Almost Tight Bounds for Rumour Spreading with Conductance*

Flavio Chierichetti, Silvio Lattanzi, Alessandro Panconesi
{chierichetti,lattanzi,ale}@di.uniroma1.it

Dipartimento di Informatica
Sapienza Università di Roma

May 28, 2010

Abstract

We show that if a connected graph with n nodes has conductance ϕ then rumour spreading, also known as randomized broadcast, successfully broadcasts a message within $\tilde{O}(\phi^{-1} \cdot \log n)$, many rounds with high probability, regardless of the source, by using the PUSH-PULL strategy. The $\tilde{O}(\cdots)$ notation hides a polylog ϕ^{-1} factor. This result is almost tight since there exists graph of n nodes, and conductance ϕ, with diameter $\Omega(\phi^{-1} \cdot \log n)$.

If, in addition, the network satisfies some kind of uniformity condition on the degrees, our analysis implies that both both PUSH and PULL, by themselves, successfully broadcast the message to every node in the same number of rounds.

Categories and Subject Descriptors

G.3 [Mathematics of Computing]: Probability and Statistics.

General Terms

Algorithms, Theory.

Keywords

Conductance, rumor spreading, social networks.

1 Introduction

Rumour spreading, also known as randomized broadcast or randomized gossip (all terms that will be used as synonyms throughout the paper), refers to the following distributed algorithm. Starting with one source node with a message, the protocol proceeds in a sequence of synchronous rounds with the goal of *broadcasting* the message, i.e. to deliver it to every node in the network. In round $t \geq 0$, every node that knows the message selects a neighbour uniformly at random to which the message is forwarded. This is the so-called PUSH strategy. The PULL variant is symmetric. In round $t \geq 0$ every node that does not yet have the message selects a neighbour uniformly at random and asks for the information, which is transferred provided that the queried neighbour knows it. Finally, the PUSH-PULL strategy is a combination of both. In round $t \geq 0$, each node selects a random neighbour to perform a PUSH if it has the information or a PULL in the opposite case.

These three strategies have been introduced in [5] and since then have been intensely investigated (see the related work section). One of the most studied questions concerns their completion time: how many rounds will it take for one of the above strategies to disseminate the information to all nodes in the graph, assuming a worst-case source? In this paper we prove the following two results:

*This research is partially supported by a grant of Yahoo! Research and by an IBM Faculty Award.

- If a network has conductance ϕ and n nodes, then, with high probability, PUSH-PULL reaches every node within $O\left(\frac{\log^2 \phi^{-1}}{\phi} \cdot \log n\right)$ many rounds, regardless of the source.
- If, in addition, the network satisfies the following condition for every edge uv and some constant $\alpha > 0$:

$$\max\left\{\frac{\deg(u)}{\deg(v)}, \frac{\deg(v)}{\deg(u)}\right\} \leq \alpha$$

then both PUSH and PULL, by themselves[1], reach every node within

$$O(c_\alpha \cdot \phi^{-1} \cdot \log n \cdot \log \phi^{-1})$$

many rounds with high probability regardless of the source, where c_α is a constant depending only on α.

The first result is a significant improvement with respect to best current bound of $O(\log^4 n/\phi^6)$ [4]. (The proof of [4] is based on an interesting connection with spectral sparsification [20]. The approach followed here is entirely different.) The result is almost tight because $\Omega(\log n/\phi)$ is a lower bound[2] — in particular, the bound is tight in the case of constant conductance (for instance, this is the case for the almost-preferential-attachment graphs of [15].) The second result can be proved using the same approach we use for the main one. In this extended abstract we omit the details of its proof.

Our main motivation comes from the study of social networks. Loosely stated, we are looking for a theorem of the form *"Rumour spreading is fast in social networks"*. There is some empirical evidence showing that real social networks have high conductance. The authors of [14] report that in many different social networks there exist only cuts of small (logarithmic) size having small (inversely logarithmic) conductance – all other cuts appear to have larger conductance. That is, the conductance of the social networks they analyze is larger than a quantity seemingly proportional to an inverse logarithm.

Our work should also be viewed in the context of expansion properties of graphs, of which conductance is an important example, and their relationship with rumour spreading. In particular we observe how, interestingly, the convergence time of the PUSH-PULL process on graph of conductance ϕ is a factor of ϕ smaller than the worst-case mixing time of the uniform random walk on such graphs.

Conductance is one of the most studied measures of graph expansion. Edge expansion, and vertex expansion are two other notable measures. In the case of edge expansion there are classes of graphs for which the protocol is slow (see [3] for more details), while the problem remains open for vertex expansion.

In terms of message complexity, we observe first that it has been determined precisely only for very special classes of graphs (cliques [12] and Erdös-Rényi random graphs [8]). Apart from this, given the generality of our class, it is impossible to improve the trivial upper bound on the number of messages – that is, number of rounds times number of nodes. For instance consider the "lollipop graph." Fix $\omega(n^{-1}) < \phi < o(\log^{-1} n)$, and suppose to have a path of length ϕ^{-1} connected to a clique of size $n - \phi^{-1} = \Theta(n)$. This graph has conductance $\approx \phi$. Let the source be any node in the clique. After $\Theta(\log n)$ rounds each node in the clique will have the information. Furthermore, at least ϕ^{-1} steps will be needed for the information to be sent to the each node in the path. So, at least $n - \phi^{-1} = \Theta(n)$ messages are pushed (by the nodes in the clique) in each round, for at least $\phi^{-1} - \Theta(\log n) = \Theta(\phi^{-1})$ rounds. Thus, the total number of messages sent will be $\Omega(n \cdot \phi^{-1})$. Observing that the running time is $\Theta(\phi^{-1} + \log n) = \Theta(\phi^{-1})$, we have that the total number of rounds times n is (asymptotically) less than or equal to the number of transmitted messages.

[1] We observe that the star, a graph of conductance 1, is such that both the PUSH and the PULL strategy by themselves require $\Omega(n)$ many rounds to spread the information to each node, assuming a worst case, or even uniformly random, source. That is, conductance alone is not enough to ensure that PUSH, or PULL, spread the information fast.

[2] Indeed, choose any n, and any $\phi \geq n^{-1+\epsilon}$. Take any 3-regular graph of constant vertex expansion (a random 3-regular graph will suffice) on $O(n \cdot \phi)$ nodes. Then, substitute each edge of the regular graph with a path of $O(\phi^{-1})$ new nodes. The graph obtained is easily seen to have $O(n)$ nodes, diameter $O(\phi^{-1} \cdot \log n)$ and conductance $\Omega(\phi)$.

2 Related work

The literature on the gossip protocol and social networks is huge and we confine ourselves to what appears to be more relevant to the present work.

Clearly, at least diameter-many rounds are needed for the gossip protocol to reach all nodes. It has been shown that $O(n \log n)$ rounds are always sufficient for each connected graph of n nodes [9]. The problem has been studied on a number of graph classes, such as hypercubes, bounded-degree graphs, cliques and Erdös-Rényi random graphs (see [9, 11, 17]). Recently, there has been a lot of work on "quasi-regular" expanders (i.e., expander graphs for which the ratio between the maximum and minimum degree is constant) — it has been shown in different settings [1, 6, 7, 10, 18] that $O(\log n)$ rounds are sufficient for the rumour to be spread throughout the graph. See also [13, 16]. Our work can be seen as an extension of these studies to graphs of arbitrary degree distribution. Observe that many real world graphs (e.g., facebook, Internet, etc.) have a very skewed degree distribution — that is, the ratio between the maximum and the minimum degree is very high. In most social networks' graph models the ratio between the maximum and the minimum degree can be shown to be polynomial in the graph order.

Mihail et al. [15] study the edge expansion and the conductance of graphs that are very similar to preferential attachment (PA) graphs. We shall refer to these as "almost" PA graphs. They show that edge expansion and conductance are constant in these graphs. Their result and ours together imply that rumor spreading requires $O(\log n)$ rounds on almost PA graphs.

For what concerns the original PA graphs, in [3] it is shown that rumour spreading is fast (requires time $O(\log^2 n)$) in those networks.

In [2] it is shown that high conductance implies that *non-uniform* (over neighbours) rumour spreading succeeds. By non-uniform we mean that, for every ordered pair of neighbours i and j, node i will select j with probability p_{ij} for the rumour spreading step (in general, $p_{ij} \neq p_{ji}$). This results does not extend to the case of uniform probabilities studied in this paper. In our setting (but not in theirs), the existence of a non uniform distribution that makes rumour spreading fast is a rather trivial matter. A graph of conductance ϕ has diameter bounded by $O(\phi^{-1} \log n)$. Thus, in a synchronous network, it is possible to elect a leader in $O(\phi^{-1} \log n)$ many rounds and set up a BFS tree originating from it. By assigning probability 1 to the edge between a node and its parent one has the desired non uniform probability distribution. Thus, from the point of view of this paper the existence of non uniform problem is rather uninteresting.

In [16] the authors consider a problem that at first sight might appear equivalent to ours. They consider the conductance ϕ_P of the *connection probability matrix* P, whose entry $P_{i,j}$, $1 \leq i, j \leq n$, gives the probability that i calls j in any given round. They show that if P is doubly stochastic then the running time of `PUSH-PULL` is $O(\phi_P^{-1} \cdot \log n)$. This might seem to subsume our result but this is not the case. The catch is that they consider the conductance of a doubly stochastic matrix instead of the actual conductance of the graph, as we do. Observe that the are graphs of high conductance that do not admit doubly stochastic matrices of high conductance. For instance, in the star, no matter how one sets the probabilities P_{ij}, there will always exist a leaf ℓ that will be contacted by the central node with probability $\leq \frac{1}{n-1}$. Since the matrix is doubly-stochastic this implies that ℓ will contact the central node with probability $O(n^{-1})$. Thus, at least $\Omega(n)$ rounds will be needed. Therefore their result gives too weak a bound for the uniform `PUSH-PULL` process that we analyze in this paper.

3 Preliminaries

Observe that $\frac{1}{2}\operatorname{vol}(V) = |E|$. Given $S \subseteq V$, and $v \in S$, we define

$$N_S^+(v) = \{w \mid w \in V - S \wedge \{v, w\} \in E\}$$

and $d_S^+(v) = |N_S^+(v)|$. Analogously, we define $N_S^-(w) = N_{V-S}^+(w)$ and $d_S^-(w) = |N_S^-(w)|$. Recall that the conductance (see [19]) of a graph $G(V,E)$ is:

$$\Phi(G) = \min_{\substack{S \subset V \\ \mathrm{vol}(S) \leq |E|}} \frac{\mathrm{cut}(S, V-S)}{\mathrm{vol}(S)}$$

Where $\mathrm{cut}(S, V-S)$ is the number of edges in the cut between S and $V-S$ and $\mathrm{vol}(S)$ is the volume of S.

We recall three classic concentration results for random variables using, respectively, the first moment, the second moment and every moment of a random variable X.

Theorem 1 (Markov inequality) *Let X be a random variable. Then, if*

$$\Pr\left[|X| \geq \frac{E[|X|]}{\epsilon}\right] \leq \epsilon.$$

Theorem 2 (Chebyshev inequality) *Let X be a random variable. Then, if*

$$\Pr\left[|X - E[X]| \geq \sqrt{\mathrm{Var}[X]/\epsilon}\right] \leq \epsilon,$$

where $\mathrm{Var}[X]$ *is the variance of* X*,* $\mathrm{Var}[X] = E[X^2] - E[X]^2$.

Theorem 3 (Chernoff bound) *Let $X = \sum_{i=1}^{n} X_i$, where X_i are independently distributed random variables in* $[0,1]$*. Then,*

$$\Pr\left[|X - E[X]| > \epsilon \cdot E[X]\right] \leq \exp\left(-\frac{\epsilon^2}{3} \cdot E[X]\right).$$

We now state and prove some technical lemmas that we will use in our analysis. The first one can be seen as an "inversion" of Markov's inequality.

Lemma 4 *Suppose $X_1, X_2, \ldots, X_t$ are random variables, with X_i having co-domain* $\{0, v_i\}$ *and such that $X_i = v_i$ with probability p_i. Fix $p \leq \min p_i$. Then, for each* $0 < q < p$*,*

$$\Pr\left[\sum_i X_i \geq \left(1 - \frac{1-p}{1-q}\right) \cdot \sum_i v_i\right] \geq q.$$

Proof Let $\overline{X}_i = v_i - X_i$. Observe that each X_i and each $\overline{X}_i$ is a non-negative random variable, of mean $p_i \cdot v_i$ and $(1-p_i) \cdot v_i$, respectively. We use μ to denote the expected sum of the X_i, $\mu = \sum(p_i \cdot v_i)$, and $\overline{\mu}$ to denote the expected sum of the $\overline{X}_i$, $\overline{\mu} = \sum((1-p_i) \cdot v_i)$. Observe that $\overline{\mu} \leq (1-p) \cdot \sum v_i$.

By Markov's inequality, we have

$$\begin{aligned}
\Pr\left[\sum_i X_i \leq \left(1 - \frac{1-p}{1-q}\right) \sum_i v_i\right] &\leq \\
\Pr\left[\sum_i X_i \leq \sum_i v_i - \frac{1}{1-q} \cdot \overline{\mu}\right] &= \\
\Pr\left[\sum_i \overline{X}_i \geq \frac{1}{1-q} \cdot \overline{\mu}\right] &\leq 1-q
\end{aligned}$$

Thus the claim. □

The next lemma gives some probabilistic bounds on the sum of binary random variables having close expectations.

Lemma 5 *Let $p \in (0,1)$. Suppose $X_1, \ldots, X_t$ are independent 0/1 random variables, the i-th of which such that $\Pr[X_i = 1] = p_i$, with $\frac{1}{2} \cdot p \leq p_i \leq p$. Then,*

1. *if $\frac{pt}{2} > 1$, then $\Pr[\sum X_i \geq \frac{pt}{4}] \geq \frac{1}{32}$;*
2. *if $\frac{pt}{2} \leq 1$, then $\Pr[\sum X_i \geq 1] \geq \frac{pt}{4}$;*
3. *in general, for $P = \min\left(\frac{1}{32}, \frac{pt}{4}\right)$, we have*

$$\Pr\left[\sum X_i \geq \frac{pt}{128 \cdot P}\right] \geq P.$$

Proof Let $X = \sum_{i=1}^{t} X_i$. In the first case, $E[X] \geq t \cdot \frac{p}{2}$; in particular, $E[X] \geq 1$. Therefore, by Chernoff's bound, we have

$$\begin{aligned} \Pr\left[X < \frac{1}{2} \cdot E[X]\right] &\leq e^{-\frac{1}{16}E[X]} \\ &\leq e^{-\frac{1}{16}} \leq 1 - \frac{1}{32}, \end{aligned}$$

where the last inequality follows from $e^{-x} \leq 1 - \frac{x}{2}$ if $x \in [0,1]$.

In the second case, we compute the probability that for no i, $X_i = 1$:

$$\begin{aligned} \prod_{i=1}^{t} \Pr\left[X_i = 0\right] &\leq \prod_{i=1}^{t} \left(1 - \frac{p}{2}\right) \\ &= \left(1 - \frac{p}{2}\right)^t = \left(\left(1 - \frac{p}{2}\right)^{\frac{2}{p}}\right)^{\frac{p}{2} \cdot t} \\ &\leq e^{-\frac{p}{2} \cdot t} \leq 1 - \frac{pt}{4}. \end{aligned}$$

So, with probability $\geq \frac{pt}{4}$ at least one X_i will be equal to 1.

The third case, follows directly from the former two, by choosing — respectively — $P = \frac{1}{32}$, and $P = \frac{pt}{2}$.
□

The following lemma, which we will use later in the analysis, gives a probability bound close to the one that one could obtain via Bernstein Inequality. We keep it this way, for simplicity of exposition of our later proofs.

Lemma 6 *Suppose a player starts with a time budget of B time units. At each round i, an adversary (knowledgeable of the past) chooses a number of time units $1 \leq \ell_i \leq L$. If the remaining budget of the player is at least ℓ_i then a game, lasting for ℓ_i time units, is played. The outcome of the game is determined by an independent random coin flip: with probability $p_i \geq P$ the gain is equal to ℓ_i, the length of the round, and with probability $1 - p_i$ the gain is zero. The game is then replayed.*

If $B \geq 193 \cdot \frac{L}{P} \cdot \ln \frac{\lceil \log_2 L \rceil}{\delta}$ with probability at least $1 - \delta$ the gain is at least $\frac{24}{193} \cdot B \cdot P$.

Proof Let the game go on until the end. Suppose the adversary chose games' lengths $\ell_1, \ell_2, \ldots, \ell_t$, with $\sum_{i=1}^{t} \ell_t > B - L \geq \frac{192}{193} B$.

Let X_j be the set containing all the rounds whose ℓ_i's were such that $2^j \leq \ell_i < 2^{j+1}$, $X_j = \{i \mid 2^j \leq \ell_i < 2^{j+1}\}$. The sets $X_0, X_1, \ldots, X_{\lceil \log_2 L \rceil}$ partition the rounds in $O(\log L)$ buckets.

Assign to each bucket X_j the total number $S(X_j)$ of time units "spent" in that bucket, $S(X_j) = \sum_{i \in X_j} \ell_i$.

Let $\mathcal{X}$ be the set of buckets X_j for which $S(X_j) \geq \frac{12}{P} \cdot 2^{j+1} \cdot \ln \frac{\lceil \log_2 L \rceil}{\delta}$. The total number of time units spent in buckets not in $\mathcal{X}$ will then be at most

$$\sum_{j=0}^{\lceil \log_2 L \rceil} \left(\frac{12}{P} \cdot 2^{j+1} \cdot \ln \frac{\lceil \log_2 L \rceil}{\delta} \right) = \frac{12}{P} \cdot \ln \frac{\lceil \log_2 L \rceil}{\delta} \cdot \sum_{j=0}^{\lceil \log_2 L \rceil} 2^{j+1} \leq \frac{96}{P} \cdot L \cdot \ln \frac{\lceil \log_2 L \rceil}{\delta}.$$

Therefore the total number of units spent in buckets of $\mathcal{X}$, $S(\mathcal{X}) = \sum_{X_j \in \mathcal{X}} S(X_j)$, will be at least $S(\mathcal{X}) \geq \frac{96}{193} B$. Furthermore, the number of rounds $|X_j|$ played in bucket $X_j \in \mathcal{X}$ will be at least $S(X_j) \cdot 2^{-(j+1)} \geq \frac{12}{P} \cdot \ln \frac{\lceil \log_2 L \rceil}{\delta}$. Each such round will let us gain a positive amount with probability at least P. Therefore, the expected number of rounds $E[W(X_j)]$ in bucket $X_j \in \mathcal{X}$ having positive gain will be at least

$$E[W(X_j)] \geq 2^{-(j+1)} \cdot S(X_j) \cdot P \geq 12 \cdot \ln \frac{\lceil \log_2 L \rceil}{\delta}.$$

By the Chernoff bound,

$$\begin{aligned} \Pr\left[W(X_j) < \frac{1}{2} \cdot 2^{-(j+1)} \cdot S(X_j) \cdot P\right] &\leq \Pr\left[W(X_j) < \frac{1}{2} \cdot E[W(X_j)]\right] \\ &\leq \exp\left(-\frac{1}{12} \cdot E[W(X_j)]\right) \\ &\leq \exp\left(-\ln \frac{\lceil \log_2 L \rceil}{\delta}\right) = \frac{\delta}{\lceil \log_2 L \rceil}. \end{aligned}$$

Observe that the gain $G(X_j)$ in bucket $X_j \in \mathcal{X}$ is at least $2^j \cdot W(X_j)$. By union bound, the probability that at least one bucket $X_j \in \mathcal{X}$ is such that $G(X_j) < \frac{1}{4} \cdot S(X_j) \cdot P$ is at most δ. Therefore with probability at least $1 - \delta$, the total gain is at least

$$\frac{P}{4} \sum_{X_j \in \mathcal{X}} S(X_j) = \frac{P}{4} \cdot S(\mathcal{X}) \geq \frac{24}{193} B \cdot P.$$

□

Finally we give a lemma that underlines symmetries in the PUSH-PULL strategy. Let $u \xrightarrow{t} v$ be the event that an information originated at u arrives to v in t rounds using the PUSH-PULL strategy. And let $u \xleftarrow{t} v$ the event that an information that is originally in v arrives to u in t rounds using the PUSH-PULL strategy. We have that:

Lemma 7 *Let $u, v \in V$, then*

$$Pr[u \xrightarrow{t} v] = Pr[u \xleftarrow{t} v]$$

Proof Look at each possible sequence of PUSH-PULL requests done by the nodes of G in t rounds. We define the "inverse" of some sequence, as the sequence we would obtain by looking at the sequence starting from the last round to the first. Now the probability that the information spreads from u to v (resp., from v to u) in at most t steps is equal to the sum of the probabilities of sequences of length at most t that manage to pass the information from u to v (from v to u) — given that the probability of a sequence, and that of its inverse, are the same, the claim follows. □

4 Warm-up: a weak bound

In this section we prove a completion time bound for the PUSH-PULL strategy of $O(\phi^{-2} \cdot \log n)$. Observe that this bound happens to be tight if $\phi \in \Omega(1)$. The general strategy is as follows:

- we will prove that, given any set S of informed nodes having volume $\leq |E|$, after $O(\phi^{-1})$ rounds (that we call a *phase*) the new set S' of informed vertices, $S' \supseteq S$, will have volume $\mathrm{vol}(S') \geq (1+\Omega(\phi))\cdot\mathrm{vol}(S)$ with constant probability (over the random choices performed by nodes during those $O(\phi^{-1})$ rounds) — if this happens, we say that the phase was successful; this section is devoted to proving this lemma;

- given the lemma, it follows that `PUSH-PULL` informs a set of nodes of volume larger than $|E|$, starting from any single node, in time $O(\phi^{-2} \cdot \log n)$.

  Indeed, by applying the Chernoff bound one can prove that, by flipping $c \cdot \phi^{-1} \cdot \log n$ IID coins, each having $\Theta(1)$ head probability, the number of heads will be at least $f(c) \cdot \phi^{-1} \cdot \log n$ with high probability — with $f(c)$ increasing, and unbounded, in c. This implies that we can get enough (that is, $\Theta(\phi^{-1} \cdot \log n)$) successful phases for covering more than half of the graph's volume in at most $\Theta(\phi^{-1}) \cdot \Theta(\phi^{-1} \cdot \log n) = \Theta(\phi^{-2} \cdot \log n)$ rounds;

- applying lemma 7, we can then show that each uninformed node can get the information in the same amount of steps, if a set S of volume $> |E|$ is informed — completing the proof. Recall that the probability that the information spreads from any node v to a set of nodes with more than half the volume of the graph is $1 - O(n^{-2})$. Then, with that probability the source node s spreads the information to a set of nodes with such volume. Furthermore, by lemma 7, any uninformed node would get the information from some node — after node s successfully spreads the information — with probability $1 - O(n^{-2})$. By union bound, we have that with probability $1 - O(n^{-1})$ `PUSH-PULL` will succeed in $O(\phi^{-2} \cdot \log n)$.

Our first lemma shows how one can always find a subset of nodes in the "smallest" part of a good conductance cut, that happen to hit many of the edges in the cut, and whose nodes have a large "fraction" of their degree that cross the cut.

Lemma 8 *Let $G(V, E)$ be a simple graph.*

Let $A \subseteq B \subseteq V$, with $\mathrm{vol}(B) \leq |E|$ and $\mathrm{cut}(A, V - B) \geq \frac{3}{4} \cdot \mathrm{cut}(B, V - B)$. Suppose further that the conductance of the cut $(B, V - B)$ is at least ϕ, $\mathrm{cut}(B, V - B) \geq \phi \cdot \mathrm{vol}(B)$. If we let

$$U = U_B(A) = \left\{ v \in A \middle| \frac{d_B^+(v)}{d(v)} \geq \frac{\phi}{2} \right\},$$

then $\mathrm{cut}(U, V - B) \geq \frac{1}{4} \cdot \mathrm{cut}(B, V - B)$.

Proof We prove the lemma with the following derivation:

$$\sum_{v \in A} d_B^+(v) + \sum_{v \in B-A} d_B^+(v) = \mathrm{cut}(B, V - B)$$

$$\sum_{v \in A} d_B^+(v) = \mathrm{cut}(B, V - B) - \sum_{v \in B-A} d_B^+(v)$$

$$\begin{aligned} \sum_{v \in A \cap U} d_B^+(v) + \sum_{v \in A-U} d_B^+(v) &= \mathrm{cut}(A, V - B) \\ \sum_{v \in U} d_B^+(v) + \sum_{v \in A-U} d_B^+(v) &\geq \frac{3}{4} \cdot \mathrm{cut}(B, V - B), \end{aligned}$$

then,

$$
\begin{aligned}
\sum_{v \in U} d_B^+(v) &\geq \frac{3}{4} \cdot \mathrm{cut}(B, V - B) - \sum_{v \in A-U} d_B^+(v) \\
\sum_{v \in U} d_B^+(v) &\geq \frac{3}{4} \cdot \mathrm{cut}(B, V - B) - \sum_{v \in A-U} \left(\frac{\phi}{2} \cdot d(v) \right) \\
\sum_{v \in U} d_B^+(v) &\geq \frac{3}{4} \cdot \mathrm{cut}(B, V - B) - \frac{\phi}{2} \cdot \mathrm{vol}(B) \\
\sum_{v \in U} d_B^+(v) &\geq \frac{3}{4} \cdot \mathrm{cut}(B, V - B) - \frac{1}{2} \cdot \mathrm{cut}(B, V - B) \\
\sum_{v \in U} d_B^+(v) &\geq \frac{1}{4} \cdot \mathrm{cut}(B, V - B).
\end{aligned}
$$

□

Given $v \in U = U_B(A)$, we define $N_U^{\circ\!\to}(v)$ (to be read "N-push-U-v") and $N_U^{\circ\!\leftarrow}(v)$(to be read "N-pull-U-v") as follows

$$N_U^{\circ\!\to}(v) = \{u \in N_U^+(v) \mid d(u) \geq \frac{1}{3} d^+(v)\}$$

and

$$N_U^{\circ\!\leftarrow}(v) = \{u \in N_U^+(v) \mid d(u) < \frac{1}{3} d^+(v)\}.$$

Then,

$$U^{\circ\!\to} = \{v \in U \mid \left|N_B^{\circ\!\to}(v)\right| \geq \left|N_B^{\circ\!\leftarrow}(v)\right|\}$$

and

$$U^{\circ\!\leftarrow} = \{v \in U \mid \left|N_B^{\circ\!\leftarrow}(v)\right| > \left|N_B^{\circ\!\to}(v)\right|\}.$$

Observe that $U^{\circ\!\to} \cap U^{\circ\!\leftarrow} = \varnothing$ and $U^{\circ\!\to} \cup U^{\circ\!\leftarrow} = U$. In particular, (at least) one of $\mathrm{vol}(U^{\circ\!\to}) \geq \frac{1}{2} \cdot \mathrm{vol}(U)$ and $\mathrm{vol}(U^{\circ\!\leftarrow}) \geq \frac{1}{2} \cdot \mathrm{vol}(U)$ holds. In the following, if $\mathrm{vol}(U^{\circ\!\to}) \geq \frac{1}{2} \cdot \mathrm{vol}(U)$ we will "apply" the `PUSH` strategy on U; otherwise, we will "apply" the `PULL` strategy.

Given a vertex $v \in U$, we will simulate either the `PUSH` the `PULL` strategy, for $O\left(\frac{1}{\phi}\right)$ steps over it. The "gain" $g(v)$ of node v is then the volume of the node(s) that pull the information from v, or that v pushes the information to.

Our aim is to get a bound on the gain of the whole original vertex set S. This cannot be done by summing the gains of single vertices in S, because of the many dependencies in the process. For instance, different nodes $v, v' \in S$ might inform (or could be asked the information by) the same node in $V - S$.

To overcome this difficulty, we use an idea similar in spirit to the deferred decision principle. First of all, let us remark that, given a vertex set S having the information, we will run the `PUSH-PULL` process for $O\left(\frac{1}{\phi}\right)$ rounds. We will look at what happens to the neighbourhoods of different nodes in S sequentially, by simulating the $O\left(\frac{1}{\phi}\right)$ steps (which we call a phase) of each $v \in S$ and some of its peers in $N_S^+(v) \subseteq V - S$. Obviously, we will make sure that no node in $V - S$ performs more than $O\left(\frac{1}{\phi}\right)$ `PULL` steps in a single phase.

Specifically, we consider Process 1 with a phase of $k = \left\lceil \frac{10}{\phi} \right\rceil$ steps.

Observe that — in Process 1 with $k = \left\lceil \frac{10}{\phi} \right\rceil$ — no vertex in $V - S$ will make more than $O\left(\frac{1}{\phi}\right)$ `PULL` steps in a single phase. Indeed, each time we run point 3 in the process, we *only* disclose whether some node $u \in V - S$ actually makes, or does not make, a `PULL` from v. If the `PULL` does not go through, and node u

Process 1 The expansion process of the $O\left(\frac{\log n}{\phi^2}\right)$ bound, with a phase length of k steps.

at step i, we consider the sets A_i, B_i; at the first step, $i = 0$, we take $A_0 = B_0 = S$ and $H_0 = \varnothing$;
if $\mathrm{cut}(A_i, V - B_i) < \frac{3}{4} \cdot \mathrm{cut}(B_i, V - B_i)$, or $\mathrm{vol}(B_i) > |E|$, we stop; otherwise, apply lemma 8 to A_i, B_i, obtaining set $U_i = U_{B_i}(A_i)$;
we take a node v out of U_i, and we consider the effects of either the push or the pull strategy, repeated for k steps, over v and $N^+_{B_i}(v)$;
$H_{i+1} \leftarrow H_i$;
each node $u \in N^+_{B_i}(v)$ that gets informed (either by a push of v, or by a pull from v) is added to the set of the "halted nodes" H_{i+1}; v is also added to the set H_{i+1};
let $A_{i+1} = A_i - \{v\}$, and $B_{i+1} = B_i \cup H_{i+1}$; observe that $B_{i+1} - A_{i+1} = H_{i+1}$;
iterate the process.

later tries to make a `PULL` to another node v', the probability of this second batch of `PULL`'s (and in fact, of any subsequent batch) to succeed is actually larger than the probability of success of the first batch of `PULL`'s of u (since at that point, we already know that the previous `PULL` batches made by u *never* reached any previous candidate node $v \in S$).

The next lemma summarize the gain, in a single step, of a node $v \in U_i$.

Lemma 9 *If $v \in U_i^{\circ\!\!\to}$, then*

$$\Pr\left[g(v) \geq \frac{1}{3} \cdot d^+_{B_i}(v)\right] \geq \frac{\phi}{4}.$$

On the other hand, if $v \in U_i^{\circ\!\!\leftarrow}$ then

$$\Pr\left[g(v) \geq \frac{1}{20} \cdot d^+_{B_i}(v)\right] \geq \frac{1}{10}.$$

In general, if $v \in U_i$,

$$\Pr\left[g(v) \geq \frac{1}{20} \cdot d^+_{B_i}(v)\right] \geq \frac{\phi}{10}.$$

Proof Suppose that $v \in U_i^{\circ\!\!\to}$. Then, at least $\frac{1}{2} d^+_{B_i}(v)$ of the neighbours of v that are not in B_i have degree $\geq \frac{1}{3} d^+_{B_i}(v)$. Since $v \in U_i$, we have that $\frac{d^+_{B_i}(v)}{d(v)} \geq \frac{\phi}{2}$. Thus, the probability that v pushes the information to one of its neighbours of degree $\geq \frac{1}{3} d^+_{B_i}(v)$ is $\geq \frac{\frac{1}{2} d^+_{B_i}(v)}{d(v)} \geq \frac{\phi}{4}$.

Now, suppose that $v \in U_i^{\circ\!\!\leftarrow}$. Recall that $g(v)$ is the random variable denoting the gain of v; that is,

$$g(v) = \sum_{u \in N^{\circ\!\!\leftarrow}_{B_i}(v)} g_u(v),$$

where $g_u(v)$ is a random variable equal to $d(u)$ if u pulls the information from v, and 0 otherwise.

Observe that $\mathrm{E}[g_u(v)] = 1$, so that $\mathrm{E}[g(v)] = \left|N^{\circ\!\!\leftarrow}_{B_i}(v)\right|$, and that the variance of $g_u(v)$ is

$$\begin{aligned}
\mathrm{Var}[g_u(v)] &= \mathrm{E}[g_u(v)^2] - \mathrm{E}[g_u(v)]^2 \\
&= \frac{1}{d(u)} \cdot d(u)^2 - 1 \\
&= d(u) - 1.
\end{aligned}$$

Since the $g_{u_1}(v), g_{u_2}(v), \ldots$ are independent, we have

$$\begin{aligned}
\mathrm{Var}[g(v)] &= \sum_{u \in N^{\circledast}_{B_i}(v)} \mathrm{Var}[g_u(v)] = \sum_{u \in N^{\circledast}_{B_i}(v)} (d(u)-1) \\
&\leq \mathrm{vol}(N^{\circledast}_{B_i}(v)) \\
&\leq \frac{1}{3}\left(d^+_{B_i}(v)\right)\left|N^{\circledast}_{B_i}(v)\right|.
\end{aligned}$$

In the following chain of inequalities we will apply Chebyshev's inequality to bound the deviation of $g(v)$, using the variance bound we have just obtained:

$$\begin{aligned}
\Pr\left[g(v) \leq \frac{1}{20}\cdot d^+_{B_i}(v)\right] &\leq \\
\Pr\left[g(v) \leq \mathrm{E}[g(v)] - \mathrm{E}[g(v)] + \frac{1}{20}\cdot d^+_{B_i}(v)\right] &\leq \\
\Pr\left[-g(v) + \mathrm{E}[g(v)] \geq \left|N^{\circledast}_{B_i}(v)\right| - \frac{1}{20}\cdot d^+_{B_i}(v)\right] &\leq \\
\Pr\left[\left|g(v) - \mathrm{E}[g(v)]\right| \geq \left|N^{\circledast}_{B_i}(v)\right| - \frac{1}{20}\cdot d^+_{B_i}(v)\right] &\leq \\
\Pr\left[\left|g(v) - \mathrm{E}[g(v)]\right| \geq \frac{\left|N^{\circledast}_{B_i}(v)\right| - \frac{1}{20}d^+_{B_i}(v)}{\sqrt{\frac{1}{3}d^+_{B_i}(v)\left|N^{\circledast}_{B_i}(v)\right|}}\cdot\sqrt{\mathrm{Var}[g(v)]}\right] &\leq \\
\frac{\frac{1}{3}d^+_{B_i}(v)\left|N^{\circledast}_{B_i}(v)\right|}{\left(\left|N^{\circledast}_{B_i}(v)\right| - \frac{1}{20}d^+_{B_i}(v)\right)^2} &\leq \\
\frac{\frac{1}{6}\left(d^+_{B_i}(v)\right)^2}{\left(\frac{1}{2}d^+_{B_i}(v) - \frac{1}{20}d^+_{B_i}(v)\right)^2} &\leq \\
\left(\frac{20}{9\cdot\sqrt{6}}\right)^2 \leq \frac{9}{10}.
\end{aligned}$$

This concludes the proof of the second claim. The third one is combination of the other two. □

Now we focus on v, and its neighbourhood $N^+_{B_i}(v)$, for $\left\lceil\frac{10}{\phi}\right\rceil$ many steps. What is the gain $G(v)$ of v in these many steps?

Lemma 10 $\Pr[G(v) \geq \frac{1}{20}\cdot d^+_{B_i}(v)] \geq 1 - e^{-1}$.

Proof Observe that the probability that the event "$g(v) \geq \frac{1}{20}\cdot d^+_{B_i}(v)$" happens at least once in $\left\lceil\frac{10}{\phi}\right\rceil$ independent trials is lower bounded by

$$1 - \left(1 - \frac{\phi}{10}\right)^{\left\lceil\frac{10}{\phi}\right\rceil} \geq 1 - \left(1 - \frac{\phi}{10}\right)^{\frac{10}{\phi}} \geq 1 - e^{-1}.$$

The claim follows. □

We now prove the main theorem of the section:

Theorem 11 *Let S be the set of informed nodes,* $\mathrm{vol}(S) \leq |E|$. *Then, if S' is the set of informed nodes after $\Omega(\phi^{-1})$ steps, then with $\Omega(1)$ probability,*

$$\mathrm{vol}(S') \geq (1 + \Omega(\phi))\cdot\mathrm{vol}(S).$$

Proof Consider Process 1 with a phase of length $k = \left\lceil \frac{10}{\phi} \right\rceil$. For the process to finish, at some step t it must happen that either $\mathrm{vol}(B_t) > |E|$ (in which case, we are done — so we assume the contrary), or $\mathrm{cut}(A_t, V - B_t) < \frac{3}{4} \cdot \mathrm{cut}(B_t, V - B_t)$. Analogously,

$$\frac{1}{4} \cdot \mathrm{cut}(B_t, V - B_t) \leq \mathrm{cut}(B_t - A_t, V - B_t) = \mathrm{cut}(H_t, V - B_t).$$

But then,

$$\begin{aligned}
\frac{1}{4} \cdot \phi \cdot \mathrm{vol}(S) &\leq \frac{1}{4} \cdot \phi \cdot \mathrm{vol}(B_t) \\
&\leq \frac{1}{4} \cdot \mathrm{cut}(B_t, V - B_t) \\
&\leq \mathrm{cut}(H_t, V - B_t) \\
&= \sum_{v \in H_t \cap S} d^+_{B_t}(v) + \sum_{v \in H_t \cap (V-S)} d^+_{B_t}(v) \\
&\leq \sum_{v \in H_t \cap S} d^+_{B_t}(v) + \sum_{v \in H_t \cap (V-S)} \mathrm{vol}(v)
\end{aligned}$$

Consider the following two inequalities (that might, or might not, hold):

(a) $\sum_{v \in H_t \cap (V-S)} \mathrm{vol}(v) \geq \frac{1}{1000} \cdot \phi \cdot \mathrm{vol}(S)$, and

(b) $\sum_{v \in H_t \cap S} d^+_{B_t}(v) \geq \frac{249}{1000} \cdot \phi \cdot \mathrm{vol}(S)$.

At least one of (a) and (b) has to be true. We call *two-cases property* the disjunction of (a) and (b). If (a) is true, we are done, in the sense that we have captured enough volume to cover a constant fraction of the cut induced by S.

We lower bound the probability of (a) to be false given the truth of (b), since the negation of (b) implies the truth of (a).

Recall lemma 10. It states that — for each $v_i \in H_t \cap S$ — we had probability at least $1 - e^{-1}$ of gaining at least $\frac{1}{20} \cdot d^+_{B_i}(v_i) \geq \frac{1}{20} \cdot d^+_{B_t}(v_i)$, since $i \leq t$ implies $B_i \subseteq B_t$.

For each v_i, let us define the random variable X_i as follows: with probability $1 - e^{-1}$, X_i has value $\frac{1}{20} \cdot d^+_{B_t}(v_i)$, and with the remaining probability it has value 0. Then, the gain of v_i is a random variable that dominates X_i. Choosing $q = 1 - 2e^{-1}$, in lemma 4, we can conclude that

$$\Pr\left[\sum_{i:\, v_i \in H_t \cap S} X_i \geq \frac{1}{2} \cdot \sum_{v_i \in H_t \cap S} \left(\frac{1}{20} \cdot d^+_{B_t}(v_i)\right)\right] \geq 1 - 2e^{-1}.$$

Thus, with constant probability ($\geq 1 - 2e^{-1}$) we gain at least $\frac{1}{40} \cdot \sum_{v_i} d^+_{B_t}(v_i)$, which in turn, is at least

$$\frac{1}{40} \cdot \sum_{v_i \in H_t \cap S} d^+_{B_t}(v_i) \geq \frac{1}{40} \cdot \frac{249}{1000} \cdot \phi \cdot \mathrm{vol}(S) \geq \frac{6}{1000} \cdot \phi \cdot \mathrm{vol}(S).$$

Hence, (a) is true with probability at least $1 - 2e^{-1}$. So with constant probability there is a gain of $\frac{1}{1000} \cdot \phi \cdot \mathrm{vol}(S)$ in $\frac{1}{\phi}$ steps. Thus using the proof strategy presented at the beginning of the section we get a $O\left(\phi^{-2} \log n\right)$ bound on the completion time. $\square$

5 A tighter bound

In this section we will present a tighter bound of

$$O\left(\frac{\log^2 \frac{1}{\phi}}{\phi} \cdot \log n\right) = \tilde{O}\left(\frac{\log n}{\phi}\right).$$

Observe that, given the already noted diametral lower bound of $\Omega\left(\frac{\log n}{\phi}\right)$ on graphs of conductance $\phi \geq \frac{1}{n^{1-\epsilon}}$, the bound is almost tight (we only lose an exponentially small factor in ϕ^{-1}).

Our general strategy for showing the tighter bound will be close in spirit to the one we used for the $O\left(\frac{\log n}{\phi^2}\right)$ bound of the previous section.

The new strategy is as follows:

- we will prove in this section that, given any set S of informed nodes having volume at most $|E|$, for some $p = p(S) \geq \Omega(\phi)$, after $O(p^{-1})$ rounds (that we call a *p-phase*) the new set S' of informed vertices, $S' \supseteq S$, will have volume $\mathrm{vol}(S') \geq \left(1 + \Omega\left(\frac{\phi}{p \cdot \log^2 \phi^{-1}}\right)\right) \cdot \mathrm{vol}(S)$ with constant probability (over the random choices performed by nodes during those $O(p^{-1})$ rounds) — if this happens, we say that the phase was successful;

- using the previous statement we can show that PUSH-PULL informs a set of nodes of volume larger than $|E|$, starting from any single node, in time $T \leq O(\phi^{-1} \cdot \log^2 \phi^{-1} \cdot \log n)$ with high probability.

  Observe that at the end of a phase one has a multiplicative volume gain of

  $$1 + \Omega\left(\frac{\phi}{p \cdot \log^2 \phi^{-1}}\right)$$

  with probability lower bounded by a positive constant c. If one averages that gain over the $O(p^{-1})$ rounds of the phase, one can say that with constant probability c, each round in the phase resulted in a multiplicative volume gain of $1 + \Omega\left(\frac{\phi}{\log^2 \phi^{-1}}\right)$.

  We therefore apply lemma 6 with $L = \Theta\left(\frac{\log^2 \phi^{-1}}{\phi}\right)$, $B = \Theta\left(\frac{\log^2 \phi^{-1}}{\phi} \cdot \log n\right)$, $P = c$ and δ equal to any inverse polynomial in n, $\delta = n^{-\Theta(1)}$. Observe that $B \geq \Theta\left(\frac{L}{P} \log \frac{\log L}{\delta}\right)$. Thus, with probability $1 - \delta = 1 - n^{-\Theta(1)}$, we have $\Theta(B \cdot P) = \Theta\left(\frac{\log^2 \phi^{-1}}{\phi} \cdot \log n\right)$ successful steps. Since each successful step gives a multiplicative volume gain of $1 + \Omega\left(\frac{\phi}{\log^2 \phi^{-1}}\right)$, we obtain a volume of

  $$\left(1 + \Omega\left(\frac{\phi}{\log^2 \phi^{-1}}\right)\right)^{\Theta\left(\frac{\log^2 \phi^{-1}}{\phi} \cdot \log n\right)} = e^{\Theta(\log n)},$$

  which, by choosing the right constants, is larger than $|E|$.

- by applying lemma 7, we can then then show by symmetry that each uninformed node can get the information in T rounds, if a set S of volume $> |E|$ is informed — completing the proof.

Given $v \in U = U_B(A)$, we define $\widehat{N}_U^{\circ\!\rightarrow}(v)$ (to be read "N-hat-push-U-v") and $\widehat{N}_B^{\circ\!\leftarrow}(v)$(to be read "N-hat-pull-U-v") as follows

$$\widehat{N}_B^{\circ\!\rightarrow}(v) = \{u \in N_B^+(v) \mid d(u) \geq \frac{1}{3} \cdot \phi^{-1} \cdot d_B^+(v)\}$$

and

$$\widehat{N}_B^{\circ\!\!\leftarrow}(v) = \{u \in N_B^+(v) \mid d(u) < \frac{1}{3} \cdot \phi^{-1} \cdot d_B^+(v)\}.$$

Then, we define,

$$\widehat{U}^{\circ\!\!\rightarrow} = \left\{v \in U \mid \left|\widehat{N}_B^{\circ\!\!\rightarrow}(v)\right| \geq \left|\widehat{N}_B^{\circ\!\!\leftarrow}(v)\right|\right\}$$

and

$$\widehat{U}^{\circ\!\!\leftarrow} = \left\{v \in U \mid \left|\widehat{N}_B^{\circ\!\!\leftarrow}(v)\right| > \left|\widehat{N}_B^{\circ\!\!\rightarrow}(v)\right|\right\}.$$

As before, if $\mathrm{vol}(\widehat{U}^{\circ\!\!\rightarrow}) \geq \frac{1}{2} \cdot \mathrm{vol}(U)$ we "apply" the `PUSH` strategy on U; otherwise, we "apply" the `PULL` strategy.

The following lemma is the crux of our analysis. It is a strengthening of lemma 9. A corollary of the lemma is that there exists a $p = p_v \geq \Omega(\phi)$, such that after p^{-1} rounds, with constant probability, node v lets us gain a new volume proportional to $\Theta\left(\frac{d^+_{B_i}(v)}{p \cdot \log \phi^{-1}}\right)$.

Lemma 12 *Assume $v \in U_i$. Then, $\widehat{N}_{B_i}^{\circ\!\!\leftarrow}(v)$ can be partitioned in at most $6 + \log \phi^{-1}$ parts, $S_1, S_2, \ldots$, in such a way that for each part i it holds that, for some $P_{S_i} \in \left(2^{-i+1}, 2^{-i}\right]$,*

$$\Pr\left[G_{S_i}(v) \geq \frac{|S_i|}{256 \cdot P_{S_i}}\right] \geq 1 - 2e^{-1},$$

where $G_{S_i}(v)$ is the total volume of nodes in S_i that perform a `PULL` from v in $P_{S_i}^{-1}$ rounds.

Lemma 9, that we used previously, only stated that with probability $\Omega(\phi)$ we gained a new volume of $\Theta(d^+_{B_i}(v))$ in a single step. If we do not allow v to go on for more than one step than the bounds of lemma 9 are sharp[3].

The insight of lemma 12 is that different nodes might require different numbers of rounds to give their "full" contribution in terms of new volume, but the more we have to wait for, the more we gain.

We now prove Lemma 12.

Proof of Lemma 12. We bucket the nodes in $\widehat{N}_{B_i}^{\circ\!\!\leftarrow}(v)$ in $K = K_{B_i}(v) = \left\lceil \lg \frac{d^+_{B_i}(v)}{3 \cdot \phi} \right\rceil$ buckets in a power-of-two manner. That is, for $j = 1, \ldots, K$, R_j contains all the nodes u in $\widehat{N}_{B_i}^{\circ\!\!\leftarrow}(v)$ having degree $2^{j-1} \leq d(u) < 2^j$. Observe that the R_j's are pairwise disjoint and that their union is equal to $\widehat{N}_{B_i}^{\circ\!\!\leftarrow}(v)$.

Consider the buckets R_j, with $j > \lg \phi^{-1}$. We will empty some of them, in such a way that the total number of nodes removed from the union of the bucket is an ϵ fraction of the total (that is, of $d^+_{B_i}(v)$). This node removal step is necessary for the following reason: the buckets R_j, with $j > \lg \phi^{-1}$, contain nodes with a degree so high that any single one of them will perform a `PULL` operation on v itself with probability strictly smaller than ϕ. We want to guarantee that the probability of a "gain" is at least ϕ we are forced to remove nodes having too high degree, if their number is so small that — overall — the probability of any single one of them to actually perform a `PULL` on v is smaller than ϕ.

[3]To prove this, we give two examples. In the first one, we show that the probability of informing *any* new node might be as small as $O(\phi)$. In the second, we show that a single step the gain might be only a ϕ fraction of the volume of the informed nodes. Lemma 12 implies that these two phenomena cannot happen together.

So, for the first example, take two stars: a little one with $\Theta\left(\phi^{-1}\right)$ leaves, and a large one with $\Theta(n)$ leaves. Connect the centers of the two stars by an edge. The graph will have conductance $\Theta(\phi)$. Now, suppose that the center and the leaves of the little star are informed, while the nodes of the large star are not. Then, the probability of the information to spread to any new node (that is, to the center of the large star), will be $O(\phi)$.

For the second example, keep the same two stars, but connect them with a path of length 2. Again, inform only the nodes of the little star. Then, in a single step, only the central node in the length-2 path can be informed. The multiplicative volume gain is then only $1 + O(\phi)$.

The node removal phase is as follows. If R'_j is the set of nodes in the j-th bucket after the node removal phase, then

$$R'_j = \begin{cases} R_j & |R_j| \geq \frac{1}{16} \cdot 2^j \cdot \phi \\ \varnothing & \text{otherwise} \end{cases}$$

Observe that the total number of nodes we remove is upper bounded by

$$\begin{aligned} \sum_{i=1}^{K} \left(\frac{1}{16} \cdot 2^i \cdot \phi \right) &\leq \frac{\phi}{16} \sum_{i=0}^{K} 2^i \leq \frac{\phi}{8} 2^{K+1} \\ &\leq \frac{\phi}{8} \cdot 4 \cdot \frac{d^+_{B_i}(v)}{3 \cdot \phi} = \frac{1}{6} \cdot d^+_{B_i}(v). \end{aligned}$$

Therefore, $\sum_j |R'_j| \geq \frac{1}{3} d^+_{B_i}(v)$, since $\sum_j |R_j| \geq \frac{1}{2} d^+_{B_i}(v)$.

Consider the random variable $g(v)$, which represent the total volume of the nodes in the different R'_j's that manage to pull the information from v. If we denote by $g_j(v)$ the contribution of the nodes in bucket R'_j to $g(v)$, we have $g(v) = \sum_{j=1}^{K} g_j(v)$.

Take any non-empty bucket R'_j. We want to show, via lemma 5, that

$$\Pr\left[g_j(v) \geq \frac{1}{128} \cdot \frac{|R'_j|}{p_j} \right] \geq p_j.$$

(If this event occurs, we say that bucket j succeeds.)

This claim follows directly from 5 by creating one X_i in the lemma for each $u \in R'_j$, and letting $X_i = 1$ iff node u pulls the information from v. The probability of this event is $p_i \in \left(2^{-j}, 2^{-j+1}\right]$, so we can safely choose the p of lemma 5 to be $p = 2^{-j+1}$.

Consider the different p_j's of the buckets. Fix some j. If p_j came out of case 2 then, since $|R'_j| \geq \frac{1}{16} \cdot 2^j \cdot \phi$, we have $p_j \geq \frac{1}{32} \cdot \phi$. If p_j came out of case 1, then $p_j = \frac{1}{32}$. In general, $p_j \geq \frac{\phi}{32}$, and $p_j \leq 1$.

Let us divide the unit into segments of exponentially decreasing length: that is, $\left[1, \frac{1}{2}\right), \left[\frac{1}{2}, \frac{1}{4}\right), \ldots, \left[2^{-j+1}, 2^{-j}\right), \ldots$. For each j, let us put each bucket R'_j into the segment containing its p_j. Observe that there are at most $\left\lceil \lg \frac{32}{\phi} \right\rceil \leq 6 + \lg \phi^{-1}$ segments.

Fix any non-empty segment ℓ. Let S_ℓ be the union of the buckets in segment ℓ. Observe that if we let the nodes in the buckets of S_ℓ run the process for 2^ℓ times, we have that, for each bucket R'_j,

$$\Pr\left[G_j(v) \geq \frac{1}{128} \cdot \frac{|R'_j|}{p_j} \right] \geq 1 - (1 - p_j)^{2^\ell} \geq 1 - e^{-1},$$

where $G_j(v)$ is the total volume of nodes in R'_j that perform a `PULL` from v in 2^ℓ rounds.

Now, we can apply lemma 4, choosing X_j to be equal $v_j = \frac{1}{128} \cdot \frac{|R'_j|}{p_j}$ if bucket R'_j in segment ℓ (buckets can be ordered arbitrarily) is such that $G_j(v) \geq v_j$, and 0 otherwise. Choosing $p = 1 - e^{-1}$ and $q = 1 - 2e^{-1}$, we obtain:

$$\Pr\left[G_{S_\ell}(v) \geq \frac{1}{256} \cdot \frac{|S_\ell|}{2^{-\ell}} \right] \geq 1 - 2e^{-1}.$$

□

The following corollary follows from lemma 12. (We prove in Appendix A that, constants aside, it is the best possible.)

Corollary 13 *Assume $v_i \in U_i$. Then, there exists $p_i \in \left(\frac{\phi}{64}, 1\right]$ such that*

$$\Pr\left[G(v_i) \geq \frac{d^+_{B_i}(v_i)}{5000 \cdot p_i \cdot \lg \frac{2}{\phi}}\right] \geq 1 - 2e^{-1}.$$

where $G(v_i)$ is the total volume of nodes in $N^+_{B_i}(v_i)$ that perform a `PULL` *from v_i, or that v_i pushes the information to, in p_i^{-1} rounds.*

Proof If $v_i \in \widehat{U}_i^{\circ\!\!\rightarrow}$, the same reasoning of lemma 9 applies. If $v_i \in \widehat{U}_i^{\circ\!\!\leftarrow}$, then we apply lemma 12 choosing the part S with the largest cardinality. By the bound on the number of partitions, we will have

$$S \geq \frac{d^+_{B_i}(v_i)}{3} \cdot \frac{1}{6 + \lg \phi^{-1}} \geq \frac{d^+_{B_i}(v_i)}{18 \cdot \lg \frac{2}{\phi}},$$

which implies the corollary. $\square$

We now prove the main theorem of the section:

Theorem 14 *Let S be the set of informed nodes, $\mathrm{vol}(S) \leq |E|$. Then, if S' is the set of informed nodes then there exists some $\Omega(\phi) \leq p \leq 1$ such that, after $O\left(p^{-1}\right)$ steps, then with $\Omega(1)$ probability,*

$$\mathrm{vol}(S') \geq \left(1 + \Omega\left(\frac{\phi}{p \cdot \log^2 \frac{1}{\phi}}\right)\right) \cdot \mathrm{vol}(S).$$

Corollary 13 is a generalization of lemma 9, which would lead to our result if we could prove an analogous of the *two-cases property* of the previous section. Unfortunately, the final gain we might need to aim for, could be larger than the cut — this inhibits the use of the *two-cases property*. Still, by using a strenghtening of the two-cases property, we will prove Theorem 14 with an approach similar to the one of Theorem 11.

Proof We say that an edge in the cut $(S, V - S)$ is *easy* if its endpoint w in $V - S$ is such that $\frac{d^-_S(w)}{d(w)} \geq \phi$. Then, to overcome the just noted issue, we consider two cases separately: (a) at least half of the edges in the cut are easy, or (b) less than half of the edges in the cut are easy.

In case (a) we bucket the easy nodes in $\Gamma(S)$ (the neighbourhood of S) in $\left\lceil \lg \frac{1}{\phi} \right\rceil$ buckets in the following way. Bucket D_i, $i = 1, \ldots, \left\lceil \lg \frac{1}{\phi} \right\rceil$, will contain all the nodes w in $\Gamma(S)$ such that $2^{-i} < \frac{d^-_S(w)}{d(w)} \leq 2^{-i+1}$. Now let D_j be the bucket of highest volume (breaking ties arbitrarily).

For any node $v \in D_j$ we have that its probability to pull the information in one step is at least 2^{-j}. So, the probability of v to pull the information in 2^j rounds is at least $1 - e^{-1}$. Hence, by applying lemma 4, we get that with probability greater than or equal to $1 - 2e^{-1}$ we gain a set of new nodes of volume at least $\frac{\mathrm{vol}(D_j)}{2}$ in 2^j rounds. But,

$$\frac{\mathrm{vol}(D_j)}{2} \geq 2^j \cdot \frac{\mathrm{cut}(S, D_j)}{2} \geq 2^j \cdot \frac{\mathrm{cut}(S, \Gamma(S))}{2\left\lceil \lg \frac{1}{\phi} \right\rceil} \geq 2^j \cdot \frac{\phi \cdot \mathrm{vol}(S)}{2\left\lceil \lg \frac{1}{\phi} \right\rceil}.$$

Thus in this first case we gain with probability at least $1 - 2e^{-1}$ a set of new nodes of volume at least $2^j \cdot \frac{\phi \cdot \mathrm{vol}(S)}{2\left\lceil \lg \frac{1}{\phi} \right\rceil}$ in 2^j rounds. By the reasoning presented at the beginning of section 5 the claim follows.

Now let us consider the second case, recall that in this case half of the edges in the cut point to nodes u in $\Gamma(S)$ having $\frac{d^-(u)}{d(u)} \geq \frac{1}{\phi}$.

We then replace the original *two-cases property* with the *strong two-cases property*:

(a') $\sum_{v \in H_t \cap (V-S)} d^-_{B_t}(v) \geq \frac{1}{1000} \cdot \mathrm{cut}(S, V-S)$, and

(b') $\sum_{v \in H_t \cap S} d^+_{B_t}(v) \geq \frac{249}{1000} \cdot \mathrm{cut}(S, V-S)$.

As before, at least one of (a') and (b') has to be true. If (a') happens to be true then we are done since the total volume of the new nodes obtains will be greater than or equal $\sum_{v \in H_t \cap (V-S)} d(v) \geq \phi^{-1} \sum_{v \in H_t \cap (V-S)} d^-_{B_t}(v) \geq \frac{1}{1000} \cdot \phi^{-1} \cdot \mathrm{cut}(S, V-S)$. By Corollary 13, we will wait at most w rounds for the cut $(S, V-S)$, for some $w \leq O\left(\phi^{-1}\right)$. Thus, if (a') holds, we are guaranteed to obtain a new set of nodes of total volume $\Omega\left(\frac{\mathrm{cut}(S,V-S)}{w}\right)$ in w rounds. Which implies our main claim.

We now show that if (b') holds, then with $\Theta(1)$ probability our total gain will be at least $\Omega\left(\frac{\mathrm{cut}(S,V-S)}{w \log^2 \phi^{-1}}\right)$ in w rounds, for some $w \leq \left(\phi^{-1}\right)$.

Observe that each $v_i \in H_t \cap S$, when it was considered by the process, was given some probability $p_i \in \left(\frac{\phi}{64}, 1\right]$ by Corollary 13. We partition $H_t \cap S$ in buckets according to probabilities p_i. The j-th bucket will contain all the nodes v_i in $H_t \cap S$ such that $2^{-j} < p_i \leq 2^{-j+1}$. Recalling that B_i is the set of informed nodes when node v_i is considered, we let F be the bucket that maximizes $\sum_{v_i \in F} d^+_{B_t}(v_i)$.

Then,

$$\sum_{v_i \in F} d^+_{B_t}(v_i) \geq \frac{249}{1000 \left\lceil \lg \frac{64}{\phi} \right\rceil} \cdot \mathrm{cut}(S, V-S) \tag{1}$$

By Corollary 13, we have that for each $v_i \in F$, there exists $p = p(F) \geq \frac{\phi}{64}$, such that with probability at least $1 - 2e^{-1}$, after $\left\lceil \frac{2}{p} \right\rceil$ round, we gain at least $\frac{1}{5000 \cdot p \cdot \lg \frac{2}{\phi}} \cdot d^+_{B_i}(v_i) \geq \frac{1}{5000 \cdot p \cdot \lg \frac{2}{\phi}} \cdot d^+_{B_t}(v_i)$ (since $i \leq t$ implies $B_i \subseteq B_t$). For each v_i, let us define the random variable X_i as follows: with probability $1 - 2e^{-1}$, X_i has value $\frac{1}{5000 \cdot p \cdot \lg \frac{2}{\phi}} \cdot d^+_{B_t}(v_i)$, and with the remaining probability it has value 0. Then, the gain of v_i is a random variable that dominates X_i. Choosing $q = 1 - \frac{5}{2}e^{-1}$, in lemma 4, we can conclude that

$$\begin{aligned} \Pr \quad & \left[\sum_{i:\, v_i \in F} X_i \geq \frac{4}{5} \cdot \sum_{v_i \in F} \left(\frac{1}{5000 \cdot p \cdot \lg \frac{2}{\phi}} \cdot d^+_{B_t}(v_i)\right)\right] \\ \geq \quad & 1 - \frac{5}{2}e^{-1} \geq 0.08. \end{aligned}$$

Thus, in $\left\lceil \frac{2}{p} \right\rceil$ rounds, with constant probability we gain at least

$$\frac{1}{6250 \cdot \lg \frac{2}{\phi}} \cdot \sum_{v_i \in F} \frac{d^+_{B_t}(v_i)}{p},$$

which by equation 1, it is lower bounded by

$$\begin{aligned} & \tfrac{1}{6250 \cdot p \cdot \lg \frac{2}{\phi}} \cdot \tfrac{249}{1000 \left\lceil \lg \frac{64}{\phi} \right\rceil} \cdot \mathrm{cut}(S, V-S) \geq \\ & \geq \Omega\left(\tfrac{\mathrm{cut}(S,V-S)}{p \cdot \log^2 \phi^{-1}}\right) \end{aligned}$$

Thus applying the reasoning presented at the beginning of the section the claim follows. $\square$

6 Push and Pull by themselves

We now comment on how one can change our analysis to get a bound of $O\left(c_\alpha \cdot \phi^{-1} \cdot \log \phi^{-1} \cdot \log n\right)$ on the completion time of PUSH or PULL by themselves. The main observation is that, if degrees of neighboring nodes have a constant ratio, then the probability that a single node $v_i \in S$ (in our analysis) to make a viable PUSH or to be hit by a PULL is $\Theta(\alpha\phi)$ (indeed, v will have at least $\phi \cdot d(v)$ neighboring nodes in $V - S$, each having at most its degree times α — an easy calculation shows that the probability bound for both PUSH and PULL is then $\Theta(\alpha\phi)$).

The following lemma is the analogous of Corollary 13. The main claim of this section follows using the same approach we used in the previous section. We only observe that since Lemma 15 does not lose a factor of $\Theta\left(\log \phi^{-1}\right)$ in the gain $G(v)$, the total number of rounds will be a $\Theta\left(\log \phi^{-1}\right)$ factor smaller than the PUSH-PULL bound of the previous section.

Lemma 15 *No matter which of the* PUSH*, or the* PULL*, strategy is used then for each $v \in U_i$ there exists a $p \geq \frac{\phi}{4\alpha}$ such that*

$$\Pr\left[G(v) \geq \frac{d^+(v)}{2p\alpha^2}\right] \geq 1 - e^{-1},$$

where $G(v)$ is the volume of new nodes informed by v in $\lceil p^{-1} \rceil$ rounds.

Proof By the uniformity condition (a) each of the $d^+_{B_i}(v)$ neighbours of v, that are outside of B_i, have degree within $\alpha^{-1} \cdot d(v)$ and $\alpha \cdot d(v)$; furthermore, since $v \in U_i$, (b) it holds that $\frac{d^+_{B_i}(v)}{d(v)} \geq \frac{\phi}{2}$.

Suppose the PUSH strategy is used. By (b) the probability that v pushes the information to some neighour outside B_i (obtaining a gain $g(v)$ of at least $\alpha^{-1} \cdot d(v)$, by (a)) is $\frac{d^+_{B_i}(v)}{d(v)} \geq \frac{\phi}{2}$. By setting p to this quantity we get that the probability of obtaining a gain of at least $\frac{d(v)}{\alpha} \geq \frac{d(v)}{2p\alpha} = \frac{d^+_{B_i}(v)}{2p\alpha^2}$ is at least p.

Suppose, instead, the PULL strategy is used. Then the probability that some neighbour of v outside B_i performs a PULL from v is

$$1 - \prod_{u \in N_{B_i}(v)} \left(1 - \frac{1}{d(u)}\right) \geq 1 - \prod_{u \in N_{B_i}(v)} \left(1 - \frac{1}{\alpha \cdot d(v)}\right) = 1 - \left(1 - \frac{1}{\alpha \cdot d(v)}\right)^{d^+_{B_i}(v)} \geq 1 - e^{-\frac{d^+_{B_i}(v)}{\alpha \cdot d(v)}} \geq \frac{d^+_{B_i}(v)}{2\alpha \cdot d(v)},$$

where the first inequality is justified by (a), and the latter two are classic algebraic manipulations. Using (a) again, we obtain that the probability of having a gain $g(v)$ of at least $\alpha^{-1} \cdot d(v)$ is at least $\frac{d^+_{B_i}(v)}{2\alpha \cdot d(v)}$; we then set p to $p = \frac{d^+_{B_i}(v)}{2\alpha \cdot d(v)} \geq \frac{\phi}{4\alpha}$.

For both strategies, with probability at least $p \geq \frac{\phi}{4}$, in a round the volume of the new nodes informed by v is at least $\frac{d^+_{B_i}(v)}{2p\alpha^2}$. Therefore in $\lceil p^{-1} \rceil$ rounds, v will inform node of volume $G(v) \geq \frac{d^+_{B_i}(v)}{2p\alpha^2}$ with probability at least $1 - e^{-1}$. $\square$

References

[1] P. Berenbrink, R. Elsässer, and T. Friedetzky. Efficient randomized broadcasting in random regular networks with appliations in peer-to-peer systems. In *Proceedings of the Twenty-Seventh Annual ACM Symposium on Principles of Distributed Computing, PODC*, pages 155–164, 2008.

[2] S. P. Boyd, A. Ghosh, B. Prabhakar, and D. Shah. Gossip algorithms: design, analysis and applications. *IEEE Transactions on Information Theory*, 52:1653–1664, 2006.

[3] F. Chierichetti, S. Lattanzi, and A. Panconesi. Rumor spreading in social networks. In *Proceedings of Automata, Languages and Programming, 36th Internatilonal Colloquium, ICALP*, pages 375–386, 2009.

[4] F. Chierichetti, S. Lattanzi, and A. Panconesi. Rumour spreading and graph conductance. In *Proceedings of Proceedings of the Twenty-first Annual ACM-SIAM Symposium on Discrete Algorithms, SODA*, 2010.

[5] A. J. Demers, D. H. Greene, C. Hauser, W. Irish, J. Larson, S. Shenker, H. E. Sturgis, D. C. Swinehart, and D. B. Terry. Epidemic algorithms for replicated database maintenance. In *Proceedings of the Sixth Annual ACM Symposium on Principles of Distributed Computing, PODC*, pages 1–12, 1987.

[6] B. Doerr, T. Friedrich, and T. Sauerwald. Quasirandom broadcasting. In *Proceedings of the Nineteenth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA*, pages 773–781, 2008.

[7] B. Doerr, T. Friedrich, and T. Sauerwald. Quasirandom rumor spreading: Expanders, push vs. pull, and robustness. In *Proceedings of Automata, Languages and Programming, 36th International Colloquium, ICALP*, pages 366–377, 2009.

[8] R. Elsässer. On the communication complexity of randomized broadcasting in random-like graphs. In *Proceedings of the 18th Annual ACM Symposium on Parallel Algorithms and Architectures, SPAA*, pages 148–157, 2006.

[9] U. Feige, D. Peleg, P. Raghavan, and E. Upfal. Randomized broadcast in networks. *Algorithms*, 1:128–137, 1990.

[10] T. Friedrich and T. Sauerwald. Near-perfect load balancing by randomized rounding. In *Proceedings of the 41st Annual ACM Symposium on Theory of Computing, STOC*, pages 121–130, 2009.

[11] A. Frieze and G. Grimmett. The shortest-path problem for graphs with random arc-lengths. *Discrete Applied Mathematics*, 10:57–77, 1985.

[12] R. Karp, C. Schindelhauer, S. Shenker, and B. Vöcking. Randomized rumor spreading. In *Proceedings of FOCS*, 2000.

[13] D. Kempe, A. Dobra, and J. Gehrke. Gossip-based computation of aggregate information. In *Proceedings of the 41st Annual Symposium on Foundations of Computer Science, FOCS*, pages 565–574, 2003.

[14] J. Leskovec, K. J. Lang, A. Dasgupta, and M. W. Mahoney. Statistical properties of community structure in large social and information networks. In *Proceedings of the 17th International Conference on World Wide Web, WWW*, pages 695–704, 2008.

[15] M. Mihail, C. H. Papadimitriou, and A. Saberi. On certain connectivity properties of the internet topology. *J. Comput. Syst. Sci.*, 72(2):239–251, 2006.

[16] D. Mosk-Aoyama and D. Shah. Fast distributed algorithms for computing separable functions. *Transactions on Information Theory*, 54:2997–3007, 2008.

[17] B. Pittel. On spreading a rumor. *SIAM Journal on Applied Mathematics*, 47:213–223, 1987.

[18] T. Sauerwald. On mixing and edge expansion properties in randomized broadcasting. In *Proceedings of Algorithms and Computation, 18th International Symposium, ISAAC*, pages 196–207, 2007.

[19] A. Sinclair and M. Jerrum. Approximate counting, uniform generation and rapidly mixing markov chains. *Information and Computation*, 82:93–133, 1989.

[20] D. A. Spielman and S.-H. Teng. Nearly-linear time algorithms for graph partitioning, graph sparsification, and solving linear systems. In *Proceedings of the 36th Annual ACM Symposium on Theory of Computing, STOC*, pages 81–90, 2004.

A Optimality of Corollary 13

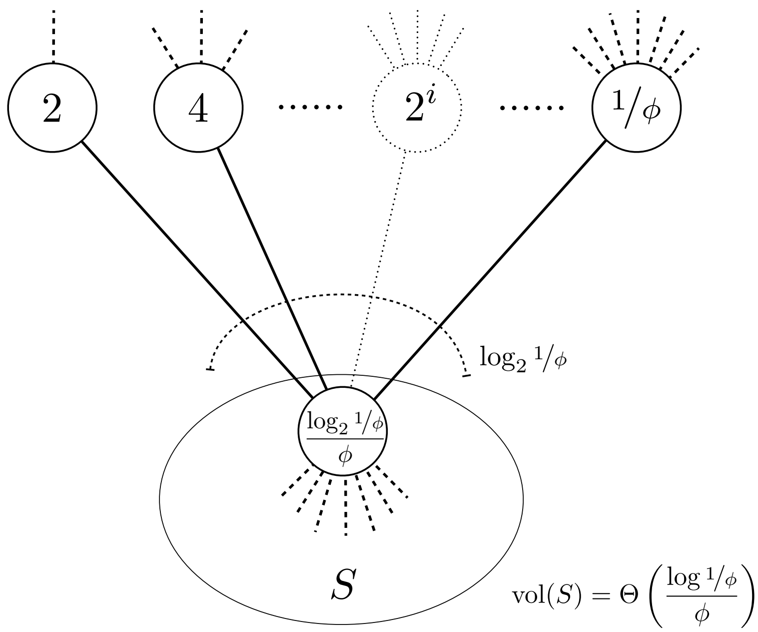

Figure 1: A construction showing that Corollary 13 is sharp. Each node is labeled with its degree.

Consider the cut in Figure 1, with $\phi = 2^{-t}$, for some integer $t \geq 1$. The set of informed nodes S is a star; its central node (shown in the figure), having degree $\frac{\log_2 1/\phi}{\phi}$, is connected to $\frac{\log_2 1/\phi}{\phi} - \log_2 1/\phi$ leaves inside S, and $\log_2 1/\phi$ nodes outside S. The volume of S is then $\mathrm{vol}(S) = \left(\frac{2}{\phi} - 1\right) \log_2 1/\phi$, and the conductance of the cut is then $\frac{\phi}{2+\phi} = \Omega(\phi)$. (It follows that, for any sufficiently large order, there exists a graph of that order with conductance $\Theta(\phi)$ that contains the graph in Figure 1 as a subgraph.) Finally, the i-th neighbour of S, $i = 1, \ldots, \log_2 1/\phi$, has degree 2^i.

Corollary 13, applied to our construction, gives that there exists some p, $\Omega(\phi) \leq p \leq 1$, such that the gain in p^{-1} rounds is $\Omega(p^{-1})$ with constant probability. (One can get a direct proof of this statement by analyzing the `PULL` performance.) We will show that Corollary 13 is sharp in the sense that, for each fixed constant $c > 0$, and for any p in the range, the probability of having a gain of at least cp^{-1} with no more than ϵp^{-1} rounds is $O(\epsilon)$.

Observe that the claim is trivial if $p > \epsilon$, since no gain can be obtained in zero rounds. We will then assume $p \leq \epsilon$. Because of the $O(\cdot)$ notation, we can also assume $\epsilon \leq c/8$. Therefore we prove the statement for $p \leq \frac{c}{8}$.

Let us analyze the `PUSH` strategy. Observe that the probability of performing a `PUSH` from S to the outside in $\epsilon\phi^{-1} \geq \Omega(\epsilon p^{-1})$ rounds is $(1-\phi)^{\epsilon\phi^{-1}} \leq \epsilon$. Therefore the probability of gaining anything with the `PUSH` strategy is at most ϵ.

Now let us analyze the `PULL` strategy. Fix any $\Omega(\phi^{-1}) \leq p \leq 1$. Let A be the set of the neighbours of S having degree less than $\frac{c}{2}p^{-1}$, and let B be the set of remaining neighbours. Then, the total volume of (and, thus, the total `PULL` gain from) nodes in A is not more than $cp^{-1} - 1$. Therefore to obtain the required gain, we need a node in B to make a `PULL` from S.

The probability that some node in B makes a `PULL` from S in one round is upperbounded by

$$\sum_{i=\lceil \log_2 \frac{c}{2}p^{-1}\rceil}^{\log_2 1/\phi} 2^{-i} = 2 \cdot \left(1 - 2^{-\log_2 1/\phi - 1} - 1 + 2^{-\lceil \log_2 \frac{c}{2}p^{-1}\rceil}\right) \leq \frac{2}{2^{\lceil \log_2 \frac{c}{2}p^{-1}\rceil}} \leq \frac{4p}{c}.$$

It follows that the probability P that some node in B performs a `PULL` from S in $k = \left\lceil \frac{\epsilon}{p} \right\rceil$ rounds is at most $P \geq 1 - \left(1 - \frac{4p}{c}\right)^k$. Since $p \leq c/8$ we have that $\left(1 - \frac{4p}{c}\right)^{\frac{c}{4p}} \geq \left(\frac{1}{2}\right)^2 = \frac{1}{4}$. Therefore,

$$P \geq 1 - 4^{-k \cdot \frac{4p}{c}} \geq 1 - 4^{-\epsilon \cdot \frac{1}{p} \cdot \frac{4p}{c}} = 1 - 4^{-\epsilon \cdot \frac{4}{c}} = \Theta\left(\frac{\epsilon}{c}\right) = \Theta(\epsilon).$$

The claim is thus proved.