

Reliable Broadcasting in Random Networks and the Effect of Density

Nikolaos Fountoulakis Anna Huber Konstantinos Panagiotou
Max-Planck-Institute for Informatics
Stuhlsatzenhausweg 85, Campus E1.4
Saarbrücken, D-66123 Germany

Abstract—Broadcasting algorithms are of fundamental importance for distributed systems engineering. In this paper we revisit the classical and well-studied *push* protocol for message broadcasting and we investigate a faulty version of it. Assuming that initially only one node has some piece of information, at each stage every one of the informed nodes chooses randomly and independently one of its neighbors and passes the message to it with some probability q ; that is, it fails to do so with probability $1 - q$.

The performance of the push protocol on a fully connected network, where each node is joined by a link to every other node, with $q = 1$ is very well understood. In particular, Frieze and Grimmett proved that with probability $1 - o(1)$ the push protocol completes the broadcasting of the message within $(1 \pm \varepsilon)(\log_2 n + \ln n)$ stages, where n is the number of nodes in the network. However, there are no tight bounds for the broadcast time on networks that are significantly sparser than the complete graph.

In this work we consider random networks on n nodes, where every edge is present with probability p , independently of every other edge. We show that if $p \geq \frac{\alpha(n) \ln n}{n}$, where $\alpha(n)$ is any function that tends to infinity as n grows, then the push protocol with faulty transmissions broadcasts the message within $(1 \pm \varepsilon) \left(\log_{1+q} n + \frac{1}{q} \ln n \right)$ stages with probability $1 - o(1)$. In other words, in almost every network of density d such that $d \geq \alpha(n) \ln n$, the push protocol broadcasts a message as fast as in a fully connected network and the speed is only affected by the success probability q . This is quite surprising in the sense that the time needed remains essentially *unaffected* by the fact that most of the links are missing. Our results are accompanied by experimental evaluation.

I. INTRODUCTION

We consider the problem of spreading information in large random networks with small average degree. Randomized broadcasting is among the most fundamental and well-studied communication primitives in distributed computing, and has also applications in several other disciplines, like e.g. in mathematical theories of epidemics. A particularly popular example [1] is the maintenance of consistency in a distributed database, which is replicated at many hundreds or thousands of sites in a large, heterogeneous network. Obviously, efficient broadcasting algorithms are crucial in order to ensure that all copies of the database converge quickly and effectively to the same content.

There is an enormous amount of literature devoted to the theoretical and experimental evaluation of broadcasting algorithms on several different underlying networks. Our interest in

considering *random* networks is motivated, among other reasons, by P2P (peer-to-peer) systems. The idea of using random graphs appears in some “real-life” networks, like the popular *Gnutella* network [2], or the *Juxtapose* protocol [3], which was originally developed by Sun Microsystems. Meanwhile, a considerable amount of work by several research groups aimed at designing many diverse networks for P2P systems that resemble properties of random graphs, see e.g. [4], [5], [6], and at developing protocols that perform efficiently on random (nearly) regular networks [7], [8], [9].

The most relevant properties of P2P networks, and more generally, of communication networks, are high expansion, connectivity, small average degree, and, (approximate) regularity of the degrees of the nodes. The random graph model considered in this paper has these properties. In particular, we investigate the classical Erdős-Rényi graph $G_{n,p}$, which is obtained by including each of the possible $\binom{n}{2}$ edges that connect any two out of n labeled vertices with probability p , independently of all other edges [10].

The Push Model

A classical protocol in the context of randomized broadcasting, which is also the main topic of our study, is the *push model* [11], [1]. Here we describe a general version of it, where errors may occur during the transmissions. Initially some information is placed on one of the nodes. In each succeeding stage, every informed node passes the information to another node, that it chooses uniformly at random and independently among its neighbors, with probability $q > 0$. The crucial question now is: how long does it take until all nodes have received the information? There are several advantages of considering a broadcast algorithm like this: it is simple, local, and scalable, and thus independent of the network topology. Moreover, it is robust against network and link failures, which makes it highly reliable.

In the case where the underlying network is the complete graph on n vertices, Frieze and Grimmett [11] proved that with high probability¹ (w.h.p.) the push protocol with $q = 1$ completes the broadcasting of the message within $(1 \pm \varepsilon)(\log_2 n + \ln n)$ stages. In other words, if a node can “talk” to *any other* node in the network, then the broadcast time will be almost surely very close to $\log_2 n + \ln n$. This bound

¹with probability tending to 1 when $n \rightarrow \infty$

was later improved by Pittel [12] to $\log_2 n + \ln n + \alpha(n)$, where $\alpha(n)$ is any function that tends to ∞ when $n \rightarrow \infty$. Feige et al. considered in [13] networks that are different from the complete graph. Among other results, they showed that if the underlying network is a random graph $G_{n,p}$, where $p \geq \frac{(1+\varepsilon)\ln n}{n}$, then the message will arrive at all nodes with high probability within $\Theta(\ln n)$ stages. Moreover, they also showed that the protocol is efficient on hypercubes, and derived bounds that hold for arbitrary graphs. Elsässer and Sauerwald determined in [14] similar bounds for several classes of Cayley graphs, thus generalizing upon [13].

Our contribution

Let $G = (V, E)$ be a graph on n vertices, where we will assume that $V = \{1, \dots, n\}$. We define $T(G)$ as the number of stages needed by the push protocol until all vertices have been informed, if the information is initially placed on node 1. In the remainder of the paper, we will be using the terms “node” and “vertex” without distinction. Note that *regardless* of the underlying network topology and the value of q we have $T(G) \geq \log_2 n$, as the number of informed vertices can at most double in each round. Consequently, all the results mentioned above state that the push model is, up to multiplicative constants, an optimal protocol for disseminating information in several important types of networks.

However, it is not at all well-understood *how much* the structure of the underlying network affects the performance of the push model. Although, for example, we know from the results in [13] that with $q = 1$ on a random graph $G_{n,p}$ the protocol requires with high probability at most $C \ln n$ rounds, for some $C > 0$, we have *a priori* no bounds that quantify how slower (or faster?) the protocol is compared to the case where the network is the complete graph. In particular, it is not clear in which way the average degree of the underlying graph influences the speed of the protocol. Our main result states that the number of stages is essentially *unaffected* by the density of the underlying graph, thus confirming the robustness and the efficiency of the push model.

Theorem I.1: Let $0 < \alpha(n) \leq \ln^{1/9} n$ be any function with the property $\lim_{n \rightarrow \infty} \alpha(n) = \infty$. Let $p \geq \frac{\alpha(n)\ln n}{n}$. Then w.h.p.

$$\left| T(G_{n,p}) - \left(\log_{1+q} n + \frac{1}{q} \ln n \right) \right| < \alpha(n)^{-1/7} \ln n.$$

For $q = 1$, the broadcast time *coincides* with the broadcast time on the complete graph, which was shown in [11] to be very close to $\log_2 n + \ln n$. If $q < 1$, the recent analysis in [15] regarding the complete graph has shown that the broadcast time is at least $\log_{1+q} n + \frac{1}{q} \ln n - o(\ln n)$. Our result shows that this bound is tight, and demonstrates that the number of stages needed is not influenced by the fact that most of the links are missing.

To avoid any confusion, we want to note that in Theorem I.1 the term “w.h.p.” refers to three independent probability spaces: first, the space from which we sample the underlying network, second, the space of the random choices performed by the

nodes, and third, the random faults that occur during the transmission.

Proof Ideas & Techniques

Before we proceed with a detailed exposition of our proof, let us mention a few words about the general strategy. Theorem I.1 is proved by bounding for each stage performed by the push model simultaneously from above and from below the number of informed nodes. In particular, we show that in the first $(1 - o(1)) \log_{1+q} n$ stages, the number of informed nodes increases roughly by a factor of $1 + q$ in each stage. As a result, we are able to show that after $(1 + o(1)) \log_{1+q} n$ rounds there will be εn informed nodes in total, where $\varepsilon = \varepsilon(n) \rightarrow 0$ will be defined during our proof. Then things evolve very fast: only after a small number of stages, the number of nodes having the information will be already at least $(1 - \varepsilon)n$. After that, we show that additionally approximately $\frac{1}{q} \ln n$ stages are necessary and sufficient to spread the information to everybody.

The analysis of the last stages is particularly challenging from a technical point of view, as the number of informed nodes increases only slowly towards the end of the process. In such cases, it is typically difficult to control the deviations of several involved random variables from their expectations. To this end, we exploit a modern and powerful tool from probability theory called *Talagrand’s inequality* [16], which – to our knowledge – has not been applied in the context of distributed computing problems. We believe that it could be widely applicable to the analysis of existing or future randomized protocols with several different degrees of dependency.

Outline

The next section presents some selected numerical results that we have obtained from extensive simulations of the push model on random graphs. Then, Section III introduces the main tools from probability theory that we will use, and in particular Talagrand’s inequality. In Section IV we collect and prove the basic properties of $G_{n,p}$ that will be important in the proof of Theorem I.1, and introduce some necessary notation that will be used throughout. Finally, we present the “core” of the proofs, where the general strategy described above is converted to a rigorous argument.

II. EXPERIMENTAL EVALUATION

In this section we present some simulation results that confirm the accuracy of our main result. In particular, we provide experimental evidence that the broadcast time is essentially independent of the density of the underlying random network, and that our lower bound on the average degree is an unavoidable ingredient of our proof.

Increasing densities: In our first experiment we generate a sequence G_0, G_1, \dots, G_{30} of random graphs, where G_i is distributed like G_{n,p_i} and $p_i = \frac{\ln^2 n}{n} + \frac{i}{30} \left(1 - \frac{\ln^2 n}{n}\right)$. In particular, the average degree of G_0 is whp. (close to) $\ln^2 n$, and G_{30} is just the complete graph on n vertices. Moreover, we set $q = 1$.

For each $i = 0, \dots, 30$, we execute the push protocol 500 times (each time independently) on G_i , and we determine the mean broadcast time T_i of these executions and the standard deviation σ_i . In all our figures, the values of the former are depicted by the dots and the standard deviation in each experiment is represented by a vertical interval of length $2\sigma_i$ centered at the dot that corresponds to T_i .

In Figure 1 we present the results of our first experiment, where we chose $n = 10^4$. We also depict the standard deviation σ_i against the error bounds $\pm \ln^{1/2} n$ from the value $\log_2 n + \ln n$ (the two enveloping lines). We observe that rarely the standard deviation violates these bounds. In other words, the broadcast time is barely affected by the density of the underlying network.

Increasing number of vertices: In our second experiment we study the effect of the network size and the success probability q on the broadcast time. Set $n_i = 1000 + 500i$, where $i = 0, \dots, 28$, and let G_i be a random graph with n_i vertices and edge probability $p_i = \frac{\ln^2 n_i}{n_i}$. That is, we cover a range from 1000 to 15000 vertices.

In Figure 2 we depict the average and the standard deviation out of 500 independent executions of the broadcast time on each of the G_i 's. This set of experiments has been performed for $q = 1, \frac{3}{4},$ and $\frac{1}{2}$. Moreover, the green (dashed) lines in Figure 2 represent the values predicted by Theorem I.1.

Sparse random graphs: The last experiment addresses our assumption on the average degree of the considered random graphs. In particular, we sampled graphs with $n = 10^4$ vertices and number of edges ranging from $5 \cdot 10^4$ to $5 \cdot 10^5$. Moreover, we used $q = 1$. In Figure 3 we present the average running time of the push algorithm on these graphs as well as the standard deviation of each set of runs. Note that when the number of edges decreases, the average running time becomes significantly larger than the $\log_2 n + \ln n$ estimate. More importantly, we observe larger fluctuations of the running times of the push algorithm, which actually increase as the density decreases. This suggests that the performance becomes gradually more and more uncertain for sparser networks.




Fig. 1. Simulation for $n = 10^4$ vertices and edge probabilities $p \geq \frac{\ln^2 n}{n}$.




Fig. 2. The results of simulations for n_i vertices.

III. PRELIMINARIES

A basic tool that we will use in the following proofs is the Chernoff bound. This provides exponentially small bounds for the probability that a binomially distributed random variable deviates significantly from its expected value.

Theorem III.1 (Chernoff Bounds): Let X be a binomially distributed random variable and let $x > 0$. Then

$$\mathbb{P}(|X - \mathbb{E}(X)| > x) \leq 2 \exp\left(-\frac{x^2}{2(\mathbb{E}(X) + x/3)}\right).$$

A more general tool that we shall apply several times is the inequality by Azuma and Hoeffding. Intuitively, it provides strong bounds for the probability that a function defined on a set of independent random variables deviates significantly from its expectation, when the value of the function is not affected much by small changes in each one of its arguments.

Theorem III.2 (Azuma-Hoeffding's Inequality): Let Z_1, \dots, Z_N be independent random variables taking values in the sets $\Lambda_1, \dots, \Lambda_N$ respectively. Let $\Lambda = \Lambda_1 \times \dots \times \Lambda_N$. Let $f : \Lambda \rightarrow \mathbb{R}$ be a function and set $X = f(Z_1, \dots, Z_N)$. Assume that there are quantities $c_k, k = 1, \dots, N$ satisfying the following condition.




Fig. 3. Simulations for sparse random graphs on $n = 10^4$ vertices.

- a. If $z, z' \in \Lambda$ differ only in the k th coordinate, then $|f(z) - f(z')| \leq c_k$.

Then, for every $x \geq 0$ we have that

$$\mathbb{P}(|X - \mathbb{E}(X)| \geq x) \leq 2 \exp\left(-\frac{x^2}{2 \sum_{i=1}^N c_i^2}\right). \quad (\text{III.1})$$

Note that the above inequality gives meaningful bounds only if the expectation of X is much larger than $(\sum_{i=1}^N c_i^2)^{1/2}$. This condition is unfortunately not always given in our intended applications. In such cases, we will use an estimate given by Talagrand (see the following theorem), which gives a much stronger tail bound, provided that an additional assumption is satisfied. Intuitively, the statement claims that if the value of X is “witnessed” by only a “small” number of its arguments, then X is sharply concentrated. However, there is a small caveat: the concentration is not guaranteed to be around the expectation, but instead around a *median* of X . (A median is a number m such that $\mathbb{P}(X < m) \leq \frac{1}{2}$ and $\mathbb{P}(X > m) \leq \frac{1}{2}$.) As we shall see below, this is not a significant problem as in general a median is very close the expected value.

Theorem III.3 (Talagrand’s Inequality, Thm. 2.29 in [17]): Suppose that the preconditions of Theorem III.2 are satisfied. Additionally, assume that there is an increasing function ψ satisfying the following condition.

- b. Let $z \in \Lambda$ and $r \in \mathbb{R}$ such that $f(z) \geq r$. Then there exists a set $J \subseteq \{1, \dots, N\}$ with $\sum_{i \in J} c_i^2 \leq \psi(r)$, such that for all $y \in \Lambda$ with $y_i = z_i$ when $i \in J$, we have $f(y) \geq r$.

Then, if m is a median of X , for every $x \geq 0$ we have

$$\mathbb{P}(|X - m| \geq x) \leq 4 \exp\left(-\frac{x^2}{4\psi(m+x)}\right). \quad (\text{III.2})$$

The next statement gives a sufficient condition that ensures that a median is very close to the expected value.

Proposition III.4 (Example 2.33 in [17]): Let X be a random variable that satisfies the preconditions of Theorem III.3 with $\psi(r) \leq \lceil r \rceil$. Then

$$|m - \mathbb{E}(X)| = O(\sqrt{\mathbb{E}(X)}). \quad (\text{III.3})$$

The presentation of the above inequalities is as in [17], where also many applications are presented.

IV. PROPERTIES OF $G_{n,p}$

For any graph G with vertex set V let $\Gamma_G(v)$ be the set of neighbors of v in G . If the underlying graph is clear, we will write $\Gamma(v)$. Moreover, for $S, S' \subseteq V$ we will denote by $e_G(S, S')$ the number of edges with one endpoint in each of S, S' . Finally, for two real numbers a, b we will write $a \pm b$ for the interval of reals $(a - b, a + b)$, and with slight abuse of notation we will write $X = a \pm b$ to denote $X \in a \pm b$.

Let $\alpha(n) \leq \frac{n}{\ln n}$ be any function with $\lim_{n \rightarrow \infty} \alpha(n) = \infty$ and let $p \geq \frac{\alpha(n) \ln n}{n}$. In this section we collect a few properties of $G_{n,p}$ that we will use in the proof of Theorem I.1.

Note that for any $S \subset V$, the expected number of neighbors of any $v \in V \setminus S$ in S is $p|S|$. The next lemma says that in

$G_{n,p}$ for all large enough S almost all vertices have roughly the right degree in S .

Lemma IV.1: The random graph $G_{n,p}$ has w.h.p. the following property. For any subset S of its vertices satisfying $|S| \geq \frac{n}{\alpha(n)}$, there is a set $X_S \subset V \setminus S$ that contains at most $\frac{4(\ln \alpha(n) + 2)n}{\ln n}$ vertices such that

$$\forall v \in (V \setminus S) \setminus X_S : |\Gamma_{G_{n,p}}(v) \cap S| = (1 \pm \alpha(n)^{-1/2})p|S|.$$

Proof: Let S be any fixed subset of the vertices such that $|S| \geq \frac{n}{\alpha(n)}$. Let $\varepsilon := \alpha(n)^{-1/2}$. We call a vertex $v \in V \setminus S$ *violating* with respect to S , if the number of its neighbors in S is $> (1 + \varepsilon)p|S|$ or $< (1 - \varepsilon)p|S|$. Assume there exist more than $t := \frac{4(\ln \alpha(n) + 2)n}{\ln n}$ vertices that are violating, and denote by X_S the set consisting of those vertices.

Note that the expected number of neighbors in S of a vertex is $p|S|$. By applying the Chernoff bounds, we obtain that the probability that a vertex is violating is for large n at most $e^{-\varepsilon^2 p|S|/4}$. Moreover, the events that two distinct vertices are violating are independent, which implies that the probability that there are t violating vertices is bounded from above by $e^{-\frac{\varepsilon^2 p|S|}{4}t}$. Hence, as there are $\binom{n}{k} \leq \left(\frac{en}{k}\right)^k \leq (e\alpha(n))^k$ ways to choose a set S of size k , the probability that there is a set of size k such that there are t violating vertices with respect to it is at most

$$\begin{aligned} & \exp\left\{k(\ln \alpha(n) + 1) - \frac{\varepsilon^2 pk}{4} \cdot t\right\} \\ & = \exp\left\{k\left(\ln \alpha(n) + 1 - \frac{\varepsilon^2 p}{4} \cdot \frac{4(\ln \alpha(n) + 2)n}{\ln n}\right)\right\}. \end{aligned}$$

This, combined with the bound $p \geq \frac{\alpha(n) \ln n}{n}$, can be estimated from above by e^{-k} . The proof is completed by summing this expression up for all $k \geq \frac{n}{\alpha(n)}$. ■

The next statement considers a similar setting as before, with the difference that now S might be very small. Here we show that the number of vertices that have many neighbors in S is only $o(|S|)$.

Lemma IV.2: For any $\varepsilon \geq \alpha(n)^{-1/2}$, the random graph $G_{n,p}$ has w.h.p. the following property. For any subset S of its vertices such that $|S| \leq \frac{n}{\alpha(n)}$ there is a set X_S containing at most $|S|\varepsilon^{-1}\alpha(n)^{-1}$ vertices, such that

$$\forall v \in (V \setminus S) \setminus X_S : |\Gamma_{G_{n,p}}(v) \cap S| \leq \varepsilon pn.$$

Proof: The proof is similar to the proof of Lemma IV.1, except that here we have to deal with small sets S . We give the whole proof for the sake of completeness. We assume that $|S| \geq \varepsilon pn$, for otherwise the statement holds trivially.

Let S be any fixed subset of the vertices such that $|S| \leq \frac{n}{\alpha(n)}$. We call a vertex $v \in V \setminus S$ *violating* with respect to S , if the number of its neighbors in S is $> \varepsilon pn$. Assume there exist more than $t := \frac{|S|}{\varepsilon \alpha(n)}$ vertices that are violating, and denote by X_S the set consisting of those vertices.

The expected number of neighbors in S of a vertex $v \in V \setminus S$ is $p|S| = o(\varepsilon pn)$. A straightforward application of the Chernoff bounds then implies that the probability that a

vertex is violating is for large n at most $e^{-\varepsilon pn}$. Hence, as the events that distinct vertices are violating with respect to S are independent, the probability that there are t such vertices is at most $e^{-\varepsilon pn \cdot t}$.

Note that the number of ways to choose S of size k is $\binom{n}{k} \leq (\frac{en}{k})^k$. In conclusion, the probability that there is an S with t violating vertices is at most

$$\begin{aligned} & \left(\frac{e}{k}\right)^k \exp\{k \ln n - \varepsilon pn \cdot t\} \\ & \leq \left(\frac{e}{k}\right)^k \exp\{k(\ln n - \ln n)\} = \left(\frac{e}{k}\right)^k. \end{aligned}$$

The proof then completes by summing this expression up for all $\varepsilon pn \leq k \leq \frac{n}{\alpha(n)}$. ■

Finally, we need the following statement about the distribution of the edges in $G_{n,p}$. The proof is a straightforward application of Chernoff's bounds, and quite standard in the classical random graph theory. We include a short proof for completeness.

Lemma IV.3: The following holds w.h.p.

$$\forall S \subseteq V : e_{G_{n,p}}(S, V \setminus S) = |S|(n - |S|)p \left(1 \pm \sqrt{\frac{8}{\alpha(n)}}\right).$$

Proof: It is sufficient to show the statements for S such that $|S| \leq n/2$. For any fixed such S , the quantity $e_{G_{n,p}}(S, V \setminus S)$ is binomially distributed with expectation $|S|(n - |S|)p$. Call S bad, if $e_{G_{n,p}}(S, V \setminus S)$ deviates from its expected value by more than $\sqrt{4|S|^2(n - |S|)p \ln n}$. Note that

$$\begin{aligned} \frac{\sqrt{4|S|^2(n - |S|)p \ln n}}{|S|(n - |S|)p} &= \sqrt{\frac{4 \ln n}{np(1 - |S|/n)}} \\ &\leq \sqrt{\frac{8 \ln n}{np}} \leq \sqrt{\frac{8}{\alpha(n)}}. \end{aligned}$$

By applying the Chernoff bounds we obtain that the probability that S is bad is with plenty of room to spare for large n at most

$$\exp\left\{-\frac{4|S|^2(n - |S|)p \cdot \ln n}{3|S|(n - |S|)p}\right\} = \exp\left\{-\frac{4}{3}|S| \ln n\right\}.$$

As the number of ways to choose S of size k is at most n^k , we infer by summing over all $1 \leq k \leq n/2$ that w.h.p. there is no bad set S in $G_{n,p}$. ■

Note that in the special case that $|S| = 1$ in the above lemma, i.e., S contains just a single vertex v , we obtain that

$$|\Gamma_{G_{n,p}}(v)| = e_{G_{n,p}}(S, V \setminus S) = (1 \pm 3\alpha(n)^{-1/2})pn.$$

We will use this fact without further reference.

V. BROADCASTING ON RANDOM GRAPHS

Let G be any graph with vertex set V and let $p \geq \frac{\alpha(n) \ln n}{n}$, where $\alpha(n) \leq \ln^{1/9} n$ is any positive function such that $\lim_{n \rightarrow \infty} \alpha(n) = \infty$. Fix

$$\varepsilon := \alpha(n)^{-1/2}.$$

We say that G is p -typical if it satisfies the following three conditions:

(I) For any $S \subseteq V$ such that $|S| \geq \frac{n}{\alpha(n)}$ there is a $X_S \subset V \setminus S$ such that $|X_S| \leq \frac{4(\ln \alpha(n) + 2)n}{\ln n}$ and

$$\forall v \in (V \setminus S) \setminus X_S : |\Gamma_G(v) \cap S| = (1 \pm \varepsilon)p|S|.$$

(II) For any $S \subseteq V$ such that $|S| \leq \frac{n}{\alpha(n)}$ there is a $X_S \subset V \setminus S$ such that $|X_S| \leq \frac{|S|}{\varepsilon \alpha(n)}$ and

$$\forall v \in (V \setminus S) \setminus X_S : |\Gamma_G(v) \cap S| \leq \varepsilon pn.$$

(III) For all $S \subseteq V$

$$e_G(S, V \setminus S) = |S|(n - |S|)p \left(1 \pm \sqrt{8\varepsilon}\right).$$

We will denote by $\mathcal{T}_n(p)$ the set of p -typical graphs on V . Note that Lemmas IV.1-IV.3 guarantee that $G_{n,p}$ is w.h.p. p -typical. Hence, we shall restrict our attention only to graphs in $\mathcal{T}_n(p)$.

Let us denote by $T_1(G)$ the first point in time where at least εn vertices are informed and $T_2(G)$ the first point in time where at least $(1 - \varepsilon)n$ vertices are informed. Our aim is to give bounds on $T(G)$ by bounding $T_1(G)$, $T_2(G) - T_1(G)$ and $T(G) - T_2(G)$ uniformly for every $G \in \mathcal{T}_n(p)$. The following three lemmas do so. In the proofs we will several times assume that n is sufficiently large so that the claimed inequalities hold, without explicitly mentioning that.

Lemma V.1: Uniformly for $G \in \mathcal{T}_n(p)$, with probability $1 - o(1)$ it holds that

$$|T_1(G) - \log_{1+q} n| \leq 18\sqrt{\varepsilon} \log_{1+q} n.$$

Proof: Let \mathcal{I}_t be the random set of informed vertices after t stages, and set $I_t := |\mathcal{I}_t|$. Note that our definitions imply that $\mathcal{I}_0 = \{1\}$.

Let $T'_1 = T'_1(G) \leq T_1(G) = T_1$ be the first time where the number of informed vertices exceeds $\ln^{1/4} n$. By Property (III) we have that if n is large enough, then for any $0 \leq t < T'_1$, each vertex in \mathcal{I}_t has at least $\frac{1}{2} \ln n$ neighbours outside \mathcal{I}_t . For n large enough, the number of stages needed until a new vertex is informed is stochastically bounded from above by a geometric random variable with parameter $q/2$. Therefore, the probability that at least k stages are needed until a new vertex is added in the set of informed vertices is at most $(1 - q/2)^k$. With $b := 1/(1 - q/2)$ and $k := \lceil \log_b \ln n \rceil$, this probability is at most $1/\ln n$. We call the set of stages between two consecutive stages at which the set of informed vertices increases an *idle period*. Therefore, the probability that each one of the idle periods before T'_1 lasts for less than k stages is at least $(1 - \frac{1}{\ln n})^{\ln^{1/4} n} \geq 1 - \frac{1}{\ln^{3/4} n}$, as there are no more than $\ln^{1/4} n$ idle periods before T'_1 . Consequently, $T'_1 \leq k \ln^{1/4} n$ with probability at least $1 - \ln^{-3/4} n$. Observe that for n that is sufficiently large the bound on T'_1 is at most $9\sqrt{\varepsilon} \log_{1+q} n$.

Next, we will show that with probability $1 - o(1)$, we have $T_1 - T'_1 \in (1 \pm 9\sqrt{\varepsilon}) \log_{1+q} n$, thus yielding the bound of the

lemma. In particular, we will show that

$$\mathbb{P}(I_{t+1} \in (1 + q \pm 7\sqrt{\varepsilon})I_t \mid \ln^{1/4} n \leq I_t < \varepsilon n) \geq 1 - o((\ln n)^{-1}). \quad (\text{V.1})$$

The proof of the lemma then completes by a repeated application of the above inequality. In particular, either there is a $t < (1 + 8\sqrt{\varepsilon}) \log_{1+q} n + T'_1$ such that $I_t \geq \varepsilon n$, in which case there is nothing to show, or, with probability $1 - o(1)$,

$$I_{\lceil (1+8\sqrt{\varepsilon}) \log_{1+q} n \rceil + T'_1} \geq (1 + q - 7\sqrt{\varepsilon})^{(1+8\sqrt{\varepsilon}) \log_{1+q} n} \geq \varepsilon n.$$

For all $t < (1 - 8\sqrt{\varepsilon}) \log_{1+q} n + T'_1$ we have $I_{t+1} \in (1 + q \pm 7\sqrt{\varepsilon})I_t$ with probability $(1 - o((\ln n)^{-1}))^{(1-8\sqrt{\varepsilon}) \log_{1+q} n} = 1 - o(1)$, and therefore

$$\begin{aligned} I_{\lfloor (1-8\sqrt{\varepsilon}) \log_{1+q} n \rfloor + T'_1} \\ \leq 2 \ln^{1/4} n (1 + q + 7\sqrt{\varepsilon})^{(1-8\sqrt{\varepsilon}) \log_{1+q} n} < \varepsilon n. \end{aligned}$$

So we showed that

$$\begin{aligned} (1 - 9\sqrt{\varepsilon}) \log_{1+q} n \leq \lfloor (1 - 8\sqrt{\varepsilon}) \log_{1+q} n \rfloor \leq T_1 - T'_1 \\ \leq \lceil (1 + 8\sqrt{\varepsilon}) \log_{1+q} n \rceil \leq (1 + 9\sqrt{\varepsilon}) \log_{1+q} n. \end{aligned}$$

In the remainder we prove (V.1). For every vertex $v \in I_t$ we define an indicator random variable N_v that equals 1 if v informs a vertex in $V \setminus \mathcal{I}_t$. Moreover, for every pair of distinct vertices $v, v' \in \mathcal{I}_t$ let $C_{v,v'}$ be the indicator variable that is equal to 1 if v and v' inform the same vertex in $V \setminus \mathcal{I}_t$. Finally, denote by \mathcal{N}_t the random set of vertices in $V \setminus \mathcal{I}_t$ that will be informed in stage $t + 1$ by the vertices in \mathcal{I}_t . By the inclusion-exclusion principle we obtain that

$$\sum_{v \in \mathcal{I}_t} N_v - \sum_{v, v' \in \mathcal{I}_t, v \neq v'} C_{v,v'} \leq |\mathcal{N}_t| \leq \sum_{v \in \mathcal{I}_t} N_v.$$

Note that

$$\begin{aligned} \mathbb{E}(N_v) &= q \frac{|\Gamma(v) \cap (V \setminus \mathcal{I}_t)|}{|\Gamma(v)|} \quad \text{and} \\ \mathbb{E}(C_{v,v'}) &= q^2 \frac{|\Gamma(v) \cap \Gamma(v') \cap (V \setminus \mathcal{I}_t)|}{|\Gamma(v)| |\Gamma(v')|}. \end{aligned} \quad (\text{V.2})$$

We shall now show that $|\mathcal{N}_t| \in (1 \pm 7\sqrt{\varepsilon})I_t q$ holds with the desired probability, which will complete the proof of (V.1). Before we proceed, let us make two auxiliary preparatory remarks. Note that by Property (III) of G we obtain for sufficiently large n that

$$\forall v \in V : |\Gamma(v)| = (1 \pm 3\varepsilon)pn.$$

This, together with (V.2) implies with a double counting argument that

$$\begin{aligned} \sum_{\substack{v, v' \in \mathcal{I}_t \\ v \neq v'}} \mathbb{E}(C_{v,v'}) &= \sum_{\substack{v, v' \in \mathcal{I}_t \\ v \neq v'}} q^2 \frac{|\Gamma(v) \cap \Gamma(v') \cap (V \setminus \mathcal{I}_t)|}{(1 \pm 7\varepsilon)(pn)^2} \\ &= \frac{(1 \pm 8\varepsilon)}{(pn)^2} q^2 \sum_{u \in V \setminus \mathcal{I}_t} \binom{|\Gamma(u) \cap \mathcal{I}_t|}{2}. \end{aligned}$$

We will use these facts in the remainder without further reference.

Recall that $I_t \geq \ln^{1/4} n$. We will first give tight bounds on the expectation of $|\mathcal{N}_t|$, and then apply the Azuma-Hoeffding inequality to show that $|\mathcal{N}_t|$ is sufficiently sharply concentrated around $\mathbb{E}(|\mathcal{N}_t|)$. By using (V.2) we obtain with plenty of room to spare for large n that

$$\begin{aligned} \mathbb{E}\left(\sum_{v \in \mathcal{I}_t} N_v\right) &= q \sum_{v \in \mathcal{I}_t} \frac{|\Gamma(v) \cap (V \setminus \mathcal{I}_t)|}{(1 \pm 3\varepsilon)pn} \\ &= q \frac{(1 \pm 4\varepsilon)e_G(\mathcal{I}_t, V \setminus \mathcal{I}_t)}{pn} \stackrel{(III)}{=} (1 \pm 8\varepsilon)I_t q. \end{aligned} \quad (\text{V.3})$$

Moreover, recall that

$$\sum_{v, v' \in \mathcal{I}_t, v \neq v'} \mathbb{E}(C_{v,v'}) = \frac{q^2(1 \pm 8\varepsilon)}{(pn)^2} \cdot \sum_{u \in V \setminus \mathcal{I}_t} \binom{|\Gamma(u) \cap \mathcal{I}_t|}{2}. \quad (\text{V.4})$$

We are going to estimate the last sum from above as follows. As $G \in \mathcal{T}_n(p)$ we may infer the following.

- If $I_t \leq \frac{n}{\alpha(n)}$, then, by (II), there is $\mathcal{X} \subset V \setminus \mathcal{I}_t$ such that $|\mathcal{X}| \leq \sqrt{\varepsilon}I_t$ and

$$\forall v \in (V \setminus \mathcal{I}_t) \setminus \mathcal{X} : |\Gamma(v) \cap \mathcal{I}_t| \leq \varepsilon pn.$$

- If $\frac{n}{\alpha(n)} \leq I_t \leq \varepsilon n = \frac{n}{\alpha(n)^{1/2}}$, then, by (I), there is $\mathcal{X} \subset V \setminus \mathcal{I}_t$ such that $|\mathcal{X}| \leq \frac{4(\ln \alpha(n) + 2)n}{\ln n}$ and

$$\forall v \in (V \setminus \mathcal{I}_t) \setminus \mathcal{X} : |\Gamma(v) \cap \mathcal{I}_t| \leq (1 + \varepsilon)pI_t \leq 2\varepsilon pn.$$

So, in both cases we have for all $v \in (V \setminus \mathcal{I}_t) \setminus \mathcal{X}$ that $|\Gamma(v) \cap \mathcal{I}_t| \leq 2\varepsilon pn$, and $|\mathcal{X}| \leq \sqrt{\varepsilon}I_t$. Moreover, by exploiting property (III) of G we obtain that for all $v \in \mathcal{X}$ it holds $|\Gamma(v) \cap \mathcal{I}_t| \leq 2pn$. Using this, we can bound from above the sum in (V.4) by splitting it into two parts as follows:

$$\begin{aligned} &\sum_{u \in V \setminus \mathcal{I}_t} |\Gamma(u) \cap \mathcal{I}_t|^2 \\ &= \sum_{u \in (V \setminus \mathcal{I}_t) \setminus \mathcal{X}} |\Gamma(u) \cap \mathcal{I}_t|^2 + \sum_{u \in \mathcal{X}} |\Gamma(u) \cap \mathcal{I}_t|^2 \\ &\leq \sum_{u \in (V \setminus \mathcal{I}_t) \setminus \mathcal{X}} |\Gamma(u) \cap \mathcal{I}_t|^2 + |\mathcal{X}| (2pn)^2 \\ &\leq \sum_{u \in (V \setminus \mathcal{I}_t) \setminus \mathcal{X}} |\Gamma(u) \cap \mathcal{I}_t|^2 + \sqrt{\varepsilon}I_t (2pn)^2. \end{aligned}$$

Note that $0 \leq |\Gamma(u) \cap \mathcal{I}_t| \leq 2\varepsilon pn$ for every $u \in (V \setminus \mathcal{I}_t) \setminus \mathcal{X}$. Moreover, it is easily seen that $\sum_{u \in V \setminus \mathcal{I}_t} |\Gamma(u) \cap \mathcal{I}_t| = e_G(\mathcal{I}_t, V \setminus \mathcal{I}_t)$. By the convexity of x^2 , an upper bound on the above sum can be obtained if we choose $|\Gamma(u) \cap \mathcal{I}_t| = 2\varepsilon pn$ for $e_G(\mathcal{I}_t, V \setminus \mathcal{I}_t)/(2\varepsilon pn)$ different u 's, and $|\Gamma(u) \cap \mathcal{I}_t| = 0$ otherwise. We obtain

$$\begin{aligned} &\sum_{u \in V \setminus \mathcal{I}_t} |\Gamma(u) \cap \mathcal{I}_t|^2 \\ &\leq \frac{I_t(n - I_t)p(2\varepsilon pn)^2}{\varepsilon pn} + \sqrt{\varepsilon}I_t (2pn)^2 \leq \frac{9}{2} \sqrt{\varepsilon} p^2 n^2 I_t. \end{aligned}$$

By plugging this into (V.4) we obtain that $\sum_{v, v' \in \mathcal{I}_t, v \neq v'} \mathbb{E}(C_{v,v'}) \leq 5q^2 \sqrt{\varepsilon}I_t$. Finally, combined

with (V.3) this gives with lots of room to spare that

$$(1 - 6\sqrt{\varepsilon})I_t q \leq \mathbb{E}(|\mathcal{N}_t|) \leq (1 + 6\sqrt{\varepsilon})I_t q.$$

To complete the proof we will bound the probability that $|\mathcal{N}_t| \notin I_t q(1 \pm 7\sqrt{\varepsilon})$ by using the Azuma-Hoeffding's inequality. Note that $|\mathcal{N}_t|$ can change by at most 1, if we modify one of the choices made by some vertex in \mathcal{I}_t . So, by applying Theorem III.2 with $c_i = 1$ and $N = I_t$ we obtain

$$\begin{aligned} & \mathbb{P}(|\mathcal{N}_t| \notin I_t q(1 \pm 7\sqrt{\varepsilon})) \\ & \leq \mathbb{P}(|\mathcal{N}_t| \notin \mathbb{E}(|\mathcal{N}_t|) \pm \sqrt{\varepsilon} I_t q) \leq 2e^{-\frac{q^2 \varepsilon}{2} \ln^{1/4} n}. \end{aligned}$$

In the next lemma we will consider the ‘‘intermediate’’ phase of the push model between $T_1(G)$ and $T_2(G)$ for $G \in \mathcal{T}_n(p)$. Our general strategy is to bound the number N_t of vertices that get informed in the current stage t from below. For this, we first estimate $\mathbb{E}(N_t)$ and then we use concentration inequalities (Theorem III.2) to show that with sufficiently high probability N_t is very close to $\mathbb{E}(N_t)$.

Lemma V.2: Uniformly for all $G \in \mathcal{T}_n(p)$, with probability $1 - o(1)$ it holds that

$$T_2(G) - T_1(G) \leq 9\varepsilon^{-1} \ln \varepsilon^{-1},$$

and there are at least $\varepsilon n / (2e)$ uninformed vertices at $T_2(G)$.

Proof: Let \mathcal{I}_t be the random set of informed vertices after t stages, and as before $I_t = |\mathcal{I}_t|$. We will show that for $T_1 = T_1(G) \leq t < T_2(G) = T_2$

$$I_{t+1} \geq I_t \left(1 + \frac{q\varepsilon}{4}\right), \quad (\text{V.5})$$

with probability at least $1 - e^{-q\varepsilon^6 n/8}$. Let us abbreviate $b = 8\varepsilon^{-1} \ln \varepsilon^{-1}$. To see that this is sufficient for the first claim, note that if ‘‘ $T_2 - T_1 \leq b$ ’’, then there is nothing to prove. On the other hand, if ‘‘ $T_2 - T_1 > b$ ’’, then with (conditional) probability at least $(1 - e^{-q\varepsilon^6 n/8})^{b+1} = 1 - o(1)$, for $\lceil b \rceil$ consecutive steps after T_1 the recursion (V.5) holds. In turn, this implies with $1 + x > e^{x/2}$, which is valid for small enough $x > 0$, that

$$I_{T_1 + \lceil b \rceil} \geq I_{T_1} \cdot \left(1 + \frac{\varepsilon}{4}\right)^{\lceil b \rceil} > \varepsilon n e^{b\varepsilon/8} > (1 - \varepsilon)n.$$

Thus we obtain with plenty of room to spare that, say, $T_2 - T_1 \leq b + 1 \leq 9\varepsilon^{-1} \ln \varepsilon^{-1}$.

Now we turn to the proof of (V.5). Let t be such that $T_1 \leq t < T_2$, and denote by \mathcal{N}_t the set of vertices in $V \setminus \mathcal{I}_t$ that will be informed by the vertices in \mathcal{I}_t in stage $t + 1$. Again, we write $N_t = |\mathcal{N}_t|$. We will show that N_t is not much smaller than its expected value. But first let us calculate $\mathbb{E}(N_t)$. The definition of the push model implies that the probability that any $v \in V \setminus \mathcal{I}_t$ does not belong to \mathcal{N}_t is precisely

$$\prod_{u \in \Gamma(v) \cap \mathcal{I}_t} \left(1 - \frac{q}{|\Gamma(u)|}\right).$$

Next we make use of property (I) in the definition of $\mathcal{T}_n(p)$: All vertices in $V \setminus \mathcal{I}_t$, apart from an exceptional set $\mathcal{X} = \mathcal{X}_t \subset V \setminus \mathcal{I}_t$ that contains at most $4(\ln \alpha(n) + 2)n / \ln n$ vertices, have

$(1 \pm \varepsilon)pI_t$ neighbors in \mathcal{I}_t . Using this and the above fact we infer that $\mathbb{E}(N_t)$ is in

$$\sum_{v \in (V \setminus \mathcal{I}_t) \setminus \mathcal{X}} \left(1 - \prod_{u \in \Gamma(v) \cap \mathcal{I}_t} \left(1 - \frac{q}{|\Gamma(u)|}\right)\right) \pm |\mathcal{X}|.$$

Next we derive tight bounds for the product above. Firstly, observe that property (III) implies that for all $u \in \mathcal{I}_t$

$$|\Gamma(u)| = np(1 \pm 3\varepsilon). \quad (\text{V.6})$$

Also, the definition of \mathcal{X} implies for $v \in (V \setminus \mathcal{I}_t) \setminus \mathcal{X}$

$$|\Gamma(v) \cap \mathcal{I}_t| = (1 \pm \varepsilon)pI_t. \quad (\text{V.7})$$

Recall that for $x > 0$ small enough we have $e^{-x-x^2} \leq 1 - x \leq e^{-x}$. So the bounds in (V.6) and (V.7) imply that for $v \in (V \setminus \mathcal{I}_t) \setminus \mathcal{X}$ we have

$$\begin{aligned} & \prod_{u \in \Gamma(v) \cap \mathcal{I}_t} \left(1 - \frac{q}{|\Gamma(u)|}\right) \\ & = \exp\left(-\frac{qI_t}{n}(1 \pm 5\varepsilon)\right) \left(1 - O\left(\frac{1}{np}\right)\right). \end{aligned}$$

As $|(V \setminus \mathcal{I}_t) \setminus \mathcal{X}| = (n - I_t)(1 \pm \varepsilon)$, we obtain

$$\mathbb{E}(N_t) = n \left(1 - \frac{I_t}{n}\right) \left(1 - e^{-\frac{qI_t}{n}}\right) (1 \pm O(\varepsilon)). \quad (\text{V.8})$$

We will bound the probability that $|N_t - \mathbb{E}(N_t)| > \varepsilon \mathbb{E}(N_t)$ using the Azuma-Hoeffding inequality. Firstly, note that as $\varepsilon < \frac{I_t}{n} \leq 1 - \varepsilon$, we have $\mathbb{E}(N_t) \geq q\varepsilon^2 n / 2$, for n sufficiently large. Moreover, if we change only one of the random choices of the vertices in \mathcal{I}_t , then N_t changes by at most 1. Thus, a simple application of Theorem III.2 with $c_k = 1$ and $N = I_t$ yields

$$\begin{aligned} & \mathbb{P}(|N_t - \mathbb{E}(N_t)| > \varepsilon \mathbb{E}(N_t)) \\ & \leq 2 \exp\left(-\frac{\varepsilon^2 \mathbb{E}^2(N_t)}{2I_t}\right) \leq 2 \exp\left(-\frac{q^2 \varepsilon^6 n}{8}\right). \end{aligned}$$

So, for n sufficiently large, with plenty of room to spare we obtain that

$$N_t = n \left(1 - \frac{I_t}{n}\right) \left(1 - e^{-\frac{qI_t}{n}}\right) (1 \pm \sqrt{\varepsilon}) \quad (\text{V.9})$$

with probability at most $2 \exp\left(-\frac{q^2 \varepsilon^6 n}{8}\right)$. This identity enables us to write a recursive formula concerning the evolution of the number of informed vertices. Recall that for all $0 < x < 1$, we have $1 - e^{-x} \geq x/2$. (V.9) implies that

$$\begin{aligned} I_{t+1} & \geq I_t + n \left(1 - \frac{I_t}{n}\right) \frac{qI_t}{2n} (1 - \sqrt{\varepsilon}) \\ & = I_t \left(1 + \frac{q}{2} \left(1 - \frac{I_t}{n}\right) (1 - \sqrt{\varepsilon})\right). \end{aligned} \quad (\text{V.10})$$

Since $I_t \leq (1 - \varepsilon)n$, it follows that for n large enough

$$\frac{q}{2} \left(1 - \frac{I_t}{n}\right) (1 - \sqrt{\varepsilon}) \geq \frac{\varepsilon q}{2} (1 - \sqrt{\varepsilon}) \geq \frac{\varepsilon q}{4}.$$

By substituting this bound into (V.10) we obtain (V.5).

What remains is to show the second statement of the lemma. This follows readily from (V.9). Indeed, if U_t denotes the number of uninformed vertices after t stages, then observe first that $n(1 - I_t/n) = U_t$. So, for n large enough

$$\begin{aligned} U_{T_2} &= U_{T_2-1} - N_{T_2-1} \\ &\stackrel{(V.9)}{\geq} U_{T_2-1} e^{-qI_{(T_2-1)}/n} (1 - e\sqrt{\varepsilon}) \geq \frac{\varepsilon n}{2e^q}. \end{aligned}$$

Finally, we proceed by bounding $T(G) - T_2(G)$, for $G \in \mathcal{T}_n(p)$. Recall that the main strategy in the previous argument was to show that the number N_t of vertices that become informed by \mathcal{I}_t in stage $t+1$ is close to its expected value. To achieve this, we exploited the fact that in G , apart from a small exceptional set, all vertices have the ‘‘right’’ degree in \mathcal{I}_t . This argument is unfortunately not applicable in the proof of the next lemma: for $t > T_2(G)$, the set $V \setminus \mathcal{I}_t$ of not yet informed vertices can become much smaller than \mathcal{X} , which makes our bounds useless. So we need to argue somehow differently. An additional difficulty is that we are not able to apply the Azuma-Hoeffding inequality in a meaningful way. Note that for $t > T_2(G)$ the quantity I_t is already of linear order, but the number N_t of newly informed vertices at stage $t+1$ may become very small. In this case, the Azuma-Hoeffding inequality gives a trivial upper bound and thus the need for a stronger concentration inequality.

Lemma V.3: Uniformly for all $G \in \mathcal{T}_n(p)$, with probability $1 - o(1)$

$$|(T(G) - T_2(G)) - \frac{1}{q} \ln n| \leq \varepsilon^{1/3} \ln n.$$

Proof: We will split the interval between $T_2 = T_2(G)$ and $T = T(G)$ into two subintervals. In particular, let $T' = T'(G)$ be the first time after T_2 where at most $\ln^{1/2} n$ uninformed vertices remain. We will give separate bounds for $T' - T_2$ and $T - T'$.

Let us start with the latter case, as it is the easier of the two. Let $t \geq T'$. Since $np > \alpha(n) \ln n$, it follows from property (III) that for every $v \in V \setminus \mathcal{I}_t$ we have for n large enough $|\Gamma(v) \cap \mathcal{I}_t| \geq np(1 - 3\varepsilon) - \ln^{1/2} n \geq np(1 - 4\varepsilon)$. So, the probability that a given uninformed vertex remains uninformed in the next stage is for large n at most

$$\left(1 - \frac{q}{np(1 + 3\varepsilon)}\right)^{np(1 - 4\varepsilon)} \leq e^{-q \frac{1 - 4\varepsilon}{1 + 3\varepsilon}} \leq \frac{2}{e^q}.$$

Therefore, the probability that such a vertex remains uninformed for at least $\ln^{1/2} n$ steps after T' is at most $(2/e^q)^{\ln^{1/2} n}$. This implies that the expected number of vertices that remain uninformed for at least $\ln^{1/2} n$ stages after T' is at most $\ln^{1/2} n (2/e^q)^{\ln^{1/2} n} \leq (2/e^q)^{\ln^{1/3} n} = o(1)$. That is, with probability at least $1 - (2/e^q)^{\ln^{1/3} n}$, we have $T - T' < \ln^{1/2} n$.

The bound on $T' - T_2$ is significantly more complex. Let U_t denote the number of vertices that are still uninformed after

the t th stage. We will show that if t is such that $U_t > \ln^{1/2} n$, then

$$U_{t+1} = U_t e^{-q} (1 \pm 50\sqrt{\varepsilon}), \quad (\text{V.11})$$

with probability at least $1 - e^{-q\varepsilon \ln^{1/2} n/40} \geq 1 - e^{-\ln^{1/3} n}$. So if $T' - T_2 > \frac{1}{q}(\ln n + 55\sqrt{\varepsilon} \ln n) =: \frac{b_1}{q}$, then with conditional probability at least $(1 - e^{-\ln^{1/3} n})^{\frac{b_1}{q}} = 1 - o(1)$ we have

$$U_{T_2 + \lceil \frac{b_1}{q} \rceil} \leq U_{T_2} e^{-b_1} (1 + 50\sqrt{\varepsilon})^{\frac{1}{q}(b_1+1)}.$$

For large n

$$(1 + 50\sqrt{\varepsilon})^{\frac{1}{q}(b_1+1)} \leq e^{\frac{55}{q}\sqrt{\varepsilon} \ln n}.$$

Also, $U_{T_2} \leq \varepsilon n$, which together with the above facts implies that $U_{T_2 + \frac{1}{q}(b_1+1)} \leq \varepsilon n$. So, we may conclude that $T' < T_2 + \frac{1}{q}(b_1 + 1)$.

Similarly, if we assume that $T' - T_2 < \frac{1}{q}(\ln n - 55\sqrt{\varepsilon} \ln n) =: \frac{b_2}{q}$, then with conditional probability at least $(1 - e^{-q\varepsilon \ln^{1/2} n/40})^{\frac{b_2}{q}+1} = 1 - o(1)$ we have

$$U_{T_2 + \lceil \frac{b_2}{q} \rceil} \geq U_{T_2} e^{-b_2-1} (1 - 50\sqrt{\varepsilon})^{\frac{b_2}{q}+1}.$$

A similar calculation as above, and the fact $U_{T_2} \geq \frac{\varepsilon n}{2e}$, which is guaranteed by Lemma V.2 to hold with probability $1 - o(1)$, show that

$$U_{T_2 + \lceil \frac{b_2}{q} \rceil} \geq \varepsilon e^{\frac{\sqrt{\varepsilon}}{q} \ln n} \gg \ln^{1/2} n.$$

Thus, $|T' - T_2 - \frac{1}{q} \ln n| \leq \frac{1}{q} 55\sqrt{\varepsilon} \ln n + 2$, which concludes the proof of the lemma.

It remains to show (V.11). As an auxiliary preparation we will show that ‘‘most’’ vertices in $V \setminus \mathcal{I}_t$ have the ‘‘right’’ degree in \mathcal{I}_t , by arguing that if this was not the case, then there would be a significant deviation in the number of edges between \mathcal{I}_t and $V \setminus \mathcal{I}_t$. More precisely, let

$$\mathcal{X} = \{v \in V \setminus \mathcal{I}_t : |\Gamma(v) \cap \mathcal{I}_t| < (1 - 3\sqrt{\varepsilon})pn\}.$$

In what follows, we argue that

$$|\mathcal{X}| \leq 3\sqrt{\varepsilon}(n - I_t). \quad (\text{V.12})$$

Indeed, as we assumed that $G \in \mathcal{T}_n(p)$, property (III) guarantees that $e_G(\mathcal{I}_t, V \setminus \mathcal{I}_t) \geq I_t(n - I_t)p(1 - 3\varepsilon)$. Moreover, it also implies that every vertex v has degree at most $(1 + 3\varepsilon)pn$. Therefore

$$e_G(\mathcal{I}_t, V \setminus \mathcal{I}_t) < |\mathcal{X}|(1 - 3\sqrt{\varepsilon})pn + (1 + 3\varepsilon)(n - I_t - |\mathcal{X}|)pn.$$

By putting the upper and the lower bounds together we obtain

$$I_t(n - I_t)p(1 - 3\varepsilon) \leq -3|\mathcal{X}|pn(\sqrt{\varepsilon} + \varepsilon) + (1 + 3\varepsilon)(n - I_t)pn,$$

which implies with $I_t \geq (1 - \varepsilon)n$ that

$$1 - 4\varepsilon \leq -3 \frac{|\mathcal{X}|}{n - I_t} (\sqrt{\varepsilon} + \varepsilon) + (1 + 3\varepsilon).$$

An elementary calculation shows that the claim (V.12) holds.

Now let $v \in (V \setminus \mathcal{I}_t) \setminus \mathcal{X}$. The probability that v becomes informed in the next stage is

$$\begin{aligned} & 1 - \prod_{u \in \Gamma(v) \cap \mathcal{I}_t} \left(1 - \frac{q}{|\Gamma(u)|} \right) \\ &= 1 - \left(1 - \frac{q}{pn(1 \pm 3\sqrt{\varepsilon})} \right)^{pn(1 \pm 3\sqrt{\varepsilon})} = 1 - \frac{1}{e^q} (1 \pm 7\sqrt{\varepsilon}). \end{aligned}$$

Denote by \mathcal{N}_t the set of vertices in $V \setminus \mathcal{I}_t$ that will be informed by the vertices in \mathcal{I}_t in stage $t+1$. Moreover, write $N_t = |\mathcal{N}_t|$. So, by linearity of expectation, for n large enough we obtain

$$\begin{aligned} \mathbb{E}(N_t) &= (n - I_t) \left(1 - \frac{1}{e^q} \right) (1 - 7\sqrt{\varepsilon}) \pm 3\sqrt{\varepsilon}(n - I_t) \\ &= (n - I_t) \left(1 - \frac{1}{e^q} \right) (1 \pm 14\sqrt{\varepsilon}). \end{aligned} \tag{V.13}$$

Next we will show that N_t is with sufficiently high probability close to its expected value. Note that the Azuma-Hoeffding inequality does not give any meaningful bounds, as the number of the independent random variables is $I_t \geq (1 - \varepsilon)n$, while the expected value of N_t is proportional only to $n - I_t$. The latter will eventually become so small that the exponent in the Azuma-Hoeffding inequality is $o(1)$, thus yielding a trivial bound. To bypass this problem, we will use Talagrand's inequality (Theorem III.3). Note first that the bounded differences condition is satisfied, that is, changing one random choice can change N_t by at most 1. Regarding the second condition, note that if $N_t = r$, then there must be at least r vertices in \mathcal{I}_t that have informed the vertices in \mathcal{N}_t . Therefore, we may take $\psi(r) = \lceil r \rceil$ and with $m(N_t)$ denoting the median of N_t , we deduce for any $x > 0$ that

$$\begin{aligned} \mathbb{P}(|N_t - m(N_t)| > x) &\leq 4 \exp\left(-\frac{x^2}{4(m(N_t) + x)}\right) \\ &\leq 4 \exp\left(-\frac{x^2}{4(2\mathbb{E}(N_t) + x)}\right), \end{aligned} \tag{V.14}$$

where in the last inequality we have used the fact that $\mathbb{E}(N_t) \geq m(N_t)\mathbb{P}(N_t > m(N_t)) \geq m(N_t)/2$, which implies that $m(N_t) \leq 2\mathbb{E}(N_t)$. However, we need to argue about the distance of $m(N_t)$ from $\mathbb{E}(N_t)$. We will use (III.3). The triangle inequality yields:

$$\begin{aligned} |N_t - \mathbb{E}(N_t)| &= |N_t - m(N_t) + m(N_t) - \mathbb{E}(N_t)| \\ &\leq |N_t - m(N_t)| + |m(N_t) - \mathbb{E}(N_t)| \\ &\stackrel{(III.3)}{=} |N_t - m(N_t)| + O\left(\sqrt{\mathbb{E}(N_t)}\right). \end{aligned}$$

Since $\alpha(n) \leq \ln^{1/9} n$, we have $\sqrt{\mathbb{E}(N_t)} = o(\sqrt{\varepsilon}\mathbb{E}(N_t))$. Therefore, for sufficiently large n

$$|N_t - \mathbb{E}(N_t)| > x \implies |N_t - m(N_t)| > x - \sqrt{\varepsilon}\mathbb{E}(N_t).$$

Using this in (V.14) with $x = \sqrt{\varepsilon}\mathbb{E}(N_t)$ we obtain

$$\mathbb{P}(|N_t - \mathbb{E}(N_t)| > 2\sqrt{\varepsilon}\mathbb{E}(N_t)) \leq 4 \exp\left(-\frac{\varepsilon\mathbb{E}(N_t)}{4(2 + \sqrt{\varepsilon})}\right).$$

Since $n - I_t \geq \ln^{1/2} n$, by (V.13) we obtain that, say, $\mathbb{E}(N_t) \geq \frac{q \ln^{1/2} n}{3}$. So, for large n , the above bound becomes

$$\mathbb{P}(|N_t - \mathbb{E}(N_t)| > 2\sqrt{\varepsilon}\mathbb{E}(N_t)) \leq \exp\left(-\frac{q\varepsilon \ln^{1/2} n}{40}\right).$$

By putting everything together we obtain that with probability at least $1 - e^{-q\varepsilon \ln^{1/2} n/40}$

$$N_t = (n - I_t) \left(1 - \frac{1}{e^q} \right) (1 \pm 16\sqrt{\varepsilon}).$$

So there remain (very generously) $(n - I_t)e^{-q}(1 \pm 50\sqrt{\varepsilon})$ vertices uninformed in $V \setminus \mathcal{I}_t$. This proves (V.11). ■

Finally, note that the bounds obtained in Lemmas V.1–V.3 imply Theorem I.1, thus concluding our proof.

REFERENCES

- [1] A. Demers, D. Greene, C. Hauser, W. Irish, J. Larson, S. Shenker, H. Sturgis, D. Swinehart, and D. Terry, "Epidemic algorithms for replicated database maintenance," in *PODC '87: Proceedings of the 6th annual ACM Symposium on Principles of distributed computing*, New York, NY, USA, 1987, pp. 1–12.
- [2] C. Inc., "Gnutella: To the bandwidth barrier and beyond," Published online, 2000.
- [3] S. Botros and S. Waterhouse, "Search in jxta and other distributed networks," in *1st IEEE International Conference on Peer-to-Peer Computing (P2P '01)*. Los Alamitos, CA, USA: IEEE Computer Society, 2001, pp. 30–35.
- [4] S. Jagannathan, G. Pandurangan, and S. Srinivasan, "Query protocols for highly resilient peer-to-peer networks," in *Proceedings of the ISCA 19th International Conference on Parallel and Distributed Computing Systems (ISCA PDCS '06)*, 2006, pp. 247–252.
- [5] G. Pandurangan, P. Raghavan, and E. Upfal, "Building low-diameter peer-to-peer networks," *IEEE Journal on Selected Areas in Communications*, vol. 21, no. 6, pp. 995–1002, 2003.
- [6] C. Law and K.-Y. Siu, "Distributed construction of random expander networks," in *In IEEE Infocom*, 2003, pp. 2133–2143.
- [7] P. Berenbrink, R. Elsässer, and T. Friedetzky, "Efficient randomised broadcasting in random regular networks with applications in peer-to-peer systems," in *PODC '08: Proceedings of the twenty-seventh ACM symposium on Principles of distributed computing*, 2008, pp. 155–164.
- [8] R. Elsässer, "On randomized broadcasting in power law networks," in *Proceedings of the 20th International Symposium on Distributed Computing (DISC '06)*, 2006, pp. 370–384.
- [9] T. S. R. Elsässer, L. Gasieniec, "On radio broadcasting in random geometric graphs," in *Proceedings of the 22nd International Symposium on Distributed Computing (DISC '08)*, 2008, pp. 212–226.
- [10] P. Erdős and A. Rényi, "On the evolution of random graphs," *Magyar Tud. Akad. Mat. Kutató Int. Közl.*, vol. 5, pp. 17–61, 1960.
- [11] A. Frieze and G. Grimmett, "The shortest-path problem for graphs with random arc-lengths," *Discrete Applied Mathematics*, vol. 10, pp. 57–77, 1985.
- [12] B. Pittel, "On spreading a rumor," *SIAM Journal on Applied Mathematics*, vol. 47, no. 1, pp. 213–223, 1987.
- [13] U. Feige, D. Peleg, P. Raghavan, and E. Upfal, "Randomized broadcast in networks," *Random Structures and Algorithms*, vol. 1, no. 4, pp. 447–460, 1990.
- [14] R. Elsässer and T. Sauerwald, "Broadcasting vs. mixing and information dissemination on Cayley graphs," in *24th Int. Symp. on Theor. Aspects of Computer Science (STACS)*, 2007, pp. 163–174.
- [15] B. Doerr, A. Huber, and A. Levavi, "Strong robustness of randomized rumor spreading protocols," in *Proceedings of the 20th International Symposium on Algorithms and Computation (ISAAC 2009)*, 2009, pp. 812–821.
- [16] M. Talagrand, "Concentration of measure and isoperimetric inequalities in product spaces," *Inst. Hautes Études Sci. Publ. Math.*, vol. 81, pp. 73–205, 1995.
- [17] S. Janson, T. Łuczak, and A. Ruciński, *Random graphs*, ser. Wiley-Interscience Series in Discrete Mathematics and Optimization. Wiley-Interscience, New York, 2000.