

Randomized Rumor Spreading: the Effect of the Network Topology

Konstantinos Panagiotou ^{*} Xavier Pérez-Giménez [†] Thomas Sauerwald [‡]
He Sun [§]

January 30, 2013

Abstract

We consider the popular and well-studied *push model*, which is used to spread information in a given network with n vertices. Initially, some vertex owns a rumor and passes it to one of its neighbors, which is chosen randomly. In each of the succeeding rounds, every vertex that knows the rumor informs a random neighbor. It has been shown on various network topologies that this algorithm succeeds in spreading the rumor within $O(\log n)$ rounds. However, many analyses are quite coarse and involve huge constants that do not allow for a direct comparison between different network topologies. We prove the following results.

- Let G be a graph with maximum degree Δ and minimum degree δ such that $\Delta/\delta = 1 + o(1)$ (as the number of vertices n grows large). If in addition the spectral expansion of G is $o(\Delta)$, rumor spreading completes after $\log_2 n + \log n + o(\log n)$ rounds with high probability. This is the first result that exhibits a general graph class for which rumor spreading is essentially as fast as on complete graphs. We also demonstrate that both assumptions are best possible.
- For a random graph $G(n, p)$ with $p = c \log n/n$, c being a constant larger than 1, the runtime is $\log_2 n + \gamma(c) \log n$, where $\gamma(c) = c \log(\frac{c}{c-1})$. This shows that the assumption $\Delta/\delta = 1 + o(1)$ used in our first result is necessary.
- For a hypercube, the runtime is at least $(1 + \beta) \cdot (\log_2 n + \log n)$, where β is a constant larger than 0. Together with the upper bound of $O(\log n)$ from [12], this implies with our second result that the push model is faster on hypercubes than on random graphs with comparable density, but slower than on complete graphs.

Mathematics Subject Classification [MSC2010]: 05C85, 68Q87, 68W20.

^{*}Mathematisches Institut, Universität München, Theresienstr. 39, 80333 München, Germany.
Email: kpanagio@math.lmu.de

[†]Max Planck Institute for Informatics, Campus E1.4, 66123 Saarbrücken, Germany.
Email: xperez@mpi-inf.mpg.de

[‡]Max Planck Institute for Informatics, Campus E1.4, 66123 Saarbrücken, Germany.
Email: sauerwal@mpi-inf.mpg.de

[§]Max Planck Institute for Informatics, Campus E1.4, 66123 Saarbrücken, Germany.
Email: hsun@mpi-inf.mpg.de

1 Introduction

Randomized broadcasting is one of the most important communication primitives in large networks. A classical and well-studied protocol is the following algorithm, which is known in the literature as the *push model* or *randomized rumor spreading*. Initially, some rumor is placed on one of the n vertices of a given network G . Then, in succeeding rounds, every vertex that knows the rumor selects one of its neighbors uniformly at random, and passes the information to it. The crucial question is: how many rounds are needed until every vertex becomes informed?

The push model has been studied in many works, and its performance on several different families of graphs is well-understood. In one of the first papers dealing with this topic, Frieze and Grimmett [15] proved that if the underlying graph is the complete graph with n vertices, then asymptotically almost surely¹ (a.a.s.) the broadcast is completed within $(1 + o(1))(\log_2 n + \log n)$ rounds, where $\log n$ will denote the natural logarithm. The problem has been subsequently studied on a number of graph classes, such as hypercubes, bounded-degree graphs, and Erdős-Rényi random graphs, see e.g. the works [12, 21]. In particular, for hypercubes and Erdős-Rényi random graphs $G(n, p)$ with $p = c \log n/n$, $c > 1$, a runtime bound of $O(\log n)$ was shown in [12]. In [10], the authors proved a lower bound of $\log_2 n + \log n - o(\log n)$ which holds for any regular graph. Neglecting low order terms, this implies that complete graphs have the fastest broadcast time among regular graphs.

Most of the existing bounds for the performance of the push model on general graphs show that a.a.s. the number of rounds needed is $O(f(G) \text{polylog}(n))$, where $f(G)$ is some graph parameter. In particular, Giakkoupis et al. [18] showed that $f(G)$ can be chosen as the inverse vertex expansion of G . Moreover, Giakkoupis [17] and Chierichetti, Lattanzi and Panconesi [6] proved, among other things, that $f(G) \leq c_\alpha \phi^{-1}$, where ϕ is the conductance of G and c_α a quantity depending on the ratio α of the maximum and minimum degree (see also [25] for a result for regular graphs). All the above results imply that if the parameter in question is within reasonable bounds, then we obtain a guaranteed logarithmic broadcast time. However, almost no work addresses the issue of *how exactly* the network topology affects the performance of the push model. In other words, we are interested in the structural properties that may have a favorable or a disadvantageous effect on the broadcast time. Resolving such questions is a fundamental issue in network design, and for practical applications, it is important to study the constants that are hidden in the $O(\cdot)$ -notation. Unfortunately, many theoretical analyses are rough, so that the involved constants are typically huge.

In this context, a precise analysis of the push model was performed by Fountoulakis, Huber and Panagiotou [14], who studied the case where G is a Erdős-Rényi random graph. We will denote by $G(n, p)$ a graph on n vertices that is obtained by including each possible edge independently with probability p . Among other results, they showed that if the average degree is $\omega(\log n)$, then a.a.s. the broadcast time coincides asymptotically with the broadcast time on the complete graph. So, the performance of the push model remains essentially unaffected by the fact that most edges are missing. Moreover, Fountoulakis and Panagiotou studied in [13] the case where $G(n, d)$ is a random d -regular graph, where $d = O(1)$, and determined the exact effect of d on the broadcast time.

Our Results. In this paper we perform a precise analysis of the push protocol on several graphs. Our first result addresses the broadcast time of the push model on *expander graphs*, which have found numerous applications in computer science and mathematics, see the survey [19]. The crucial

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

property of an expander graph is that every set of vertices is connected to the rest of the graph by a large number of edges. Here we focus on a spectral characterization of such graphs, which is related to the spectral gap of their adjacency matrix.

Let G be a connected graph with n vertices that has minimum degree equal to δ and maximum degree equal to Δ . Let $\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_n$ be the eigenvalues of the adjacency matrix of G , and set $\lambda = \max_{2 \leq i \leq n} |\lambda_i|$. We will say that G is an $(n, \delta, \Delta, \lambda)$ -graph. Our statements about G should be interpreted in the context of a sequence of graphs indexed by n , where Δ , δ and λ are functions of n . We are interested in asymptotic results as n tends to infinity. Our main result about expander graphs states that the broadcast time of the push model is asymptotically the same as on the complete graph. In other words, expander graphs belong to the same “universality class” with respect to the performance of the push model. We will denote in the remainder by $T(G)$ the (random) broadcast time if the underlying graph is G , and the rumor is placed initially on the vertex with label 1 (or equivalently, the rumor is placed on a vertex chosen u.a.r.).

Theorem 1.1. *Let G be an $(n, \delta, \Delta, \lambda)$ -graph, where $\Delta/\delta = 1 + o(1)$ and $\lambda = o(\Delta)$, as $n \rightarrow \infty$. Then, a.a.s.*

$$|T(G) - (\log_2 n + \log n)| = o(\log n).$$

Note that the above theorem, together with some well-known facts about random graphs, see [16, 26], also implies the main result in [14], where the same bounds on the broadcast time were shown for the special case that $G = G(n, p)$ and $p = \omega(\frac{\log n}{n})$. Moreover, Theorem 1.1 also applies to random d -regular graphs, where d can be any increasing function of n such that $d = o(\sqrt{n})$, as for this range of d , $\lambda = O(\sqrt{d})$ was shown in Broder et al. [5].

We now demonstrate that the two conditions in Theorem 1.1 are best possible in the sense that if we replace any of two $o(\cdot)$ terms in the statement by $O(\cdot)$, then there are graphs that do not satisfy the conclusion. For instance, consider a random d -regular graph with constant d . It trivially satisfies the first condition, but satisfies only $\lambda \leq 3\sqrt{d}$ [5]. Since G is a constant-degree, regular graph, it follows by [10, Theorem 1 & Lemma 1] that $T(G(n, d)) = \log_2 n + \log n + \Omega(\log n)$, a.a.s. Let us now address the condition on Δ/δ . It is well known ([16, 26], cf. Lemma A.1 and Lemma 3.2) that a random graph $G(n, p)$ with $p = \frac{c \log n}{n}$ for $c > 1$ satisfies $\lambda = o(\Delta)$ and $\Delta/\delta = \Theta(1)$. However, the next theorem implies that $T(G(n, p)) = \log_2 n + \log n + \Omega(\log n)$, a.a.s.

Theorem 1.2. *Let $c > 1$ be any constant and $p = \frac{c \log n}{n}$. Set $\gamma(c) = c \log(\frac{c}{c-1}) > 1$. Then, a.a.s.*

$$|T(G(n, p)) - (\log_2 n + \gamma(c) \log n)| = o(\log n).$$

Moreover, this result settles an important question that was left open by the previous results by extending the analysis of the performance of the push model to sparser random graphs with $p = \frac{c \log n}{n}$ for a constant $c > 1$ (see [12, 14] and Table 1). In particular, it shows that the broadcast time for a constant $c > 1$ is larger than for $c = \omega(1)$. Moreover, if $c < 1$, then $G(n, p)$ is a.a.s. not connected, see e.g. [20]. So, a complete broadcast is not possible in this case. Observe that $\gamma(c) \rightarrow 1$ when $c \rightarrow \infty$, which nicely matches the result of Theorem 1.1. On the other hand, $\gamma(c) \rightarrow \infty$ when $c \rightarrow 1$.

Our final result addresses the performance of the push model on hypercubes, which constitute popular topologies for the analysis of algorithms in distributed systems. We show that the push model on the hypercube is slower than on the complete graph, but faster than on a random graph $G(n, p)$ with $p = c \log n/n$ where c is sufficiently large. In other words, the regular degree distribution of the hypercube seems to be of more help for rumor spreading than the higher expansion of a random graph. Interestingly, our analysis reveals that on random graphs, the number of informed vertices nearly doubles in each round, as long as their number is $o(n)$. In contrast to that,

Graph	Asymptotic Broadcast Time	Reference
complete graph	$\log_2 n + \log n$	[15]
$G(n, p)$, $p = \omega(\log n)/n$	$\log_2 n + \log n$	[14]
$G(n, d)$, $d = O(1)$	$\log_{2(1-1/d)} n + \log_{(1-1/d)-d} n$	[13]
d -reg. graph, $\lambda = O(\sqrt{d})$, $d = \omega(\sqrt{n})$	$\log_2 n + \log n$	[13]
Hypercube	$O(\log n)$	[12]
$G(n, p)$, $p = c \log n/n$, $c > 1$	$\log_2 n + \gamma(c) \log n$, $\gamma(c) > 1$	Thm. 1.2
$G(n, d)$, $d = \omega(1)$	$\log_2 n + \log n$	Thm. 1.1
Graph with $\Delta/\delta = 1 + o(1)$, $\lambda = o(\Delta)$	$\log_2 n + \log n$	Thm. 1.1
Hypercube	$\log_2 n + \log n + \Omega(\log n)$	Thm. 1.3

Figure 1: Overview of the previous and our new results concerning the push model.

the growth is significantly smaller on the hypercube for the same range. However, to inform the remaining vertices, the random graph suffers from its heterogeneous structure, i.e., there are too many vertices of small degree that are connected to higher degree vertices.

Theorem 1.3. *Let H_n be a hypercube with n vertices, where n is a power of 2. Then there is a constant $\beta > 0$ such that a.a.s. $T(H_n) \geq (1 + \beta) \cdot (\log_2 n + \log n)$.*

Our results, together with some earlier bounds, are summarized in Table 1.

Techniques & Methods. One of the most fundamental results in spectral graph theory is the famous *expander mixing lemma*, see [2]. Let G be an n -vertex d -regular graph such that the adjacency matrix has eigenvalues $\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_n$. Then, roughly speaking, the expander mixing lemma guarantees that the number of edges between any two subsets of the vertices is close to what one would expect in a *random* d -regular graph, if $\lambda = \max_{2 \leq i \leq n} |\lambda_i|$ is not too large. Unfortunately, we cannot apply this lemma directly to support us in the proof of Theorem 1.1, as the relevant graphs are only close to being regular. In order to overcome this limitation, we prove in Section 2 a general statement, which provides similarly sharp bounds for almost regular graphs. Based on this, we perform a tight analysis of the push model by separating the evolution of the set of informed vertices into phases, and controlling the growth in each phase individually.

In Section 3 we prove Theorem 1.2. The proof is based mainly on techniques from the theory of random graphs, which allow us to prove various facts about the degree sequence and the distribution of the edges in certain structures. The most challenging aspect in this proof is that $G(n, p)$, for the range of parameters considered here, is quite far from being a regular graph. This introduces a significant bias in the actual evolution of the set of informed vertices. However, also in this case it is possible to perform a tight analysis of the push model, by using our knowledge of the degree sequence and from the fact that a.a.s. all vertices have the property that most of their neighbours have degree close to average. We hope that similar methods might be useful for studying randomized broadcasting on other classes of random graphs with stronger fluctuations in their degree sequences.

Finally, a further important ingredient in our proofs is a general purpose technique that we use for obtaining lower bounds for the broadcast time in terms of *random walks*. While the relation between random walks and rumor spreading has been studied before (e.g., [4, 11, 25]), all the obtained relations are of asymptotic nature. Our new approach establishes a direct link between rumor spreading and the transition probabilities of random walks, which is described in more detail in Section 4. Theorem 1.3 is a consequence of this general principle.

Notation. In order to avoid ambiguities we will introduce some notation that will be used throughout the paper. Let $G = (V, E)$ be a graph, v a vertex of G , and U, W two subsets of

its vertices. We will write $N_G(v)$ for the set of neighbors of v in G , and $\deg_G(v) = |N_G(v)|$ for its degree. For any vertex set S , let $N(S) = \cup_{v \in S} N(v)$. Moreover, we will write $E_G(U, W)$ for the set of edges with one endpoint in U and one in W , and $e_G(U, W) = |E_G(U, W)|$. We will abbreviate $E_G(U, U) = E_G(U)$. Finally, we will omit the subscript if the graph in question is clear from the context. For any round t , we denote by I_t the set of informed vertices at the end of round t . Similarly, $U_t := V \setminus I_t$ is the set of uninformed vertices. Recall that $\log x$ denotes the logarithm of x to the base e .

2 Rumor Spreading on Expander Graphs

Throughout this section, let $d = 2|E|/n$ be the average degree of G . Our main goal here is to prove Theorem 1.1.

2.1 Expansion Lemmas

The first ingredient in our proofs addresses the structure of regular graphs. The following statement is a strengthened version of the famous expander mixing lemma, where in comparison to the bound below, the factors $(1 - |U|/n)(1 - |W|/n)$ do not occur.

Lemma 2.1 ([22, Theorem 2.11]). *Let $G = (V, E)$ be an (n, d, d, λ) -graph. Then for any two subsets $U, W \subseteq V$, we have $\left| e(U, W) - \frac{d|U||W|}{n} \right| \leq \lambda \sqrt{|U||W| \left(1 - \frac{|U|}{n}\right) \left(1 - \frac{|W|}{n}\right)}$.*

Remark 2.2. *An inspection of the proof of Theorem 2.11 in [22] shows that the statement continues to hold for graphs (possibly with loops) that correspond to adjacency matrices \mathbf{A} where each row sum equals d , but the diagonal elements $A_{i,i}$ can be any non-negative integer, provided that the sets U and W are disjoint.*

Let us now turn our attention to non-regular graphs.

Lemma 2.3 ([22]). *Let $G = (V, E)$ be any graph with average degree d and spectral expansion $\lambda < d$. Then for any subsets $U, W \subseteq V$ we have $\left| e(U, W) - \frac{d|U||W|}{n} \right| \leq \lambda \sqrt{|U||W|}$.*

We prove the following generalization of Lemma 2.1 to almost regular graphs.

Lemma 2.4. *Let $G = (V, E)$ be an $(n, \delta, \Delta, \lambda)$ -graph. Then for any two disjoint subsets $U, W \subseteq V$, we have $\left| e(U, W) - \frac{\Delta|U||W|}{n} \right| \leq (\lambda + (\Delta - \delta)) \sqrt{|U||W| \left(1 - \frac{|U|}{n}\right) \left(1 - \frac{|W|}{n}\right)}$.*

Proof. Define a matrix $\mathbf{B} := \mathbf{A} + \mathbf{D}$, where \mathbf{D} is the diagonal matrix defined by $\mathbf{D}_{i,i} = \Delta - \deg(i)$, and \mathbf{A} is the adjacency matrix of G . The eigenvalues of \mathbf{D} are $\varepsilon_1 \geq \varepsilon_2 \geq \dots \geq \varepsilon_n$. The matrix \mathbf{B} is an integer-valued, symmetric matrix with eigenvalues $\Delta = \beta_1 \geq \beta_2 \geq \dots \geq \beta_n$. This allows us to apply Lemma A.2, which yields $\lambda_2 \leq \beta_2 \leq \lambda_2 + \varepsilon_1$ and $\lambda_n \leq \beta_n \leq \lambda_n + \varepsilon_1$ (as all eigenvalues of \mathbf{D} are positive). Therefore with $\beta := \max_{2 \leq i \leq n} |\beta_i|$, we have

$$\begin{aligned} \beta &= \max\{|\beta_2|, |\beta_n|\} \\ &\leq \max\{|\lambda_2| + |\beta_2 - \lambda_2|, |\lambda_n| + |\beta_n - \lambda_n|\} \\ &\leq \max\{|\lambda_2|, |\lambda_n|\} + \varepsilon_1 \leq \lambda + (\Delta - \delta). \end{aligned}$$

Since every row (and column) sum of \mathbf{B} equals Δ , we can apply Lemma 2.1 to the graph induced by the matrix \mathbf{B} to conclude that

$$\begin{aligned} \left| e_B(U, W) - \frac{\Delta|U||W|}{n} \right| &\leq \beta \sqrt{|U||W| \left(1 - \frac{|U|}{n}\right) \left(1 - \frac{|W|}{n}\right)} \\ &\leq (\lambda + (\Delta - \delta)) \sqrt{|U||W| \left(1 - \frac{|U|}{n}\right) \left(1 - \frac{|W|}{n}\right)}. \end{aligned}$$

As U and W are disjoint, we have $e_B(U, W) = e(U, W)$ which proves the claim. \square

Corollary 2.5. *Let $G = (V, E)$ be an $(n, \delta, \Delta, \lambda)$ -graph satisfying the preconditions of Theorem 1.1. Then for any subset $U \subseteq V$ of size $1 \leq |U| \leq n/2$, it holds that*

$$\left| e(U, V \setminus U) - \frac{\Delta|U|(n - |U|)}{n} \right| = o(\Delta) \cdot |U|, \quad (2.1)$$

2.2 Analysis of the Algorithm (Upper Bound)

Lemma 2.6. *Let G be a graph that satisfies the preconditions of Theorem 1.1. Then all the following statements are a.a.s. true.*

- **Phase I** Suppose $1 \leq |I_t| \leq n/\log n$. Then there is $\tau = \log_2 n + o(\log n)$ such that $|I_{t+\tau}| > n/\log n$.
- **Phase II** Suppose $n/\log n \leq |I_t| \leq n - n/\log n$. Then there is $\tau = o(\log n)$ such that $|I_{t+\tau}| > n - n/\log n$.
- **Phase III** Suppose $|I_t| \geq n - n/\log n$. Then there is $\tau = \log n + o(\log n)$ such that $|I_{t+\tau}| = n$.

Proof. **Phase I** $|I_t| \in [1, n/\log n]$ By Corollary 2.5, we know that

$$e(I_t, V \setminus I_t) \geq \frac{\Delta|I_t|(n - |I_t|)}{n} - o(\Delta)|I_t| \geq \left(1 - \frac{1}{\log n} - o(1)\right) \Delta|I_t|. \quad (2.2)$$

Define $\tilde{\lambda} := \frac{\lambda}{d} + \frac{1}{\log n}$. Let $A = \{v \in N(I_t) \setminus I_t : |N(v) \cap I_t| \geq 2d\sqrt{\tilde{\lambda}}\}$. Then by definition of A , $e(A, I_t) \geq |A| \cdot 2d\sqrt{\tilde{\lambda}}$. By Lemma 2.3, we have

$$e(A, I_t) \leq \frac{d|A||I_t|}{n} + \lambda\sqrt{|A||I_t|} < \frac{d|A||I_t|}{n} + d\tilde{\lambda}\sqrt{|A||I_t|}.$$

Thus

$$|A| \cdot 2d\sqrt{\tilde{\lambda}} < \frac{d|A||I_t|}{n} + d\tilde{\lambda}\sqrt{|A||I_t|} \iff 2\sqrt{\tilde{\lambda}} - \frac{|I_t|}{n} < \tilde{\lambda}\sqrt{\frac{|I_t|}{|A|}},$$

which implies that

$$\begin{aligned} |A| &< \tilde{\lambda}^2 \cdot |I_t| \cdot \left(\frac{n}{2n\sqrt{\tilde{\lambda}} - |I_t|}\right)^2 \leq \tilde{\lambda}^2 \cdot |I_t| \cdot \left(\frac{n}{2n\sqrt{\tilde{\lambda}} - n/\log n}\right)^2 \\ &\leq \tilde{\lambda}^2 \cdot |I_t| \cdot \left(\frac{1}{2\sqrt{\tilde{\lambda}} - 1/\log n}\right)^2 < \tilde{\lambda} \cdot |I_t|. \end{aligned}$$

Define $B = (N(I_t) \setminus I_t) \setminus A$. Using Equation (2.2), we can bound $e(I_t, B)$ by

$$\begin{aligned} e(I_t, B) &= e(I_t, V \setminus I_t) - e(I_t, A) \\ &\geq (1 - o(1)) \Delta |I_t| - \frac{d|A||I_t|}{n} - d\tilde{\lambda}\sqrt{|A||I_t|} \\ &\geq (1 - o(1)) \Delta |I_t| - d \left(\frac{\lambda}{d} + \frac{1}{\log n} \right) |I_t| \\ &\geq (1 - o(1)) \Delta |I_t|. \end{aligned}$$

With the above estimate at hand, we compute the expected value of $|I_{t+1} \cap B|$. Note that for any $v \in B$, the probability that it gets informed is at least $1 - \prod_{u \in N(v) \cap I_t} (1 - \frac{1}{\Delta})$. We have

$$\mathbf{E}[|I_{t+1} \cap B|] \geq \sum_{v \in B} \left[1 - \prod_{u \in N(v) \cap I_t} \left(1 - \frac{1}{\Delta} \right) \right] = \sum_{v \in B} 1 - \left(1 - \frac{1}{\Delta} \right)^{|N(v) \cap I_t|}.$$

Using the inequality that $(1 - x)^n \leq 1 - nx + n^2x^2$, we get

$$\begin{aligned} \mathbf{E}[|I_{t+1} \cap B|] &\geq \sum_{v \in B} 1 - \left(1 - \frac{|N(v) \cap I_t|}{\Delta} + \frac{|N(v) \cap I_t|^2}{\Delta^2} \right) \\ &= \sum_{v \in B} \frac{|N(v) \cap I_t|}{\Delta} \left(1 - \frac{|N(v) \cap I_t|}{\Delta} \right) \geq \left(1 - \frac{2d\sqrt{\tilde{\lambda}}}{\Delta} \right) \frac{e(I_t, B)}{\Delta}. \end{aligned}$$

Since the last expression is at least

$$\left(1 - \frac{2d\sqrt{\tilde{\lambda}}}{\Delta} \right) \frac{e(I_t, B)}{\Delta} \geq (1 - o(1))(1 - o(1))|I_t| = (1 - o(1))|I_t|,$$

we obtain $\mathbf{E}[|I_{t+1} \setminus I_t|] \geq \mathbf{E}[|I_{t+1} \cap B|] \geq (1 - o(1))|I_t|$. Since $|I_{t+1} \setminus I_t| \leq |I_t|$, it follows by using Markov's inequality (applied to $|I_t| - |I_{t+1} \setminus I_t|$) that $\Pr[|I_{t+1}| \geq (2 - f(n))|I_t|] \geq 1 - g(n)$, where $f(n)$ and $g(n)$ are both functions that tend to zero. Hence the time to reach $|I_t| \geq n/\log n$ can be upper bounded by the sum of $\log_{2-f(n)} n$ independent, identically distributed geometric random variables with expectation at most $1 + o(1)$ each. Using the Chernoff bound from Theorem A.4 yields for $\tau := \log_2 n + o(\log n)$ that $\Pr[|I_{t+\tau}| > n/\log n] = 1 - o(1)$.

Phase II $|I_t| \in [n/\log n, n - n/\log n]$. We further divide this phase into the two cases $|I_t| \in [n/\log n, n/2]$ and $|I_t| \in [n/2, n - n/\log n]$. We start with the first case $|I_t| \in [n/\log n, n/2]$. By Lemma 2.4, we have

$$\begin{aligned} e(I_t, V \setminus I_t) &\geq \frac{\Delta |I_t| |V \setminus I_t|}{n} - (\lambda + \Delta - \delta) \cdot \frac{|I_t| |V \setminus I_t|}{n} \\ &\geq \frac{1}{2}(\delta - \lambda)|I_t| > \frac{1}{4}\delta|I_t|, \end{aligned}$$

where in the last inequality we used the assumptions that $\lambda = o(\Delta)$ and $\Delta/\delta = 1 + o(1)$. Similar to the analysis of Phase I, we can lower bound the expected number of vertices that become informed in round $t + 1$:

$$\begin{aligned} \mathbf{E}[|I_{t+1} \setminus I_t|] &\geq \sum_{u \in N(I_t) \setminus I_t} \left[1 - \prod_{v \in N(u) \cap I_t} \left(1 - \frac{1}{\Delta} \right) \right] \geq \sum_{u \in N(I_t) \setminus I_t} 1 - e^{-|N(u) \cap I_t|/\Delta} \\ &\geq \sum_{u \in N(I_t) \setminus I_t} \frac{|N(u) \cap I_t|}{2\Delta} = \frac{e(I_t, V \setminus I_t)}{2\Delta} \geq \frac{\delta}{8\Delta}|I_t|, \end{aligned}$$

where the third inequality used the fact that $e^{-x} \leq 1 - x/2$ for any $x \in (0, 1)$.

Since $|I_{t+1}| \leq 2|I_t|$, we obtain as long as $|I_t| \leq n/2$ that there are constants $\alpha, \beta > 0$ so that $\Pr[|I_{t+1}| \geq (1 + \alpha)|I_t|] \geq \beta$. Hence the time to reach $|I_t| \geq n/2$ can be upper bounded by the sum of $\log_{1+\alpha}(\log n)$ independent, identically distributed geometric random variables with expectation at most $1/\beta$ each. Using a Chernoff bound for the sum of geometric random variables (see Theorem A.4) yields that with probability $1 - o(1)$, we reach $|I_t| \geq n/2$ within at most $o(\log n)$ additional rounds.

Consider now the case $|I_t| \in [n/2, n - n/\log n]$. To analyze this case, we examine the shrinking of $U_t := V \setminus I_t$. Again, as $|U_t| \leq n/2$, by Lemma 2.4 we have

$$e(U_t, I_t) > \frac{1}{4}\delta|U_t|.$$

Let us now compute the expected number of uninformed vertices after one additional round:

$$\begin{aligned} \mathbf{E}[|U_{t+1}|] &= \sum_{u \in U_t} \prod_{v \in N(u) \cap I_t} \left(1 - \frac{1}{\deg(v)}\right) \leq \sum_{u \in U_t} e^{-|N(u) \cap I_t|/\Delta} \\ &\leq \sum_{u \in U_t} 1 - \frac{|N(u) \cap I_t|}{2\Delta} \leq |U_t| - \frac{\delta}{8\Delta}|U_t| = \left(1 - \frac{\delta}{8\Delta}\right)|U_t|. \end{aligned}$$

A simple inductive argument yields for any integer $\tau \geq 1$ that,

$$\mathbf{E}[|U_{t+\tau}|] \leq \left(1 - \frac{\delta}{8\Delta}\right)^\tau |U_t|,$$

so for $\tau := \log \log(n) / \log(1/(1 - \frac{\delta}{8\Delta})) + \omega(1)$, where $\omega(1)$ is an arbitrarily slow growing function, we have $\mathbf{E}[|U_{t+\tau}|] = o(n/\log n)$. Hence by Markov's inequality, $\Pr[|U_{t+\tau}| \geq n/\log n] = o(1)$.

Phase III $|I_t| \in [n - n/\log n, n]$ Again, we analyze the shrinking of the set U_t . Recall from equation 2.2 that $e(I_t, U_t) \geq (1 - f(n)) \cdot \Delta \cdot |U_t|$, where $f(n)$ is any function with $f(n) = o(1)$. Let $A \subseteq U_t$ be the set of vertices $v \in U_t$ for which $|N(v) \cap I_t| \leq (1 - \sqrt{f(n)}/2) \cdot \Delta$. We assume for a contradiction that $|A| > 2\sqrt{f(n)} \cdot |U_t|$ and conclude

$$\begin{aligned} e(I_t, U_t) &= \sum_{v \in A} |N(v) \cap I_t| + \sum_{v \in U_t \setminus A} |N(v) \cap I_t| \\ &\leq |A| \cdot (1 - \sqrt{f(n)}/2) \cdot \Delta + |U_t \setminus A| \cdot \Delta \\ &= |U_t| \cdot \Delta - |A| \cdot \sqrt{f(n)} \cdot \Delta/2 \\ &< (1 - f(n)) \cdot \Delta \cdot |U_t|, \end{aligned}$$

which yields the desired contradiction. Now define $B := U_t \setminus A$ so that for each $u \in B$, $|N(u) \cap I_t| > (1 - \sqrt{f(n)}/2) \cdot \Delta$ and $|B| \geq (1 - 2\sqrt{f(n)}) \cdot |U_t|$. Using linearity of expectations,

$$\begin{aligned} \mathbf{E}[|U_{t+1}|] &\leq \sum_{v \in B} \Pr[v \notin I_{t+1}] + \sum_{v \in A} \Pr[v \notin I_{t+1}] \\ &\leq \sum_{v \in B} \left(1 - \frac{1}{\Delta}\right)^{|N(v) \cap I_t|} + \sum_{v \in A} 1 \end{aligned}$$

Using the inequality $(1 - 1/\Delta) \leq e^{-1/\Delta}$, we get

$$\begin{aligned} \mathbf{E}[|U_{t+1}|] &\leq \sum_{v \in B} e^{-|N(v) \cap I_t|/\Delta} + |A| \\ &< (e^{-1} \cdot e^{\sqrt{f(n)}/2}) \cdot |B| + 2\sqrt{f(n)} \cdot |U_t| \\ &\leq \left(e^{-1} \cdot \left(1 + e^{\sqrt{f(n)}/2} \right) \right) \cdot |B| + 2\sqrt{f(n)} \cdot |U_t| \\ &\leq (e^{-1} + 3\sqrt{f(n)}) \cdot |U_t|, \end{aligned}$$

where the third inequality used the fact that $e^x \leq 1 + e \cdot x$ for any $x \in [0, 1]$. By induction, it follows that for any step $\tau > t$, $\mathbf{E}[|U_\tau|] \leq (e^{-1} + 3\sqrt{f(n)})^{\tau-t} \cdot |U_t|$. We choose

$$\tau := t + \log_{e^{-1} + 3\sqrt{f(n)}}(n) = t + \log n + o(\log n)$$

and obtain that

$$\mathbf{E}[|U_{t+\tau}|] \leq \frac{1}{n} \cdot |U_t|.$$

Since by assumption $|U_t| \leq n/\log n$, this implies $\mathbf{E}[|U_{t+\tau}|] \leq 1/\log n$, so that $\mathbf{Pr}[|U_{t+\tau}| \geq 1] \leq \mathbf{E}[|U_{t+\tau}|] \leq 1/\log n$. \square

2.3 Analysis of the Algorithm (Lower Bound)

We first note that if the graph G is regular, then the lower bound follows directly from [10, Theorem 1 & Lemma 1]. If the graph G is not regular, then it must hold that $\delta = \omega(1)$. For the proof of the lower bound, observe that after $t := \log_2 n - 1$ steps, we have $|I_t| \leq n/2$. For a lower bound on the running time of the algorithm, we may assume that each vertex $u \in U_t$ may get the rumor from any of its neighbors at any time. So we can forget about who actually knows the rumor but consider the model in which at each step each vertex in V picks a neighbor. We want a lower bound on the probability that some vertex in U_t never gets selected in $\tau := (\log n - \omega(1)) \cdot \frac{\delta - 3/4}{\Delta} = \log n - o(\log n)$ steps, where $\omega(1)$ is any slow growing function. For each $u \in U_t$ let E_u denote the event that u is never selected within those steps. We compute

$$\mathbf{Pr}[E_u] = \prod_{v \in N(u)} \left(1 - \frac{1}{|N(v)|} \right)^\tau \geq \left(1 - \frac{1}{\delta} \right)^{\Delta\tau} \geq e^{\Delta\tau/(\delta-3/4)} = e^{-\log n + \omega(1)},$$

where we used the fact that $(1 - 1/x)^{x-3/4} \geq e^{-1}$ for any $x \geq 2$. Summing over all vertices in U_t we obtain

$$\sum_{u \in U_t} \mathbf{Pr}[E_u] \geq n/2 \cdot e^{-\log n + \omega(1)} \rightarrow \infty.$$

By construction, the events $\{\overline{E_u} : u \in U_t\}$ are negatively correlated in the sense that, for any set of different vertices $u, u_1, \dots, u_k \in U_t$, we have that $\mathbf{Pr}[\overline{E_u} \mid \overline{E_{u_1}} \wedge \dots \wedge \overline{E_{u_k}}] \leq \mathbf{Pr}[\overline{E_u}]$. Therefore,

$$\mathbf{Pr} \left[\bigwedge_{u \in U_t} \overline{E_u} \right] \leq \prod_{u \in U_t} \mathbf{Pr}[\overline{E_u}] \leq e^{-\sum_{u \in U_t} \mathbf{Pr}[E_u]} = o(1).$$

To conclude, we have shown that $\mathbf{Pr}[T(G) \geq \log_2 n + \log n - o(\log n)] \geq 1 - o(1)$. This lower bound together with the upper bound obtained in Lemma 2.6 immediately yield Theorem 1.1.

3 Rumor spreading on $G(n, p)$

In this section we analyze the push protocol on the classical Erdős-Rényi random graph, and prove Theorem 1.2. Henceforth, let $G(n, p)$ denote the random graph on vertex set $V = \{1, \dots, n\}$ where each edge is selected independently with probability p . Throughout this section, we will write $f(n) \sim g(n)$ if $\lim_{n \rightarrow \infty} f(n)/g(n) = 1$.

3.1 Properties of $G(n, p)$

In this subsection we collect some basic facts about $G(n, p)$, which will be useful in the forthcoming proofs. We begin with computing tight bounds for the degree sequence of $G(n, p)$, where $p = \Theta(\frac{\log n}{n})$. The next two lemmas provide us with the required information.

Lemma 3.1. *Let $x_1 > x_0 > 0$ and $c > 1$ be constants, and let $p \sim c \log n/n$. Fix $V' \subseteq V$ with $|V'| = n - o(n)$ (where $V' = V$ is allowed). Let N_i be the number of vertices in V' with degree i in $G(n, p)$. Then, uniformly for all $i \in [x_0 \log n, x_1 \log n]$ we have that $\mathbf{E}[N_i] = n^{g(i/\log n) \pm o(1)}$, where $g(x) = x(1 + \log(c/x)) - c + 1$. Moreover, $\mathbf{Var}[N_i] \leq o(\mathbf{E}[N_i]^2) + \mathbf{E}[N_i]$.*

Proof. The degree $\deg(v)$ of any given vertex $v \in V'$ is distributed as $\text{Bin}(n-1, p)$. Then, Stirling's formula implies that

$$\mathbf{E}[N_i] = |V'| \Pr[\deg(v) = i] = (n - o(n)) \binom{n-1}{i} p^i (1-p)^{n-1-i} = n^{1-o(1)} \left(\frac{en}{i}\right)^i p^i (1-p)^{n-1-i}.$$

Hence, writing $x = i/\log n$, we get

$$\mathbf{E}[N_i] = n^{1-o(1)} \left(\frac{ec \pm o(1)}{x}\right)^i e^{-pn \pm o(1)} = n^{\log((ec/x)^x) - c + 1 \pm o(1)} = n^{g(x) \pm o(1)}.$$

Next, we estimate the probability that two different vertices v and w in V' have both degree i . To compute that, we take into account whether or not v and w share a common edge, which occurs with probability p , and then rearrange the expression obtained in terms of $\Pr[\deg(v) = i]$.

$$\begin{aligned} & \Pr[\deg(v) = i \cap \deg(w) = i] \\ &= p \binom{n-2}{i-1}^2 p^{2i-2} (1-p)^{2n-2-2i} + (1-p) \binom{n-2}{i}^2 p^{2i} (1-p)^{2n-4-2i} \\ &= \left(\frac{i^2}{p(n-1)^2} + \frac{(n-1-i)^2}{(1-p)(n-1)^2} \right) (\Pr[\deg(v) = i])^2 \sim (\Pr[\deg(v) = i])^2. \end{aligned}$$

Finally, we have

$$\begin{aligned} \mathbf{Var}[N_i] &= |V'| \Pr[\deg(v) = i] + |V'|(|V'| - 1) \Pr[\deg(v) = i \cap \deg(w) = i] - \mathbf{E}[N_i]^2 \\ &= \mathbf{E}[N_i] + (1 \pm o(1)) \mathbf{E}[N_i]^2 - \mathbf{E}[N_i]^2 \\ &\leq o(\mathbf{E}[N_i]^2) + \mathbf{E}[N_i]. \end{aligned} \quad \square$$

Lemma 3.2. *Let $c > 1$ be any constant and $p = c \log n/n$. Let δ and Δ be the minimum and maximum degrees of $G(n, p)$. Then there exist constants c_0, c_1 with $0 < c_0 < c - 1 < c < c_1$ such that a.a.s. $(1 - o(1))c_0 \log n \leq \delta \leq \Delta \leq (1 + o(1))c_1 \log n$.*

Proof. Define

$$c_0 = -\frac{c-1}{W_0(-e^{-1}(c-1)/c)} \quad \text{and} \quad c_1 = -\frac{c-1}{W_1(-e^{-1}(c-1)/c)},$$

where W_0 and W_1 are respectively the lower and upper branch of the Lambert W function on $[-e^{-1}, 0]$ (recall that each branch of the Lambert W function satisfies $W(y)e^{W(y)} = y$). Let $g(x)$ be defined as in the statement of Lemma 3.1. By direct substitution and easy computations, we can check that $g(c_0) = g(c_1) = 0$, $g(c-1) > 0$ and $g(c) > 0$. Moreover, by looking at the derivative $g'(x) = \log(c/x)$ we see that $g(x)$ is increasing in $(0, c)$ and decreasing in (c, ∞) . Thus we can conclude that the only positive solutions of $g(x) = 0$ are precisely $x = c_0$ and $x = c_1$, and moreover $0 < c_0 < c-1 < c < c_1$.

Given any constant $\varepsilon > 0$, for each vertex v we have

$$\begin{aligned} \Pr[\deg(v) \leq (c_0 - \varepsilon) \log n] &= \sum_{j=0}^{\lfloor (c_0 - \varepsilon) \log n \rfloor} \binom{n-1}{j} p^j (1-p)^{n-1-j} \\ &= O(\Pr[\deg(v) = \lfloor (c_0 - \varepsilon) \log n \rfloor]), \end{aligned}$$

since the ratio between any two consecutive terms in the above sum is at most $c_0/c < 1$. Similarly,

$$\begin{aligned} \Pr[\deg(v) \geq (c_1 + \varepsilon) \log n] &= \sum_{j=\lceil (c_1 + \varepsilon) \log n \rceil}^{n-1} \binom{n-1}{j} p^j (1-p)^{n-1-j} \\ &= O(\Pr[\deg(v) = \lceil (c_1 + \varepsilon) \log n \rceil]), \end{aligned}$$

since the corresponding ratio is at least $c_1/c > 1$. In view of these facts and by Lemma 3.1, the expected number of vertices of degree at most $(c_0 - \varepsilon) \log n$ is

$$O(n \cdot \Pr[\deg(v) = \lfloor (c_0 - \varepsilon) \log n \rfloor]) = n^{g(c_0 - \varepsilon) \pm o(1)} = o(1),$$

since $g(c_0 - \varepsilon) < 0$. A totally analogous argument shows that the expected number of vertices of degree at least $(c_1 + \varepsilon) \log n$ is $o(1)$. Since the choice of ε was arbitrary, we conclude that the required bounds on δ and Δ hold a.a.s. \square

We will also use the following property of $G(n, p)$, which essentially says that the neighborhood of all vertices contains many vertices with actual degree close to the expected degree.

Lemma 3.3. *Let $c > 1$ be any constant and $p = c \log n/n$. For each vertex u in $G(n, p)$, define*

$$\tilde{N}(u) = \left\{ v \in N(u) : \left| |N(v)| - pn \right| \leq \log^{3/4} n \right\}. \quad (3.1)$$

Then, a.a.s. for every vertex u we have $|N(u) \setminus \tilde{N}(u)| \leq \log^{3/4} n$.

Proof. In view of Lemma 3.2, there is a constant $x_1 > c$ such that a.a.s. $\Delta \leq x_1 \log n$. We call this event E , and restrict our analysis to that case. Let us fix a vertex $u \in V$. Put $N = N(u)$ and $N' = N(u) \setminus \tilde{N}(u)$. To prove the lemma, it suffices to show that $\Pr[(|N'| > \log^{3/4} n) \cap E] = o(1/n)$, and then apply a union bound over all the vertices $u \in V$.

First, we expose the neighbors of u , and condition upon the particular N obtained. We use a subscript star to denote probabilities in the conditional space. We assume that $|N| \leq x_1 \log n$ (as implied by event E). Then we expose all internal edges in N , and let X be the number of them. Since there are at most $|N|^4 \leq (x_1 \log n)^4$ possible pairs of internal edges, we have that

$$\Pr_*[X \geq 2] \leq (x_1 \log n)^4 p^2 = n^{-2+o(1)}. \quad (3.2)$$

Next we expose the edge set $E(N, V_0)$, where $V_0 = V \setminus (N \cup \{u\})$. For each $v \in N$, consider the number $X_v = |N(v) \cap V_0|$ of neighbors of v which lie in V_0 . Note that each X_v is distributed as $\text{Bin}(|V_0|, p)$, so a version of the Chernoff bounds (see e.g. Alon and Spencer [3, Theorems A.1.11 and A.1.13]) yields

$$\Pr_* \left[|X_v - pn| > \log^{3/4} n - 2 \right] = e^{-\Omega(\log^{1/2} n)}.$$

Let B be the event that there are more than $\log^{3/4} n$ vertices $v \in N$ for which $|X_v - pn| > \log^{3/4} n - 2$ holds. Note that by definition, the random variables $\{X_v : v \in N(u)\}$ are mutually independent. Hence,

$$\Pr_*[B] \leq |N| \binom{|N|}{\lceil \log^{3/4} n \rceil} \exp\left(-\Omega(\log^{1/2} n) \lceil \log^{3/4} n \rceil\right) = e^{-\Omega(\log^{5/4} n)}. \quad (3.3)$$

Since the bounds in (3.2) and (3.3) hold uniformly for all N satisfying $|N| \leq x_1 \log n$, we deduce that

$$\Pr[(X \geq 2) \cup B \cap (|N| \leq x_1 \log n)] \leq n^{-2+o(1)} = o(1/n).$$

Note that if $X < 2$ then $X_v \leq \deg(v) \leq X_v + 1$ for each $v \in N$. We conclude the proof by observing that the event $(|N'| > \log^{3/4} n) \cap E$ implies $((X \geq 2) \cup B) \cap (|N| \leq x_1 \log n)$ by construction, and thus we get the desired bound $\Pr[(|N'| > \log^{3/4} n) \cap E] = o(1/n)$. \square

The lemma below bounds the number of edges contained in small subsets of the vertices. The first statement follows immediately from Lemma 5.3 in [8], while the second statement is from [7, Property 3].

Lemma 3.4. *Let $c > 1$ be any constant and $p = c \log n/n$. A.a.s. for every subset $S \subseteq V$ of $G(n, p)$ of size $|S| = O(n/\log n)$ we have that $e(S) = o(|S| \log n)$. Moreover, for any set S with $1 \leq |S| \leq n/2$, $e(S, V \setminus S) = \Omega(|S|np)$.*

Finally, we include a lemma that addresses the typical structure of the neighborhoods of small sets.

Lemma 3.5. *Let $c > 1$ be any constant and $p = c \log n/n$. There is a function $f(n) = o(1)$ such that the following is a.a.s. true. For every set S of vertices of $G(n, p)$ such that $|S| \leq \frac{n}{\log n}$ there are at most $f(n)|S|$ vertices outside of S that have more than $f(n) \log n$ neighbors in S .*

Proof. We will show the claim for $f(n) = (\log \log n)^{-1/3}$. Let H be a fixed subset of the vertex set of $G(n, p)$, and let $v \in V \setminus H$. Then, if $|H| \leq \frac{n}{\log n}$,

$$\begin{aligned} \Pr[|N(v) \cap H| \geq f(n) \log n] &\leq \binom{|H|}{\lceil f(n) \log n \rceil} p^{\lceil f(n) \log n \rceil} \\ &\leq \left(\frac{en}{f(n) \log^2 n} \cdot \frac{c \log n}{n} \right)^{\lceil f(n) \log n \rceil} = n^{-\Omega((\log \log n)^{2/3})}. \end{aligned}$$

Hence, the probability that there are at least $f(n)s$ vertices v in $V \setminus H$ with $|N(v) \cap H| \geq f(n) \log n$ (i.e. set H contradicts the claim in the statement) is at most

$$n^{f(n)s} \cdot n^{-\Omega((\log \log n)^{2/3}) \cdot f(n)s} = n^{-\Omega((\log \log n)^{2/3}) \cdot f(n)s}.$$

By linearity of expectation, the expected number of subsets of $G(n, p)$ with $1 \leq s \leq \frac{n}{\log n}$ vertices having the property above is at most

$$n^s n^{-\Omega((\log \log n)^{2/3}) \cdot f(n)s} = o(1/n),$$

and the proof is completed by summing for all s and applying Markov's inequality. \square

3.2 Analysis of the Algorithm

We first analyze the evolution of I_t as long as $1 \leq |I_t| \leq n - n/\log^2 n$. The following lemma proves that as long as $|I_t|$ is not too large, $|I_t|$ almost doubles in each round.

Lemma 3.6. *Let $c > 1$ be any constant and set $p = c \log n/n$. Then there is some $t = \log_2 n + o(\log n)$ such that we have a.a.s. that $|I_t| \geq n - n/\log^2 n$.*

Proof. Before we proceed with the actual proof, let us recall some basic properties of $G(n, p)$. First, by applying Lemma 3.2 we infer that there are constants $0 < c_0 < c_1$ such that a.a.s. the minimum degree δ of $G_{n,p}$ satisfies $\delta = (1 - o(1))c_0 \log n$ and the maximum degree satisfies $\Delta = (1 + o(1))c_1 \log n$. Moreover, let $f(n) = o(1)$ be the function guaranteed to exist by Lemma 3.5. Set

$$A := \{v \in N(I_t) \setminus I_t : |N(v) \cap I_t| \geq f(n) \log n\}.$$

Then, a.a.s. we may assume that $|A| \leq f(n)|I_t|$.

Let us first consider the case where $1 \leq |I_t| \leq n/\log n$. By applying Lemma 3.4 we infer that a.a.s. the set I_t spans $o(|I_t| \log n)$ edges, i.e., $e(I_t) = o(|I_t| \log n)$. In the remainder we will assume that $G(n, p)$ has all these properties without further reference.

Let $B := (N(I_t) \setminus I_t) \setminus A$. Define for any vertex $u \in I_t$ an indicator random variable X_u which is one if u sends the rumor to a vertex $v \in B$ and no other vertex sends a rumor to v . We have

$$\begin{aligned} \Pr[X_u = 1] &= \sum_{v \in B} \Pr[u \text{ is the only vertex that informs } v] \\ &\geq \sum_{v \in B \cap N(u)} \frac{1}{\deg(u)} \cdot \left(1 - \frac{1}{\delta}\right)^{f(n) \log n} = (1 - o(1)) \frac{|N(u) \cap B|}{\deg(u)}. \end{aligned}$$

Let $X := \sum_{u \in I_t} X_u$. By linearity of expectation,

$$\begin{aligned} \mathbf{E}[X] &= \sum_{u \in I_t} \mathbf{E}[X_u] \geq (1 - o(1)) \cdot \sum_{u \in I_t} \frac{\deg(u) - |N(u) \cap I_t| - |N(u) \cap A|}{\deg(u)} \\ &\geq (1 - o(1)) \left(|I_t| - \frac{1}{\delta} \sum_{u \in I_t} |N(u) \cap I_t| - \frac{1}{\delta} \sum_{u \in I_t} |N(u) \cap A| \right). \end{aligned}$$

Note that $\sum_{u \in I_t} |N(u) \cap I_t| = 2e(I_t) = o(|I_t| \log n)$. Moreover, a simple double counting argument implies $\sum_{u \in I_t} |N(u) \cap A| \leq |A| \Delta$. As $|A| \leq f(n)|I_t|$ and $\delta, \Delta = \Theta(\log n)$ we infer that $\mathbf{E}[X] \geq (1 - o(1))|I_t|$, and we infer that $\mathbf{E}[|I_{t+1}| \mid |I_t| \leq n/\log n] \geq (2 - o(1)) \cdot |I_t|$. As in Lemma 2.6, we can prove that for $t = \log_2 n + o(\log n)$ a.a.s, $|I_t| \geq n/\log n$.

Consider now the case where $n/\log n \leq |I_t| \leq n - n/\log^2 n$. We can do exactly the same analysis as in Lemma 2.6 and use the fact that the ratio Δ/δ is a constant. Here, the second statement of Lemma 3.4 provides a sufficiently large lower bound on the expansion. Therefore, after $\tau := O(\log \log n)$ additional steps a.a.s, $|I_{t+\tau}| \geq n - n/\log^2 n$. \square

The next proposition analyzes the last stages of the rumour spreading algorithm when $|I_t|$ is large. Its conclusion, combined with Lemma 3.6, provides us with an upper bound on the broadcast time for $G(n, p)$.

Proposition 3.7 (Upper bound). *Let $\varepsilon > 0$, $c > 1$ be any constants, and $p = c \log n/n$. Let t be such that $|I_t| \geq n - n/\log^2 n$. Then, all the remaining vertices will a.a.s. get informed within additional $\lceil (\gamma + \varepsilon) \log n \rceil$ rounds, where $\gamma = \gamma(c) = c \log(\frac{c}{c-1})$.*

Proof. Recall from Lemma 3.2 that there exist constants $x_1 > x_0 > 0$ such that a.a.s. $x_0 \log n < \delta \leq \Delta < x_1 \log n$. Call this event E_1 . Let V_i denote the set of vertices of degree i . By Markov's inequality, the event $|V_i| \geq \log^2 n \mathbf{E}[|V_i|]$ has probability at most $1/\log^2 n$. Therefore, we can apply a union bound and also Lemma 3.1 to deduce that a.a.s. all degrees $i \in [x_0 \log n, x_1 \log n]$ satisfy $|V_i| < \log^2 n \mathbf{E}[|V_i|] \leq n^{g(i/\log n)+o(1)}$. Call this event E_2 . Let E_3 be the event that $|N(u) \setminus \tilde{N}(u)| \leq \log^{3/4} n$ for every vertex u , where $\tilde{N}(u)$ is defined as in (3.1). By Lemma 3.3, E_3 holds a.a.s.

Let us write $I = I_t$ and $U = U_t$ for simplicity, and note that the internal edges of U were not exposed during the rumor spreading, so each one still occurs independently with probability p . For any vertex u in U , using the inequality $\binom{m}{k} \leq (em/k)^k$, we obtain

$$\Pr \left[|N(u) \cap U| \geq \frac{\log n}{\log \log n} \right] \leq \left(\binom{|U|}{\lceil \frac{\log n}{\log \log n} \rceil} \right) p^{\lceil \frac{\log n}{\log \log n} \rceil} \leq \left(\frac{ec \log \log n}{\log^2 n} \right)^{\frac{\log n}{\log \log n}} \leq n^{-2+o(1)}.$$

Thus taking the union bound over all vertices in U , we get that a.a.s. each vertex in U has at most $\log n / \log \log n$ neighbors in U . Call this event E_4 .

Henceforth, assume that E_1, E_2, E_3 and E_4 hold together. We can partition U into the sets $U_i = U \cap V_i$, for $i \in [x_0 \log n, x_1 \log n]$, which must satisfy $|U_i| \leq |V_i| \leq n^{g(i/\log n)+o(1)}$ (since E_1 and E_2 hold). Pick an i in that range and a vertex $u \in U_i$. Set $x = i/\log n$. Define $N_1 = \tilde{N}(u) \cap I$ and $N_2 = (N(u) \setminus \tilde{N}(u)) \cap I$, and observe that events E_3 and E_4 imply $|N_1| \sim i = x \log n$ and $|N_2| = o(\log n)$. We can upper-bound the probability that u does not receive the rumor in $\lceil (\gamma + \varepsilon) \log n \rceil$ steps from any vertex in I by

$$\prod_{v \in N(u) \cap I} \left(1 - \frac{1}{|N(v)|} \right)^{\lceil (\gamma + \varepsilon) \log n \rceil} \leq \prod_{v \in N_1} e^{-\frac{(\gamma + \varepsilon) \log n}{|N(v)|}} \prod_{v \in N_2} e^{-\frac{(\gamma + \varepsilon) \log n}{|N(v)|}} \leq n^{-(\gamma + \varepsilon)x/c + o(1)},$$

and therefore the expected number of vertices in U_i not being informed in that time is at most $n^{g(x) - (\gamma + \varepsilon)x/c + o(1)}$. Standard analysis shows that the function $cg(x)/x$ maximizes at $x = c - 1$ and takes the value γ . Therefore, $g(x) \leq \gamma x/c$, and the expectation above is at most $n^{-\varepsilon x/c + o(1)} \leq n^{-\varepsilon x_0/c + o(1)}$. Hence taking a union bound over all degrees $i \in [x_0 \log n, x_1 \log n]$ we get that a.a.s. all vertices get informed. \square

Finally, we bound from below the time it takes to inform the last uninformed vertices of the graph, when $|I_t|$ is not too large. This will be used to obtain a lower bound on the broadcast time for $G(n, p)$.

Proposition 3.8 (Lower bound). *Let $c > 1$ and $p = c \log n/n$. Let us assume that t is such that $|I_t| \leq n/\log^2 n$. Then, given any $\varepsilon > 0$, a.a.s. after $\lceil (\gamma - \varepsilon) \log n \rceil$ steps there are still some uninformed vertices in $G(n, p)$, where $\gamma = \gamma(c) = c \log(\frac{c}{c-1})$.*

Proof. Let us write $I = I_t$ for brevity. In view of Lemma 3.2, we assume that all vertices in I have degree at most $x_1 \log n$, for some constant $x_1 > c$, and in particular $|N(I)| \leq x_1 n / \log n$. We define $U = V \setminus I$ and $U' = U \setminus N(I)$, which must satisfy $|U| \geq n - n/\log^2 n$ and $|U'| \geq n - x_1 n / \log n$. Moreover, observe that the internal edges of U have not been exposed yet by the rumor spreading algorithm.

Let $i = \lceil (c - 1) \log n \rceil$, and let U_i be the set of vertices in U' with exactly i neighbors. We wish to estimate the size of U_i . Since there are no crossing edges between U' and I , henceforth we confine our attention only to $G(n, p)$ restricted to U , which can be regarded as $G(|U|, p)$. We apply Lemma 3.1 to $G(|U|, p)$ and to the subset U' , and infer that $\mathbf{E}[|U_i|] \geq n^{g(c-1)-o(1)} \rightarrow \infty$

and $\sqrt{\mathbf{Var}[|U_i|]} = o(\mathbf{E}[|U_i|])$, so as a consequence of Chebyshev's inequality we deduce that a.a.s. $|U_i| \geq n^{g(c-1)-o(1)}$.

Summarizing, we obtained a.a.s. a set U_i of at least $n^{g(c-1)-o(1)}$ uninformed vertices with degree i . In view of Lemma 3.3, for every vertex $u \in U_i$ we have $|N(u) \setminus \tilde{N}(u)| \leq \log^{3/4} n$, where $\tilde{N}(u)$ is defined as in (3.1). We may assume for our lower bound on the running time of the algorithm that each vertex $u \in U_i$ may get the rumor from any of its neighbors at any time. So we can forget about who actually knows the rumor at a given time but consider the model in which at each step each vertex in U picks a neighbor. We want a lower bound on the probability that some vertex in U_i never gets selected in $\lceil(\gamma - \varepsilon) \log n\rceil$ steps. For each $u \in U_i$ let E_u denote the event that u is never selected within those steps. We compute

$$\begin{aligned} \Pr[E_u] &= \prod_{v \in N(u)} \left(1 - \frac{1}{|N(v)|}\right)^{\lceil(\gamma - \varepsilon) \log n\rceil} \geq \prod_{v \in N(u)} e^{-\frac{(\gamma - \varepsilon) \log n}{|N(v)|} - o(1)} \\ &= \prod_{v \in \tilde{N}(u)} e^{-\frac{(\gamma - \varepsilon) \log n}{|N(v)|} - o(1)} \prod_{v \in N(u) \setminus \tilde{N}(u)} e^{-\frac{(\gamma - \varepsilon) \log n}{|N(v)|} - o(1)} \geq n^{-(\gamma - \varepsilon)(c-1)/c - o(1)}, \end{aligned}$$

and summing over all vertices in U_i , $\sum_{u \in U_i} \Pr[E_u] \geq n^{g(c-1) - (\gamma - \varepsilon)(c-1)/c - o(1)} = n^{\varepsilon(c-1)/c - o(1)} \rightarrow \infty$. As in Section 2.3, we arrive at

$$\Pr \left[\bigwedge_{u \in U_i} \overline{E_u} \right] \leq \prod_{u \in U_i} \Pr[\overline{E_u}] \leq e^{-\sum_{u \in U_i} \Pr[E_u]} = o(1),$$

and in particular we will a.a.s. have some uninformed vertices after $\lceil(\gamma - \varepsilon) \log n\rceil$ steps. \square

At this stage, we have all the ingredients we need to prove Theorem 1.2. The upper bound on $T(G(n, p))$ follows immediately from Lemma 3.6 and Proposition 3.7. For the lower bound, we simply observe that after $t = \lfloor \log_2(n / \log^2 n) \rfloor = \log_2 n - o(\log n)$ steps we still have $|I_t| \leq n / \log^2 n$, and combine this fact with Proposition 3.8.

4 Lower Bounding Rumor Spreading Time by Random Walks

In this section, we develop a new technique for lower bounding the rumor spreading time in terms of random walk matrices. We consider (lazy) random walks on G based on the transition matrix $\mathbf{P} := \frac{1}{2} \cdot (\mathbf{I} + \mathbf{D}^{-1} \mathbf{A})$, where \mathbf{D} is the $n \times n$ -diagonal matrix with $\deg(u)$ at entry (u, u) . Note that in each step, the random walk stays at the current vertex with probability $1/2$ and otherwise moves to a neighbor chosen uniformly at random.

We first derive a general lemma that relates the probability of informing a vertex to the corresponding entry of the transition matrix \mathbf{P} . Then we apply this lemma to the hypercube and show that after $\log_2 n$ rounds, less than $5n^{0.9} \log n$ nodes are informed. Once we have established this, the lower bound of $\log_2 n + \log n + \Omega(\log n)$ is almost immediate.

Lemma 4.1. *For any step t and any pair of vertices u, v , we have $\Pr[v \in I_t \mid I_0 = \{u\}] \leq 2^t \cdot \mathbf{P}_{u,v}^t$.*

Proof. Consider the matrix $\mathbf{M} := \mathbf{I} + \mathbf{D}^{-1} \mathbf{A}$. Clearly, $\mathbf{M} = 2 \cdot \mathbf{P}$. Fix a vertex $u \in V$ which originates the rumor. We now prove the statement by induction on t . To this end, let p^0 be the unit-vector with $p_u^0 = 1$ and $p_v^0 = 0$ for any $v \neq u$. Define $p^t := p^{t-1} \cdot \mathbf{M}$. Clearly, $p^t = p^0 \mathbf{M}^t$,

and for any vertex v , $p_v^t = \sum_{w \in V} p_w^0 \cdot \mathbf{M}_{w,v}^t = \mathbf{M}_{u,v}^t$. Hence it suffices to prove that for any step t , $\Pr[v \in I_t \mid I_0 = \{u\}] \leq p_v^t$. This holds for $t = 0$ by definition. For the induction step,

$$\begin{aligned}
& \Pr[v \in I_{t+1} \mid I_0 = \{u\}] \\
& \leq \sum_{w \in N(v)} \Pr[v \text{ receives rumor from } w \text{ in round } t+1 \mid I_0 = \{u\}] + \Pr[v \in I_t \mid I_0 = \{u\}] \\
& = \sum_{w \in N(v)} \Pr[w \in I_t \mid I_0 = \{u\}] \cdot \frac{1}{\deg(w)} + p_v^t \\
& \leq \sum_{w \in N(v)} p_w^t \cdot \frac{1}{\deg(w)} + p_v^t = [p^t \cdot \mathbf{M}]_v = p_v^{t+1},
\end{aligned}$$

where the last inequality holds due to the induction hypothesis. \square

We now use the above lemma to prove Theorem 1.3.

Proof of Theorem 1.3. Assume that the rumor starts at the vertex $0 = 0^{\log_2 n}$. Our aim is to use Lemma 4.1, so we have to analyze $\mathbf{P}_{u,v}^t$. Recall that a random walk according to the matrix \mathbf{P} stays at the current vertex with probability $1/2$ or moves to a randomly chosen neighbor. This is equivalent to saying that in each step the random walk chooses a coordinate $\{1, \dots, d\}$ uniformly at random and flips the bit with probability $1/2$. Our aim is to prove that for $t := d = \log_2 n$, the random walk is unlikely to reach a vertex which has approximately $(1/2)d$ or more ones in its binary representation. To this end, we let C_t denote the number of coordinates that are chosen within the first t steps of the random walk. Moreover, F_t denotes the number of coordinates that are set to 1 at step t . Then, F_t has distribution $\text{Bin}(C_t, 1/2)$. The idea is now to prove that C_t can be approximated by $\text{Bin}(d, 1 - 1/e)$ which implies that $F_t \approx \text{Bin}(d, p)$, where $p := (1 - 1/e)/2$.

To make this more formal, let us first consider C_t . C_t can be seen as the number of non-empty bins when throwing $t = d$ balls into d bins, where each bin is chosen independently and uniformly at random. Now consider a setting where the number of balls in each bin is an independent Poisson random variable with expected value 1. Then if Y is the number of non-empty bins in this setting, Y has distribution $\text{Bin}(d, 1 - 1/e)$. From the theory of balls into bins we know that for any $0 \leq k \leq d$,

$$\Pr[C_t \geq k] \leq 4 \cdot \Pr[Y \geq k] = 4 \cdot \Pr[\text{Bin}(d, 1 - 1/e) \geq k],$$

since $\{C_t \geq k\}$ is a monotone event (see [1, Corollary 13]). Since $\Pr[\text{Bin}(i, 1/2) \geq r]$ is increasing in i , we conclude that

$$\begin{aligned}
\Pr[F_t \geq r] &= \sum_{i=0}^d \Pr[C_t = i] \cdot \Pr[\text{Bin}(i, 1/2) \geq r] \\
&\leq 4 \cdot \sum_{i=0}^d \Pr[\text{Bin}(d, 1 - 1/e) = i] \cdot \Pr[\text{Bin}(i, 1/2) \geq r] \\
&= 4 \cdot \Pr[\text{Bin}(d, p) \geq r],
\end{aligned}$$

so that $\Pr[F_t \geq r] \leq 4 \cdot \Pr[\text{Bin}(d, p) \geq r]$. By the Chernoff bound,

$$\Pr[F_t \geq (1 + \delta)pd] \leq 4 \cdot \Pr[\text{Bin}(d, p) \geq (1 + \delta)pd] \leq 4 \cdot \left(\frac{e^\delta}{(1 + \delta)^{1 + \delta}} \right)^{pd}$$

Denoting by $|v|_1$ the number of ones in the binary representation of a vertex $v \in \{0, 1\}^d$, we can now upper bound the expected number of informed nodes after $t = \log_2 n$ rounds with the help of Lemma 4.1 as follows:

$$\begin{aligned} \mathbf{E}[|I_t|] &\leq 1 + \sum_{v \in V \setminus \{0\}} \min\{1, 2^t \cdot \mathbf{P}_{0,v}^t\} \\ &\leq 1 + \sum_{v \in V \setminus \{0\}, |v|_1 \leq (1+\delta)pd} 1 + \sum_{v \in V \setminus \{0\}, |v|_1 > (1+\delta)pd} 2^t \cdot \mathbf{P}_{0,v}^t \\ &\leq 1 + \sum_{k=1}^{(1+\delta)pd} \binom{d}{k} + 2^t \cdot 4 \cdot \left(\frac{e^\delta}{(1+\delta)^{1+\delta}} \right)^{dp} \leq 1 + n^{H((1+\delta)p)} + 4n \cdot \left(\frac{e^\delta}{(1+\delta)^{1+\delta}} \right)^{dp}, \end{aligned}$$

where $H(x)$ is the binary entropy of $x \in (0, 1)$. By choosing $\delta = 0.315$, we infer that the last expression is at most $5n^{0.9}$. Using Markov's inequality, $\Pr[|I_t| \geq \log n \cdot \mathbf{E}[|I_t|]] = o(1)$. But if $|I_t| \leq 5n^{0.9} \log n$, we need at least additional $0.1 \log_2 n - o(\log n)$ rounds to reach a step $\tau = t + 0.1 \log_2 n - o(\log n)$ with $|I_\tau| \in [n/2, n/4]$, since $|I_t|$ can at most double in each step. As shown in [10, Theorem 1] (see also the end of the proof of Proposition 3.8), the number of uninformed nodes can at most decrease by a factor of $(1 - 1/d)^d \geq 1/e - o(1)$ in each round for any regular graph. Consequently, we need a.a.s. additional $\log n - o(\log n)$ rounds to inform all n nodes. This completes the proof. \square

References

- [1] M. Adler, S. Chakrabarti, M. Mitzenmacher, and L. E. Rasmussen. Parallel randomized load balancing. *Random Structures and Algorithms*, 13(2):159–188, 1998.
- [2] N. Alon and F. R. K. Chung. Explicit construction of linear sized tolerant networks. *Discrete Mathematics*, 72(1-3):15–19, 1988.
- [3] N. Alon and J. Spencer. *The Probabilistic Method*. Wiley-Interscience Series in Discrete Mathematics and Optimization. John Wiley & Sons, 3rd edition, 2008.
- [4] S. Boyd, A. Ghosh, B. Prabhakar, and D. Shah. Randomized Gossip Algorithms. *IEEE Transactions on Information Theory and IEEE/ACM Transactions on Networking*, 52(6):2508–2530, 2006.
- [5] A. Z. Broder, A. M. Frieze, S. Suen, and E. Upfal. Optimal construction of edge-disjoint paths in random graphs. *SIAM Journal on Computing*, 28(2):541–573, 1998.
- [6] F. Chierichetti, S. Lattanzi, and A. Panconesi. Almost tight bounds for rumour spreading with conductance. In *42nd Annual ACM Symposium on Theory of Computing (STOC'10)*, pages 399–408, 2010.
- [7] C. Cooper and A. M. Frieze. The cover time of sparse random graphs. *Random Structures and Algorithms*, 30(1-2):1–16, 2007.
- [8] B. Doerr, T. Friedrich, and T. Sauerwald. Quasirandom rumor spreading. In *19th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'08)*, pages 773–781, 2008. full version available at <http://arxiv.org/abs/1012.5351>.

- [9] D. Dubhashi and A. Panconesi. *Concentration of Measure for the Analysis of Randomized Algorithms*. Cambridge University Press, 2009.
- [10] R. Elsässer and T. Sauerwald. On the runtime and robustness of randomized broadcasting. *Theoretical Computer Science*, 410(36):3414–3427, 2009.
- [11] R. Elsässer and T. Sauerwald. Cover time and broadcast time. In *26th International Symposium on Theoretical Aspects of Computer Science (STACS'09)*, pages 373–384, 2009.
- [12] U. Feige, D. Peleg, P. Raghavan, and E. Upfal. Randomized broadcast in networks. *Random Structures and Algorithms*, 1(4):447–460, 1990.
- [13] N. Fountoulakis and K. Panagiotou. Rumor spreading on random regular graphs and expanders. In *14th International Workshop on Randomization and Computation (RANDOM'10)*, pages 560–573, 2010.
- [14] N. Fountoulakis, A. Huber, and K. Panagiotou. Reliable broadcasting in random networks and the effect of density. In *29th IEEE Conference on Computer Communications (INFOCOM'10)*, pages 2552–2560, 2010.
- [15] A. Frieze and G. Grimmett. The shortest-path problem for graphs with random-arc-lengths. *Discrete Applied Mathematics*, 10:57–77, 1985.
- [16] Z. Füredi and J. Kómlós. The eigenvalues of random symmetric matrices. *Combinatorica*, 3: 233–241, 1981.
- [17] G. Giakkoupis. Tight upper bounds for rumor spreading in graphs of a given conductance. In *28th International Symposium on Theoretical Aspects of Computer Science (STACS'11)*, pages 57–68. 2011.
- [18] G. Giakkoupis and T. Sauerwald. Rumor spreading and vertex expansion. In *23rd Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'12)*, pages 1623–1641, 2012.
- [19] S. Hoory, N. Linial, and A. Wigderson. Expander graphs and their applications. *Bulletin of the American Mathematical Society*, 43(4):439–561, 2006.
- [20] S. Janson, T. Łuczak, and A. Ruciński. *Random graphs*. Wiley-Interscience Series in Discrete Mathematics and Optimization. Wiley-Interscience, New York, 2000.
- [21] R. Karp, C. Schindelhauer, S. Shenker, and B. Vöcking. Randomized rumor spreading. In *41st Annual IEEE Symposium on Foundations of Computer Science (FOCS'00)*, pages 565–574, 2000.
- [22] M. Krivelevich and B. Sudakov. Pseudo-random graphs. In *More Sets, Graphs and Numbers, Bolyai Society Mathematical Studies 15*, Springer, pages 199–262, 2006.
- [23] C. McDiarmid. On the method of bounded differences. *Surveys in Combinatorics, London Mathematical Society Lecture Note Series 141*, pages 148–188, 1989.
- [24] C. D. Meyer. *Matrix Analysis and Applied Linear Algebra*. SIAM, 2000. available online <http://www.matrixanalysis.com/DownloadChapters.html>.
- [25] D. Mosk-Aoyama and D. Shah. Fast distributed algorithms for computing separable functions. *IEEE Transactions on Information Theory*, 54(7):2997–3007, 2008.
- [26] V. H. Vu. Spectral norm of random matrices. *Combinatorica*, 27(6):721–736, 2007.

Appendix

A Auxiliary Lemmas

Lemma A.1 ([16, 26]). *Let A be the adjacency matrix of $G(n, p)$, where $0 < p < 1/2$. Then with probability $1 - o(1)$, for every $2 \leq i \leq n$, we have $|\lambda_i| = O(\sqrt{pn})$.*

Lemma A.2 ([24, Example 7.5.2, page 551]). *Let \mathbf{A} and \mathbf{E} be the two $n \times n$, symmetric and real-valued matrices, and let $\mathbf{B} := \mathbf{A} + \mathbf{E}$. Let $\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_n$ be the n eigenvalues of \mathbf{A} , $\varepsilon_1 \geq \varepsilon_2 \geq \dots \geq \varepsilon_n$ be the n eigenvalues of \mathbf{E} and $\beta_1 \geq \beta_2 \geq \dots \geq \beta_n$ be the n eigenvalues of \mathbf{B} . Then for every i , the following double inequality holds:*

$$\lambda_i + \varepsilon_1 \geq \beta_i \geq \lambda_i + \varepsilon_n.$$

The following concentration inequalities are used to analyze the evolution of the informed nodes over time.

Theorem A.3 (Method of Bounded Differences, [23, Lemma 1.2]). *Let X_1, X_2, \dots, X_n be independent random variables, with X_i taking values in a set A_i for each i . Suppose that the function $f : \prod_{i=1}^n A_i \rightarrow \mathbb{R}$ satisfies $|f(x) - f(x')| \leq c_k$, whenever the vectors x and x' differ only in the k -th coordinate. Let $Y := f(X_1, X_2, \dots, X_n)$. Then for any $\rho > 0$,*

$$\Pr[|Y - \mathbf{E}[Y]| \geq \rho] \leq 2 \cdot \exp\left(-2\rho^2 / \sum_{k=1}^n c_k^2\right).$$

We note the following standard Chernoff bound for sum of geometric random variables which can be easily derived by using a Chernoff bound for a sum of Bernoulli random variables.

Theorem A.4 ([9]). *Suppose that X_1, \dots, X_n are independent geometric random variables on \mathbb{N} with parameter δ , so $\mathbf{E}[X_i] = 1/\delta$ for each i . Let $X := \sum_{i=1}^n X_i$, $\mu = \mathbf{E}[X] = n/\delta$. Then for any $\varepsilon > 0$,*

$$\Pr[X \geq (1 + \varepsilon)\mu] \leq e^{-\varepsilon^2 n / 2(1 + \varepsilon)}.$$