Faster Rumor Spreading with Multiple Calls

Konstantinos Panagiotou[1], Ali Pourmiri[2], and Thomas Sauerwald[2,3]

[1] The Ludwig Maximilian University of Munich, Germany
[2] Max Planck Institute for Informatics, Saarbrücken, Germany
[3] University of Cambridge, United Kingdom

Abstract. We consider the random phone call model introduced by Demers et al. [8], which is a well-studied model for information dissemination on networks. One basic protocol in this model is the so-called Push protocol which proceeds in synchronous rounds. Starting with a single node which knows of a rumor, every informed node calls a random neighbor and informs it of the rumor in each round. The Push-Pull protocol works similarly, but additionally every uninformed node calls a random neighbor and may learn the rumor from that neighbor.

While it is well-known that both protocols need $\Theta(\log n)$ rounds to spread a rumor on a complete network with n nodes, we are interested by how much we can speed up the spread of the rumor 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 u is chosen independently according to a probability distribution R with bounded mean determined at the beginning of the process. 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 R follows a power law distribution with exponent $\in (2, 3)$, we show that the Push-Pull protocol spreads a rumor in $\Theta(\log \log n)$ rounds.

1 Introduction

Rumor spreading is an important primitive for information dissemination in networks. The goal is to spread a piece of information, the so-called rumor, from an arbitrary node to all the other nodes. The random phone call model is 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. [11]). Therefore these protocols have been successfully applied in other contexts such as replicated databases [8], failure detection [23], resource discovery [17], load balancing [3], data aggregation [19], and analysis of the spread of computer viruses [2].

A basic protocol for spreading a rumor in the phone call model 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

L. Cai, S.-W. Cheng, and T.-W. Lam (Eds.): ISAAC2013, LNCS 8283, pp. 446–456, 2013.

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 of these protocols which is the number of rounds required until a rumor initiated by a single node reaches all other nodes.

In one of the first papers in this area, Frieze and Grimmett [14] proved that if the underlying graph is a complete graph with n nodes, then the runtime of the Push protocol is $\log_2 n + \log n \pm o(\log n)$ with high probability[1], where log denotes the natural logarithm. This result was later strengthened by Pittel [22]. For the standard Push-Pull protocol, Karp et al. [18] 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. More recently Doerr and Fouz [9] proposed a new protocol using only Push calls with runtime $(1+o(1)) \log_2 n$ using only $O(n \cdot f(n))$ calls (and messages), where $f(n)$ is an arbitrarily slow growing function.

Besides the complete graph, the randomized rumor spreading protocols mentioned above have been shown to be efficient also on other topologies. In particular, their runtime is at most logarithmic in the number of nodes n for topologies ranging from basic networks, such as random graphs [11,12] and hypercubes [11], random regular graphs [1], graphs with constant conductance [20,5,15], constant weak conductance [4] or constant vertex expansion [16], to more complex structures including preferential attachment graphs modeling social networks [10,13]. In particular, recent studies establishing a sub-logarithmic runtime on certain social network models [10,13] raise the question whether it is possible to achieve a sub-logarithmic runtime also on complete graphs. In addition to analyses on static graphs, there are also studies on mobile geometric graphs, e.g., [7,21], that have deal with strong correlations as nodes are moving according to a random walk.

Since the Push protocol, the Pull protocol and the Push-Pull protocol all 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 each node u, denoted by C_u is determined by a probability distribution R on the positive integers 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. While being a natural extension from a theoretical perspective, different C_u values could arise due to varying battery capacities, processor speeds or clock synchronizations. Our aim is to understand the impact of the distribution R on the runtime of randomized rumor spreading. In particular, we seek for conditions on R which are necessary (or sufficient) for a sublogarithmic runtime.

Our first result concerns the Push protocol for the case where R has bounded mean and bounded variance. As this is the most basic setting, our runtime bound

[1] By with high probability we refer to an event which holds with probability $1 - o(1)$ as $n \to \infty$. For simplicity, we sometimes omit the "with high probability" in the introduction.

is even tight up to low-order terms. To this end, let $T_n = \min\{t \mid \mathbf{Pr}[I_t = n] \geqslant 1 - g(n)\}$ be the first round in which all nodes are informed with probability $1 - g(n)$, where $g(n)$ is a function tending to zero as n goes to infinity (for simplicity we do not specify $g(n)$).

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

Note that by putting $R \equiv 1$, we retain the classic result by Frieze and Grimmett for the standard Push protocol. If we drop the assumption on the variance, then the theorem below provides a lower bound of $\Omega(\log n)$. 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 prob. $1 - o(1)$, the* Push *protocol needs $\Omega(\log n)$ 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 [14,22] for the standard Push-Pull protocol already imply an upper bound of $\mathcal{O}(\log n)$ rounds. Next we 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 $\beta < 3$, i.e., there are constants $0 < c_1 \leqslant c_2$ such that $c_1 z^{1-\beta} \leqslant \mathbf{Pr}[R \geqslant z] \leqslant c_2 z^{1-\beta}$ for any $z \geqslant 1$, and $\mathbf{Pr}[R \geqslant 1] = 1$. We are especially interested in power law distributions, because they are scale invariant and have been observed in a variety of settings in real life. 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 prob. $1 - o(1)$.*

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 $\Theta(\log n)$ rounds [2]. While a very similar dichotomy was shown in [13] for random graphs with a power law degree distribution, our result here concerns the spread of the rumor from one to *all* nodes (and not only to a constant fraction as in [13]).

[2] We do not consider the case $\beta \leqslant 2$, 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). The analysis of the case $\beta = 3$ is an interesting open problem.

In addition, the distribution of the edges used throughout the execution of the Push-Pull protocol is different from the distribution of the edges in a power law random graph, as the latter is proportional to the product of the two nodes weights. Therefore it seems difficult to apply the previous techniques for power law random graphs used for the average distance [6] and rumor spreading [13].

Besides the power law distribution, one may also consider a simple two point distribution, where for instance, $R = n$ with probability n^{-1} and $R = 1$ otherwise. It is then straightforward to see that with constant probability, the Push-Pull protocol informs all nodes in $\mathcal{O}(1)$ rounds. The same result also holds if $R = n^{\epsilon}$ with probability $n^{-\epsilon}$ and $R = 1$ otherwise. However, the power law distribution is arguably a more natural distribution which occurs in a variety of instances in practice.

Finally, we also show that it is crucial that the C_u's do not change over time. Instead, suppose we generate a new variable C_u^t according to the distribution R for the number of calls made by node u in each round t. Then one can prove a lower bound of $\Omega(\log n)$ for the Push-Pull protocol for any distribution R with bounded mean. Based on this lower bound it seems crucial to have a *fixed* set of powerful nodes (i.e. nodes u with large C_u) in order to obtain a sublogarithmic rumor spreading time.

2 Definitions and Notations

We now provide additional definitions and notations (note that the classic Push, Pull and Push-Pull protocols have already been defined before). Here we generalize the classic Push, Pull and Push-Pull to the following statistical model on a complete graph with n nodes.

Before the protocol starts, every node u generates a random integer $C_u \geqslant 1$ according to a distribution R. Then, the rumor is placed on a randomly chosen node[3]. Our generalized Push, Pull and Push-Pull protocol proceed like the classic ones except that every (un)informed node u calls C_u node(s) chosen independently and uniformly at random and sends (request) the rumor.

Let $\mathcal{I}_t$ be the set of all informed nodes in round t (which means after the execution of round t) and $\mathcal{U}_t$ be the complement of $\mathcal{I}_t$, i.e., the set of uninformed nodes. The size of $\mathcal{I}_t$ and $\mathcal{U}_t$ are 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 denoted by N_t. Let S_t be the number of Push calls in round $t+1$, so $S_t = \sum_{u \in \mathcal{I}_t} C_u \geqslant I_t$. Let us define $\mathcal{N}_t^{\mathsf{Pull}}$ and $\mathcal{N}_t^{\mathsf{Push}}$ to be the set of newly informed nodes by Pull and Push calls in round $t+1$, respectively. The size of $\mathcal{N}_t^{\mathsf{Pull}}$ and $\mathcal{N}_t^{\mathsf{Push}}$ are denoted by N_t^{Pull} and N_t^{Push}. The size of every set divided by n will be denoted by the

[3] This is equivalent to saying that the initial node which knows the rumor has to be chosen without knowing the sequence $C_u, u \in \mathcal{V}$. We make this assumption throughout the paper, as it is frequently needed for lower bounding the runtime, e.g., the lower bound in Theorem 1.2 may not hold if the rumor initiates from the node with the largest C_u.

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, we define the set

$$\mathcal{L}(z) := \{u \in \mathcal{V} : C_u \geqslant z\}.$$

The size of $\mathcal{L}(z)$ is denoted by $L(z)$. We define Δ to be $\max_{u\in\mathcal{V}} C_u$.

3 Push Protocol

3.1 Push Protocol with Bounded Variance (Thm. 1.1)

In this subsection we assume that the random numbers C_u's are generated according to some distribution R with bounded mean and variance. Recall that $T_n := \min\{t \mid \mathbf{Pr}[I_t = n] \geqslant 1 - o(1)\}$, i.e., the first round in which all nodes are informed with probability $1 - o(1)$. In Theorem 1.1 we show that if R is a distribution with $\mathbf{E}[R] = \mathcal{O}(1)$ and $\mathbf{Var}[R] = \mathcal{O}(1)$, then $|T_n - (\log_{1+\mathbf{E}[R]} n + \log_{e^{\mathbf{E}[R]}} n)| = o(\log n)$.

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 one informed node and ends when $I_t \geqslant \log^5 n$ and $S_t \leqslant \log^{\mathcal{O}(1)} n$. Similar to 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 $S_t \geqslant I_t$. 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 $S_t \leqslant \log^{\mathcal{O}(1)} n$.
- **The Middle Phase.** This phase starts when $\log^5 n \leqslant I_t \leqslant S_t \leqslant \log^{\mathcal{O}(1)} n$ and ends when $I_t \geqslant \frac{n}{\log\log n}$. 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 $\frac{n}{\log\log n}$ informed nodes.
- **The Final Phase.** This phase starts when $I_t \geqslant \frac{n}{\log\log n}$ 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 $\frac{n}{\log^5 n}$. Then we show the probability that an arbitrary uninformed node remains uninformed is $e^{-\mathbf{E}[R]\pm o(\frac{1}{\log n})}$, so U_t decreases by this probability. Finally, an inductive argument establishes that it takes $\log_{e^{\mathbf{E}[R]}} n \pm o(\log n)$ rounds until every node is informed.

3.2 Push Protocol with Arbitrary Variance (Thm. 1.2)

We prove that if 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. In the Push protocol, in round $t + 1$, at most S_t randomly chosen uninformed nodes are informed. Hence the total contribution of newly informed nodes to $\mathbf{E}[S_{t+1}]$ is at most $\mathbf{E}[R] \cdot S_t$. Applying the law of total expectation shows that $\mathbf{E}[S_{t+1}] \leqslant (1+\mathbf{E}[R])^t\, \mathbf{E}[R]$ which implies that $\Omega(\log n)$ rounds are necessary to inform all nodes.

4 Push-Pull Protocol

4.1 Push-Pull Protocol with Bounded Variance (Thm 1.3)

In this part we consider the case where R is a distribution with bounded mean and bounded variance. We prove that with probability at least $1-\epsilon$, the Push-Pull protocol needs at least $\Omega(\log n)$ rounds to inform all nodes. One interesting example for a distribution R with bounded mean and bounded variance is a power law distribution with parameter $\beta > 3$.

The crucial ingredient of the proof is to bound the C_u's of the nodes that become informed by using Pull, i.e., the C_u's of uninformed nodes that call an informed node. Note that the contribution of an uninformed node $u \in \mathcal{U}_t$ to $\mathbf{E}[S_{t+1}]$ is C_u times the probability that it gets informed, which is at most $C_u \cdot (I_t/n) \leqslant C_u \cdot (S_t/n)$. Hence the contribution of $u \in \mathcal{U}_t$ is at most $C_u^2 \cdot (S_t/n)$. Now using the assumption that R has bounded variance, we have that $\sum_{u\in\mathcal{V}} C_u^2 = \mathcal{O}(n)$ which implies that S_t increases only exponentially in t.

4.2 Push-Pull Protocol with Power Law Distr. $2 < \beta < 3$ (Thm. 1.4)

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 each phase takes only $\mathcal{O}(\log\log n)$ rounds. The proof of the lower bound is ommitted in this extended abstract.

Proof of the Upper Bound. The following lemmas about Push will be used throughout this section.

Lemma 4.1. *Consider the* Push *protocol and suppose that* $S_t \leqslant \log^c n$, *where* $c > 0$ *is any constant. Then with probability* $1-\mathcal{O}(\frac{\log^{2c} n}{n})$ *we have* $I_{t+1} = I_t + S_t$.

Lemma 4.2. *Consider the* Push *protocol. Then with probability* $1 - o(\frac{1}{\log n})$ *we have that* $s_t - 2s_t^2 - 2\sqrt{\frac{s_t \log\log n}{n}} \leqslant n_t \leqslant s_t$.

We will also use the following fact about Power law distributions.

Lemma 4.3. *Let* $\{C_u : u \in \mathcal{V}\}$ *be a set of* n *independent random variables and assume that each* C_u *is generated according to a power law distribution with exponent* $\beta > 2$. *Then for every* $z = \mathcal{O}(n^{\frac{1}{\beta-1}}/\log n)$, *it holds with probability* $1 - o(\frac{1}{n})$

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

The Preliminary Phase. This phase starts with just one informed node and ends when $I_t \geqslant n^{\frac{1}{\beta-1}}/(2\log n)$. Let T_1 be the number of rounds needed so that the number of informed nodes exceeds $n^{\frac{1}{\beta-1}}/(2\log n)$. We will show that with probability $1-o(1)$, $T_1 = \mathcal{O}(\log\log n)$. At first we prove the following lemma.

Lemma 4.4. *Let $c > 0$ be any constant. Then with probability $1-o(1)$, the number of rounds needed to inform $\log^c n$ nodes is bounded by $\mathcal{O}(\log\log n)$.*

Proof. In order to prove our lemma we only consider Push calls and apply Lemma 4.1 which states that as long as $S_t \leqslant \log^c n$, with probability $1-\mathcal{O}(\frac{\log^{2c} n}{n})$,

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

Thus as long as $S_t \leqslant \log^c n$, 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 $\log^c n$ nodes. □

Lemma 4.5. *With probability $1-o(1)$, $T_1 = \mathcal{O}(\log\log n)$.*

Proof. Let T_0 be the first round when $I_{T_0} \geqslant \log^{\frac{2}{3-\beta}} n$. Let us define the constant $\gamma := \frac{3-\beta}{2(\beta-2)} > 0$. Let T be the first round such that

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

Now for any $T_0 \leqslant t \leqslant T$, we can apply Lemma 4.3 and conclude that with probability $1-o(\frac{1}{n})$,

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

So,

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

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,

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

Since for any $t \geqslant T_0$, $I_t \geqslant \log^{\frac{2}{3-\beta}} n$, 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 $S_{t+1} \geqslant I_t^{1+\gamma}$. 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 $1 - o(\frac{1}{\log n})$, $S_{t+1}(1 - o(1)) \leqslant N_{t+1}$. Thus,

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

An inductive argument shows that for any integer $k \geqslant 1$ as long as $I_{T_0+2k-2}^{1+\gamma} \leqslant n^{\frac{1}{\beta-1}}/\log n$, 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 $I_{T_0+2k} \geqslant n^{\frac{1}{\beta-1}}/(2\log n)$ which means $T_1 \leqslant T_0 + 2k = \mathcal{O}(\log\log n)$ 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

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

$\square$

The Middle Phase. This phase starts with at least $n^{\frac{1}{\beta-1}}/(2\log n)$ informed nodes and ends when $I_t \geqslant \frac{n}{\log n}$. 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 4.6. *Suppose that* $I_t \geqslant n^{\frac{1}{\beta-1}}/(2\log n)$ *for some round* t. *Let* $Z_{t+1} := \frac{n\log\log n}{I_t}$. *Then with probability* $1 - o(\frac{1}{n})$,

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

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 $X_u = 1$ if u gets informed by Pull in round $t+1$ and $X_u = 0$ otherwise.

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

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

Thus $\mathbf{Pr}\left[X_u = 1\right] > \frac{3}{4}$ and $\mathbf{E}\left[X\right] = \sum_{u \in \mathcal{L}'(Z_{t+1})} \mathbf{Pr}\left[X_u = 1\right] > \frac{3}{4}|\mathcal{L}'(Z_{t+1})|$. Since $|\mathcal{L}'(Z_{t+1})| = |\mathcal{L}(Z_{t+1}) \cap \mathcal{U}_{t+1}| > \frac{3}{4}L(Z_{t+1})$, $\mathbf{E}\left[X\right] \geqslant \frac{9}{16}L(Z_{t+1})$. We know that $I_t \geqslant n^{\frac{1}{\beta-1}}/(2\log n)$ and also I_t is a non-decreasing function in t, so

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

where the last inequality holds because $\beta < 3$. Now we can apply Lemma 4.3 (see appendix) to infer that with probability $1 - o(\frac{1}{n})$,

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

Therefore, $\mathbf{E}\left[X\right] \geqslant \frac{9 \cdot c_1 \cdot \log^{\beta-1} n}{32}$. Then applying a Chernoff bound results into

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

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

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

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

Lemma 4.7. *With probability* $1 - o(1)$, $T_2 - T_1 = \mathcal{O}(\log\log n)$.

Proof. Since $I_t \geqslant n^{\frac{1}{\beta-1}}/(2\log n)$, $Z_{t+1} = \frac{n \log\log n}{I_t} < n^{\frac{1}{\beta-1}}/\log n$, using Lemma 4.6 results into a lower bound for $|\mathcal{L}(Z_{t+1}) \cap \mathcal{I}_{t+1}|$. So with probability $1 - o(\frac{1}{n})$,

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

By applying Lemma 4.3, we conclude that with probability $1 - o(\frac{1}{n})$, $L(Z_{t+1}) \geqslant \frac{n \cdot c_1 \cdot Z_{t+1}^{1-\beta}}{2}$. Therefore, with probability $1 - o(\frac{1}{n})$, $S_{t+1} \geqslant \frac{n \cdot c_1 \cdot Z_{t+1}^{2-\beta}}{8}$. 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 $1 - o(\frac{1}{\log n})$,

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

Thus,

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

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

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

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

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

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

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 [18, Theorem 2.1] for the standard Push-Pull protocol which states that once $I_t \geqslant \frac{n}{\log n}$, additional $\mathcal{O}(\log\log n)$ rounds are with probability $1 - o(1)$ sufficient to inform all n nodes.

References

1. Berenbrink, P., Elsässer, R., Friedetzky, T.: Efficient randomised broadcasting in random regular networks with applications in peer-to-peer systems. In: Proc. 27th Symp. Principles of Distributed Computing (PODC), pp. 155–164 (2008)
2. Berger, N., Borgs, C., Chayes, J., Saberi, A.: On the Spread of Viruses on the Internet. In: Proc. 16th Symp. Discrete Algorithms (SODA), pp. 301–310 (2005)
3. Boyd, S., Ghosh, A., Prabhakar, B., Shah, D.: Randomized gossip algorithms. IEEE Transactions on Information Theory 52(6), 2508–2530 (2006)
4. Censor-Hillel, K., Shachnai, H.: Partial information spreading with application to distributed maximum coverage. In: Proceedings of the 29th Annual ACM Symposium on Principles of Distributed Computing, pp. 161–170 (2010)
5. Chierichetti, F., Lattanzi, S., Panconesi, A.: Almost tight bounds for rumour spreading with conductance. In: Proc. 42nd Symp. Theory of Computing (STOC), pp. 399–408 (2010)
6. Chung, F.R.K., Lu, L.: The average distance in a random graph with given expected degrees. Internet Mathematics 1(1), 91–113 (2003)

7. Clementi, A.E.F., Pasquale, F., Silvestri, R.: Opportunistic manets: Mobility can make up for low transmission power. IEEE/ACM Trans. Netw. 21(2), 610–620 (2013)
8. Demers, A., Greene, D., Hauser, C., Irish, W., Larson, J., Shenker, S., Sturgis, H., Swinehart, D., Terry, D.: Epidemic algorithms for replicated database maintenance. In: Proc. 6th Symp. Principles of Distributed Computing (PODC), pp. 1–12 (1987)
9. Doerr, B., Fouz, M.: Asymptotically optimal randomized rumor spreading. In: Aceto, L., Henzinger, M., Sgall, J. (eds.) ICALP 2011, Part II. LNCS, vol. 6756, pp. 502–513. Springer, Heidelberg (2011)
10. Doerr, B., Fouz, M., Friedrich, T.: Social networks spread rumors in sublogarithmic time. In: Proc. 43rd Symp. Theory of Computing (STOC), pp. 21–30 (2011)
11. Feige, U., Peleg, D., Raghavan, P., Upfal, E.: Randomized broadcast in networks. Random Structures and Algorithms 1(4), 447–460 (1990)
12. Fountoulakis, N., Huber, A., Panagiotou, K.: Reliable broadcasting in random networks and the effect of density. In: Proc. 29th IEEE Conf. Computer Communications (INFOCOM), pp. 2552–2560 (2010)
13. Fountoulakis, N., Panagiotou, K., Sauerwald, T.: Ultra-fast rumor spreading in social networks. In: Proc. 23rd Symp. Discrete Algorithms (SODA), pp. 1642–1660 (2012)
14. Frieze, A., Grimmett, G.: The shortest-path problem for graphs with random-arc-lengths. Discrete Applied Mathematics 10, 57–77 (1985)
15. Giakkoupis, G.: Tight bounds for rumor spreading in graphs of a given conductance. In: Proc. 28th Symp. Theoretical Aspects of Computer Science (STACS), pp. 57–68 (2011)
16. Giakkoupis, G.: Tight bounds for rumor spreading with vertex expansion. In: Proc. 25th Symp. Discrete Algorithms, SODA (to appear, 2014)
17. Harchol-Balter, M., Leighton, F.T., Lewin, D.: Resource discovery in distributed networks. In: Proc. 18th Symp. Principles of Distributed Computing (PODC), pp. 229–237 (1999)
18. Karp, R., Schindelhauer, C., Shenker, S., Vöcking, B.: Randomized Rumor Spreading. In: Proc. 41st Symp. Foundations of Computer Science (FOCS), pp. 565–574 (2000)
19. Kempe, D., Dobra, A., Gehrke, J.: Gossip-based computation of aggregate information. In: Proc. 44th Symp. Foundations of Computer Science (FOCS), pp. 482–491 (2003)
20. Mosk-Aoyama, D., Shah, D.: Fast distributed algorithms for computing separable functions. IEEE Transactions on Information Theory 54(7), 2997–3007 (2008)
21. Peres, Y., Sinclair, A., Sousi, P., Stauffer, A.: Mobile geometric graphs: Detection, coverage and percolation. In: Proc. 22nd Symp. Discrete Algorithms (SODA), pp. 412–428 (2011)
22. Pittel, B.: On spreading a rumor. SIAM Journal on Applied Mathematics 47(1), 213–223 (1987)
23. van Renesse, R., Minsky, Y., Hayden, M.: A gossip-style failure detection service. In: Proceedings of the 15th IFIP International Conference on Distributed Systems Platforms, Middleware 1998, pp. 55–70 (1998)