boyd, A. Ghosh, B. Prabhakar, and D. Shah. Randomized gossip algorithms. *IEEE Transactions on Information Theory*, 52(6):2508–2530, 2006.
- [4] K. Censor-Hillel and H. Shachnai. Partial information spreading with application to distributed maximum coverage. In *Proc. 29th Symp. Principles of Distributed Computing (PODC)*, pages 161–170, 2010.
- [5] F. Chierichetti, S. Lattanzi, and A. Panconesi. Rumor spreading in social networks. In *Proc. 36th Intl. Coll. Automata, Languages and Programming (ICALP)*, pages 375–386, 2009.
- [6] F. Chierichetti, S. Lattanzi, and A. Panconesi. Almost tight bounds for rumour spreading with conductance. In *Proc. 42nd Symp. Theory of Computing (STOC)*, pages 399–408, 2010.

- [7] A. E. F. Clementi, F. Pasquale, and R. Silvestri. Opportunistic manets: Mobility can make up for low transmission power. *IEEE/ACM Transaction on Networking*, 21(2):610–620, 2013.
- [8] A. Demers, D. Greene, C. Hauser, W. Irish, J. Larson, S. Shenker, H. Sturgis, D. Swinehart, and D. Terry. Epidemic algorithms for replicated database maintenance. In *Proc. 6th Symp. Principles of Distributed Computing (PODC)*, pages 1–12, 1987.
- [9] B. Doerr and M. Fouz. Asymptotically optimal randomized rumor spreading. In *Proc. 38th Intl. Coll. Automata, Languages and Programming (ICALP)*, pages 502–513, 2011.
- [10] B. Doerr, M. Fouz, and T. Friedrich. Social networks spread rumors in sublogarithmic time. In *Proc. 43th Symp. Theory of Computing (STOC)*, pages 21–30, 2011.
- [11] B. Doerr, M. Fouz, and T. Friedrich. Asynchronous rumor spreading in preferential attachment graphs. In *Proc. 13th Scandinavian Workshop Algorithm Theory (SWAT)*, pages 307–315, 2012.
- [12] D. Dubhashi and A. Panconesi. *Concentration of Measure for the Analysis of Randomized Algorithms*. Cambridge University Press, 2009.
- [13] R. Elsässer and T. Sauerwald. On the runtime and robustness of randomized broadcasting. In *Proc. 17th Intl. Symp. Algorithms and Computation, (ISAAC)*, pages 349–358, 2006.
- [14] R. Elsässer and T. Sauerwald. The power of memory in randomized broadcasting. In *Proc. 19th Symp. Discrete Algorithms (SODA)*, pages 218–227, 2008.
- [15] U. Feige, D. Peleg, P. Raghavan, and E. Upfal. Randomized broadcast in networks. *Random Structures and Algorithms*, 1(4):447–460, 1990.
- [16] N. Fountoulakis, A. Huber, and K. Panagiotou. Reliable broadcasting in random networks and the effect of density. In *Proc. 29th IEEE Conf. Computer Communications (INFOCOM)*, pages 2552–2560, 2010.
- [17] N. Fountoulakis and K. Panagiotou. Rumor spreading on random regular graphs and expanders. In *Proc. 14th Intl. Workshop on Randomization and Comput. (RANDOM)*, pages 560–573, 2010.
- [18] N. Fountoulakis, K. Panagiotou, and T. Sauerwald. Ultra-fast rumor spreading in social networks. In *Proc. 23th Symp. Discrete Algorithms (SODA)*, pages 1642–1660, 2012.
- [19] A. Frieze and G. Grimmett. The shortest-path problem for graphs with random-arc-lengths. *Discrete Applied Mathematics*, 10:57–77, 1985.
- [20] G. Giakkoupis. Tight bounds for rumor spreading in graphs of a given conductance. In *Proc. 28th Symp. Theoretical Aspects of Computer Science (STACS)*, pages 57–68, 2011.
- [21] G. Giakkoupis. Tight bounds for rumor spreading with vertex expansion. *CoRR*, abs/1302.6243, 2013.

- [22] G. Giakkoupis and T. Sauerwald. Rumor spreading and vertex expansion. In *Proc. 23th Symp. Discrete Algorithms (SODA)*, pages 1623–1641, 2012.
- [23] M. Harchol-Balter, F. T. Leighton, and D. Lewin. Resource discovery in distributed networks. In *Proc. 18th Symp. Principles of Distributed Computing (PODC)*, pages 229–237, 1999.
- [24] R. Karp, C. Schindelhauer, S. Shenker, and B. Vöcking. Randomized Rumor Spreading. In *Proc. 41st Symp. Foundations of Computer Science (FOCS)*, pages 565–574, 2000.
- [25] D. Kempe, A. Dobra, and J. Gehrke. Gossip-based computation of aggregate information. In *Proc. 44th Symp. Foundations of Computer Science (FOCS)*, pages 482–491, 2003.
- [26] C. McDiarmid. Concentration. In M. Habib, C. McDiarmid, J. Ramirez-Alfonsin, and B. Reed, editors, *Probabilistic Methods for Algorithmic Discrete Mathematics*, pages 195–243. Springer-Verlag, 1998.
- [27] D. Mosk-Aoyama and D. Shah. Fast distributed algorithms for computing separable functions. *IEEE Transactions on Information Theory*, 54(7):2997–3007, 2008.
- [28] Y. Peres, A. Sinclair, P. Sousi, and A. Stauffer. Mobile geometric graphs: Detection, coverage and percolation. In *Proc. 22nd Symp. Discrete Algorithms (SODA)*, pages 412–428, 2011.
- [29] B. Pittel. On spreading a rumor. *SIAM Journal on Applied Mathematics*, 47(1):213–223, 1987.
- [30] R. van Renesse, Y. Minsky, and M. Hayden. A gossip-style failure detection service. In *Proc. 15th IFIP Intl. Conf. Distributed Systems Platforms (Middleware’98)*, pages 55–70, 1998.