Tight Bounds for Quasirandom Rumor Spreading

Spyros Angelopoulos Benjamin Doerr Anna Huber
Konstantinos Panagiotou

Max-Planck-Institut für Informatik, 66123 Saarbrücken, Germany

{sangelop, bdoerr, ahuber, kpanagio}@mpi-inf.mpg.de

Submitted: Apr 29, 2009; Accepted: Jul 30, 2009; Published: Aug 7, 2009
Mathematics Subject Classification: 68Q87, 68W20, 68R05, 60-08

Abstract

This paper addresses the following fundamental problem: Suppose that in a group of n people, where each person knows all other group members, a single person holds a piece of information that must be disseminated to everybody within the group. How should the people propagate the information so that after short time everyone is informed?

The classical approach, known as the *push model*, requires that in each round, every informed person selects some other person in the group at random, whom it then informs. In a different model, known as the *quasirandom push model*, each person maintains a cyclic list, i.e., permutation, of all members in the group (for instance, a contact list of persons). Once a person is informed, it chooses a random member in its own list, and from that point onwards, it informs a new person per round, in the order dictated by the list.

In this paper we show that with probability $1 - o(1)$ the quasirandom protocol informs everybody in $(1 \pm o(1)) \log_2 n + \ln n$ rounds; furthermore we also show that this bound is tight. This result, together with previous work on the randomized push model, demonstrates that irrespectively of the choice of lists, quasirandom broadcasting is as fast as broadcasting in the randomized push model, up to lower order terms. At the same time it reduces the number of random bits from $O(\log^2 n)$ to only $\lceil \log_2 n \rceil$ per person.

1 Introduction

Randomized Broadcast in Networks Information spreading in large networks is an important topic of study with several applications in distributed systems. Consider, for instance, the task of maintaining replicated databases on name servers in large networks [5, 9]. Here, the goal is to propagate updates (that originate at some specific vertex or vertices) to all other vertices in the network. This is typically done by means of information exchange between pairs of nodes. Namely, a pair of neighboring vertices check whether their copies of the database are in agreement, and subsequently perform the necessary updates. In order to guarantee fast dissemination of the information, it is important that these pairs of vertices are chosen suitably. Moreover, it is desirable, and many times a requirement, that the broadcasting algorithms are simple, resilient against failures, and that they operate locally, i. e., the vertices should not require knowledge of the global network topology. Similar broadcasting scenarios have been investigated in the mathematics of infectious diseases (see e.g. [12]). Here, the desired event in such a setting is that not too many nodes become infected. These requirements have also motivated the study of *gossip-based* multicast protocols in distributed networks (see, for instance, [1], and references therein).

The classical broadcasting protocol is the so-called *push model*, also known, in more informal terms, as *randomized rumor spreading*. In this protocol, the information is initially known to a single vertex and is spread iteratively to all other vertices in a randomized fashion. More specifically, the protocol proceeds in rounds (and as a consequence, we must assume some form of synchronization). We assume an underlying network, modeled by an unweighted graph, in which a vertex can transmit messages to all its neighbors, but only one message per vertex may be transmitted in any round, due to bandwidth constraints. In each round, every informed vertex chooses a neighbor uniformly at random, to which it then sends the information. This protocol is simple, local, and robust on sufficiently dense graphs. The important question is how many rounds are required in order to inform all vertices in the graph.

If the underlying network is the complete graph K_n on n vertices, Frieze and Grimmet [11] showed that $(1 + o(1))(\log_2 n + \ln n)$ rounds suffice with probability $1 - o(1)$. This was improved to $\log_2 n + \ln n + h(n)$ for every $h \in \omega(1)$ by Pittel [15]. Feige, Peleg, Raghavan, and Upfal [9] were the first to give bounds for general graphs. Moreover, they proved that $O(\log n)$ rounds suffice for the hypercube and the random graph $G_{n,p}$ with $p \geqslant (1 + \varepsilon) \ln n/n$. Elsässer and Sauerwald determined in [8] similar bounds for several Cayley graphs.

The Quasirandom Push Model Recently, Doerr, Friedrich, and Sauerwald [7] proposed a *quasirandom* model for randomized rumor spreading. The basic setup is the same as in the randomized push model, where in each time-step every informed vertex contacts one of its neighbors. However, the choices of these neighbors are not stochastically independent. Instead, each vertex has a fixed, cyclic list of its neighbors, which dictates the order in which the vertex contacts them. The first neighbor to be contacted by the vertex is determined by choosing a starting position in this cyclic list at random, and independently of the choices of the other vertices. From that point onwards, in each round, the vertex informs one new vertex per round, in the order dictated by the list. The appealing property of this protocol is that each vertex needs to make only a *single* random choice, as opposed to a new random choice in each round of the game. This adds to the simplicity and robustness of the protocol. However, it is not immediately clear whether the protocol is as efficient as the (standard) push model, in terms of the number of rounds required to guarantee full dissemination of the rumor.

The work of [7] initiated a study of the quasirandom push model, and showed that this model is efficient in the following sense. For the complete graph, the hypercube and the random graph as defined earlier, they showed that all nodes are informed within $O(\log n)$ rounds, with probability $1 - o(1)$, *independently* of the particular choice of the lists. The $O(\log n)$ bound also holds for much sparser random graphs, for which the classical model needs incurs more rounds. We refer the reader to [7] for a detailed exposition and comparison.

Our contribution: results and methods In this paper, we provide an exact evaluation of the quasirandom push model in the setting in which each node has complete knowledge of every other node (neighbor) in the network, and thus can inform all nodes, i. e., the setting can be modeled by a complete graph. This is a simple, yet fundamental model that has been studied extensively in the context of distributed networks, as well as in the context of randomized broadcast protocols in the push model. We emphasize that [7] has shown only an asymptotic analysis of the performance of the quasirandom protocol that ignores all multiplicative factors (which could be fairly large). In other words, it could be the case that the quasirandom protocol is far inferior to the standard push model. Our main result affirms the efficiency of quasirandomness:

Theorem 1. *Let S_n denote the number of rounds necessary to inform all vertices in the quasirandom model on the complete graph on n vertices. Then,* independently of the choice of the lists*, with probability* $1 - o(1)$*,*

$$|S_n - (\log_2 n + \ln n)| = o(\log n).$$

Combined with the main result in [11] on rumor spreading in complete graphs, the above theorem demonstrates that quasirandom broadcasting is asymptotically precisely as fast as random broadcasting, *irrespectively* of the choice of the lists at each vertex. An important corollary is that the quasirandom model matches the performance of the fully random model, while at the same time reducing the number of random bits from $O(\log^2 n)$ to only $\log_2 n$ at each vertex.

Theorem 1 is derived by obtaining tight lower and upper bounds on the number of required rounds (c.f. Theorem 3, and Theorem 4, respectively). The crucial difficulty in obtaining tight bounds is that the transmission of messages at various nodes is not independent, but instead may vary by very little or by very much, depending on the choice of the lists. To circumvent this difficulty, we need to introduce some techniques, which we believe could be applicable to the analysis of existing or future randomized protocols with varying degrees of dependency. For the lower bound, we define and analyze a random experiment based on the definition of the quasirandom protocol, which is much more "powerful", in the sense that it allows some vertices to be informed at no cost. This enables us to describe the process with the help of supermartingales, which makes a precise analysis possible.

For the upper bound, we use so-called *delaying techniques*, where we assume that a vertex stops informing, or that a vertex does not transfer the information on for a certain number of rounds. This makes the protocol only slower. On the other hand, it enables us to perform a very precise analysis by applying well-known tools from probability theory like Chernoff's and Azuma's inequalities.

Very recently, by using different methods, Fountoulakis and Huber [10] showed that the "$o(\log n)$ " error-term in the statement of Theorem 1 can be replaced by slower growing functions.

Related Work In this subsection we sketch briefly the connection of our work to the general concept of quasirandomness. The main underlying idea is to imitate particular properties of a random process deterministically. This concept occurs in several areas of mathematics and computer science. Prominent examples are low-discrepancy point sets and Quasi-Monte Carlo Methods (see, e. g., [14]).

A particular example that inspired our work is a quasirandom analogue of random walks introduced by [16] and later popularized by Jim Propp (see, e. g., [4, 2, 3]). To imitate the property of a random walk that many visits to a vertex result in a balanced number of moves going from it to each of its neighbors, each vertex is equipped with a rotor always pointing to a neighbor together with a cyclic permutation of the neighbors. A walk arises from leaving the current vertex in the rotor direction and then updating the rotor to the next neighbor according to the order given by the permutation. Some beautiful results exist on this model. Particularly, Cooper and Spencer [4] showed that if an arbitrary large population of particles does such a quasirandom walk on an infinite grid $\mathbb{Z}^d$, then (under some mild conditions) the number of particles on any vertex at each time deviates from the corresponding expected value by only a constant c_d. This constant is independent of the number of particles, their initial position, and the cyclic permutations used by the rotors. More specifically, for $d = 1$, that is, the graph being the infinite path, the best possible constant is $c_1 \approx 2.29$ [3]. For the two-dimensional grid, the best possible constant satisfies $c_2 < 8.03$ [6].

2 Model, Notation and Preliminaries

We denote by $G = (V, E)$ the underlying complete graph and by $n := |V|$ the number of vertices (or nodes). Each vertex $v \in V$ is associated with a (cyclic) list L_v of its neighbors that can be viewed as a permutation L_v of $V \setminus \{v\}$.

Suppose that $v \in V$ obtains the rumor for the first time in round t. Then it chooses a position p_v on its list uniformly at random as its starting point. In the following round, $t + 1$, vertex v contacts the vertex $L_v[p_v]$ and informs it, if it was not already. In each subsequent round $t + 1 + t'$, where $t' = 1, 2, \ldots$, it sends the rumor to vertex $L_v[p_v + t' \bmod |N(v)|]$. This vertex then becomes informed, if it was not already.

The focus of this work is to investigate how long it takes until some rumor initially known only to a single vertex is broadcast to all other vertices. We adopt a worst-case view in that we aim at bounds that are independent of all the lists.

As a technical tool, in order to estimate the sizes of sets of informed vertices we will apply Azuma's inequality several times. We will use it in the following form, see e. g. [13, Lemma 1.2].

Lemma 2 (Azuma's inequality)**.** *Let $X_1, \dots, X_n$ be independent random variables, with X_k taking values in a set A_k for each k. Suppose a function $f : \prod_{k=1}^n A_k \to \mathbb{R}$ is measurable and there exist quantities $c_k > 0$ such that $|f(x) - f(x')| \leqslant c_k$ whenever the vectors x and x' differ only in the kth coordinate. Let Y be the random variable $f(X_1, \dots, X_n)$. Then for any $t > 0$,*

$$\Pr\left(|Y - E(Y)| \geqslant t\right) \leqslant 2 \exp\left(-\frac{2t^2}{\sum_{k=1}^n c_k^2}\right).$$

3 The Lower Bound

This section is devoted to the proof of the lower bound in Theorem 1. In particular, we show the following statement.

Theorem 3. *Let $\varepsilon \in (0, 1)$. With high probability, for* any *choice of the lists, the number of rounds required in the quasirandom model to inform all n vertices of a complete graph is at least*

$$\log_2 n + (1 - \varepsilon) \ln n.$$

Proof. Let $\ell = \lfloor (1 - \frac{\varepsilon}{2}) \ln n \rfloor$, and denote by X an execution of the quasirandom protocol on a complete graph with $|V| = n$ vertices. Our proof will proceed by first defining an auxiliary experiment X', which is based on the execution X. This experiment is related to X, but has a considerably simpler structure, which will make the further analysis more accessible.

We begin by describing X', which consists of the following three *phases*:

- **Phase 1:** Run X until $\lceil \frac{n}{\ln^2 n} \rceil$ vertices are informed. Let $\mathcal{A}$ denote the set of these vertices, and denote by R the number of elapsed rounds.

- **Phase 2:** For each vertex $v \in \mathcal{A}$ denote by i_v the number of vertices contacted by v in Phase 1. Designate all vertices $L_v[p_v], \dots, L_v[p_v + i_v + \ell - 1 \pmod{n-1}]$ as informed. Let

$$\mathcal{B} = \bigcup_{v \in \mathcal{A}} \left\{ L_v[p_v], \dots, L_v[p_v + i_v + \ell - 1 \pmod{n-1}] \right\} \setminus \mathcal{A}.$$

  Set also $\mathcal{C} = V \setminus (\mathcal{A} \cup \mathcal{B})$.

- **Phase 3:** The last phase consists of $m = |\mathcal{B} \cup \mathcal{C}| = n - |\mathcal{A}|$ *iterations*, where we shall write $\mathcal{B} \cup \mathcal{C} = \{v_1, \dots, v_m\}$. In iteration i, vertex v_i informs simultaneously all the vertices $L_{v_i}[p_{v_i}], \dots, L_{v_i}[p_{v_i} + \ell - 1 \pmod{n-1}]$.

Informally, in the first phase X' simulates X, until a certain number of vertices has been informed. This defines the vertex set $\mathcal{A}$. In Phase 2 we allow every vertex in $\mathcal{A}$ to contact immediately ℓ additional vertices. Finally, X' considers each vertex in the set $\mathcal{B} \cup \mathcal{C}$, and contacts ℓ consecutive vertices in its list, starting at the same index as the one chosen in X.

Clearly, the random experiment X' and the execution X of the quasirandom protocol are related. However, there are two important differences: First, we allow all vertices in $\mathcal{B}$ to become informed at no cost (i.e., immediately). More importantly, in the third phase we allow a vertex $v \in \mathcal{C}$ to start informing other vertices, even though v is not informed yet. At an informal level, we may then expect that the total number of vertices that never became informed after X' is completed does not exceed the number of vertices that never became informed after the first $R + \ell$ rounds in X.

We will now formalize the latter claim. Let U_X be the random variable that denotes the number of vertices that are not informed after $\xi = R + \ell$ rounds of X. Similarly, let $U_{X'}$ denote the number of vertices that are not informed after X' terminates. We argue below that

$$U_X \geqslant U_{X'}, \tag{1}$$

i.e., for every admissible choice of the p_v's, the total number of uninformed vertices after ξ rounds in X is at least as big as the total number of uninformed vertices after X' was executed. Moreover, we will show that with high probability over the random choice of the p_v's there will be uninformed vertices after the termination of X', i.e.,

$$\Pr\left(U_{X'} > 0\right) = 1 - o(1). \tag{2}$$

With this facts in hand we can prove the theorem as follows. First, note that (1) and (2) imply that with high probability, after $R + \ell$ rounds in the execution of X, there will be at least one uninformed vertex. Second, note that $R \geqslant \log_2(\frac{n}{\ln^2 n}) = \log_2 n - o(\ln n)$ always, as the number of informed vertices at most doubles in each round. The proof is completed.

What remains is to show (1) and (2). To see (1), note that in X' every $v \in \mathcal{A}$ contacted $i_v + \ell$ vertices, while every $v \in V \setminus A$ contacted ℓ vertices (where the sets of contacted vertices consist of consecutive vertices in the corresponding lists L_v, starting at the corresponding index p_v). On the other hand, due to the definition of the quasirandom model, after ξ rounds each vertex $v \in V$ informed $\xi - Q(v)$ other vertices, where $Q(v)$ is the number of the round in which v was informed for the first time (or ξ if it was not informed during the first ξ rounds). By definition, $Q(v) = R - i_v$ for all $v \in \mathcal{A}$. Moreover, $Q(v) \geqslant R$ for all $v \in V \setminus \mathcal{A}$, due to the definition of $\mathcal{A}$. It follows that in X, the number of vertices informed by any $v \in \mathcal{A}$ is $i_v + \ell$, and otherwise it is $\leqslant \ell$. This completes the proof of (1).

Next we show (2). Let us make some auxiliary observations first. Note that for large n

$$|\mathcal{B}| \leqslant |\mathcal{A}| \cdot \ell = \Theta\left(\frac{n}{\ln n}\right) \quad \text{and hence} \quad |\mathcal{C}| \geqslant n\left(1 - \Theta\left(\frac{1}{\ln n}\right)\right). \tag{3}$$

For $0 \leqslant i \leqslant m = |\mathcal{B} \cup \mathcal{C}|$ denote by Z_i the number of uninformed vertices after the ith iteration in Phase 3 of X'. Clearly, $Z_0 = |\mathcal{C}|$. We will show that for all $i \geqslant 0$

$$\mathrm{E}\left(Z_i\right) \geqslant \alpha^i |\mathcal{C}| \quad \text{and} \quad \mathrm{E}\left(Z_i^2\right) \leqslant \beta^i |\mathcal{C}|^2 + \ell n(\alpha^i - \beta^i) + \ell, \tag{4}$$

where we abbreviated

$$\alpha = 1 - \frac{\ell}{n-1} \quad \text{and} \quad \beta = 1 - \frac{2\ell}{n-1}.$$

This, combined with the Second Moment Method, proves (2) as follows:

$$\begin{aligned} \Pr\left(U_{X'} = 0\right) = \Pr\left(Z_m = 0\right) &\leqslant \frac{\mathrm{Var}[Z_m]}{\mathrm{E}\left(Z_m\right)^2} = \frac{\mathrm{E}\left(Z_m^2\right)}{\mathrm{E}\left(Z_m\right)^2} - 1 \\ &\overset{(3),(4)}{\leqslant} (1 + o(1)) \frac{\beta^m n^2 + \ell n \alpha^m + \ell}{\alpha^{2m} n^2} - 1. \end{aligned}$$

Recall that $m = n - |\mathcal{A}| = (1 - o(1))n$. Using this, the claim then follows from the easy observations $\frac{\ell}{\alpha^m n} = \Theta(\frac{\log n}{n^{\varepsilon/2}}) = o(1)$, $\frac{\beta^m}{\alpha^{2m}} \leqslant 1$, and $\frac{\ell}{\alpha^{2m} n^2} = o(1)$.

We now prove the first statement in (4). Denote by $\mathcal{Z}_{i-1}$ the set of uninformed vertices after the first $i-1$ iterations in Phase 3 were executed, and define for $v \in \mathcal{Z}_{i-1}$ by I_v the indicator variable for the event that vertex v becomes informed in iteration i. Clearly, as there are precisely ℓ indexes such that v_i informs any $v \in \mathcal{Z}_{i-1} \setminus \{v_i\}$ we readily obtain $\mathrm{E}\left(I_v \mid \mathcal{Z}_{i-1}\right) = \frac{\ell}{n-1}$. Using that $Z_i = Z_{i-1} - \sum_{v \in \mathcal{Z}_{i-1}} I_v$ we obtain by linearity of expectation

$$\mathrm{E}\left(Z_i \mid \mathcal{Z}_{i-1}\right) \geqslant Z_{i-1} - Z_{i-1} \frac{\ell}{n-1} = \alpha Z_{i-1}.$$

By using this the first statement in (4) follows from $\mathrm{E}\left(Z_i\right) = \mathrm{E}\left(\mathrm{E}\left(Z_i \mid \mathcal{Z}_{i-1}\right)\right)$ and induction. Similarly, we obtain $\mathrm{E}\left(Z_i \mid \mathcal{Z}_{i-1}\right) \leqslant Z_{i-1} - (Z_{i-1} - 1)\frac{\ell}{n-1} = \alpha Z_{i-1} + \frac{\ell}{n-1}$, which implies

$$\mathrm{E}\left(Z_i\right) \leqslant \alpha^i |\mathcal{C}| + 1. \tag{5}$$

Finally, we show the second statement in (4). Again, using that $Z_i = Z_{i-1} - \sum_{v \in \mathcal{Z}_{i-1}} I_v$ we obtain

$$\mathrm{E}\left(Z_i^2 \mid \mathcal{Z}_{i-1}\right) = Z_{i-1}^2 - 2 Z_{i-1} \mathrm{E}\left(\sum_{v \in \mathcal{Z}_{i-1}} I_v \;\middle|\; \mathcal{Z}_{i-1}\right) + \mathrm{E}\left(\Big(\sum_{v \in \mathcal{Z}_{i-1}} I_v\Big)^2 \;\middle|\; \mathcal{Z}_{i-1}\right). \tag{6}$$

Let j_v be the position of v in the list L_{v_i}, and denote for $v, v' \in \mathcal{Z}_{i-1}$ by $d(v,v')$ the distance of the positions of v and v', i.e. $d(v,v') = \min\{|j_v - j_{v'}|, (n-1) - |j_v - j_{v'}|\}$. With this notation we may write

$$\mathrm{E}\left(I_v \cdot I_{v'} \mid \mathcal{Z}_{i-1}\right) = \frac{\max\{\ell - d(v,v'), 0\}}{n-1},$$

as there are precisely $\max\{\ell - d(v,v'),0\}$ indexes such that v_i informs both v and v'. As for every $v \in \mathcal{Z}_{i-1}$ and every $0 < d < \ell$ there are two distinct vertices v', v'' such that $d(v,v') = d(v,v'') = d$ we obtain with the above facts that

$$\mathrm{E}\left(\Big(\sum_{v\in\mathcal{Z}_{i-1}} I_v\Big)^2 \;\Big|\; \mathcal{Z}_{i-1}\right) \leqslant \sum_{v\in\mathcal{Z}_{i-1}}\sum_{0<d<\ell}\frac{2(\ell-d)}{n-1} + Z_{i-1}\frac{\ell}{n-1} \leqslant Z_{i-1}\frac{\ell^2}{n-1}.$$

By plugging this bound into (6) we obtain

$$\begin{aligned}\mathrm{E}\left(Z_i^2 \mid \mathcal{Z}_{i-1}\right) &\leqslant Z_{i-1}^2 - 2Z_{i-1}\cdot Z_{i-1}\frac{\ell}{n-1} + Z_{i-1}\frac{\ell^2}{n-1}\\ &= \left(1-\frac{2\ell}{n-1}\right)Z_{i-1}^2 + Z_{i-1}\frac{\ell^2}{n-1}.\end{aligned}$$

Using this and the upper bound in (5) we may infer that $\mathrm{E}\left(Z_i^2\right) \leqslant \beta\mathrm{E}\left(Z_{i-1}^2\right)+\alpha^{i-1}\ell^2+\frac{\ell^2}{n-1}$. A simple inductive argument then shows that

$$\mathrm{E}\left(Z_i^2\right) \leqslant \beta^i|\mathcal{C}|^2 + \ell^2\cdot\frac{\alpha^i-\beta^i}{\alpha-\beta} + \frac{\ell^2}{n-1}\cdot\frac{1-\beta^i}{1-\beta} \leqslant \beta^i|\mathcal{C}|^2 + \ell n(\alpha^i-\beta^i) + \ell,$$

and the proof of (4) is completed. □

4 The Upper Bound

In this section we will bound the broadcast time of the quasirandom model on complete graphs from above. We show the following version of the upper bound in Theorem 1, which also includes an estimate on the failure probability.

Theorem 4. *Let $\varepsilon > 0$. There exist $\eta = \eta(\varepsilon) > 0$ and $n_0 \in \mathbb{N}$ such that for all $n \geqslant n_0$ the quasirandom model on the complete graph on n vertices informs all vertices in*

$$(1+\varepsilon)(\log_2 n + \ln n)$$

time-steps with probability at least $1 - n^{-\eta}$.

In the following we give an overview of the main techniques in our proof.

Delaying techniques In the analysis of the upper bound of the quasirandom push model it will occasionally be convenient to assume that a vertex stops informing (*cut off*) or that a vertex does not transfer the information on for a certain number of rounds after receiving the rumor (*delay*). This means that other vertices receive the rumor later. Consequently, the random variable describing the broadcast time of in these ways modified models strictly dominates the broadcast time of the quasirandom model. Of course, this also holds if several vertices stop or delay the propagation of the rumor. Any quasirandom push model with arbitrary cut offs and delays is called a *delayed model.*

Lemma 5. *Any random variable describing the broadcast time of a delayed model dominates the broadcast time of the quasirandom model.*

Lazy and busy phases We split the analysis into *phases* consisting of several rounds. In each phase we assume that the only vertices that can inform are those which informed nobody else in the previous phases. Note that by invoking the delaying concept, we have an influence over which vertices have not informed other vertices. The vertices that become informed during the current phase and do not inform any other vertices are called *newly informed in this phase*.

In *lazy* phases we will assume that the action of all nodes that become informed during the phase is delayed until the next phase. This means that only the newly informed vertices from the previous phase inform, and all vertices, which become informed by them, form the set of newly informed vertices.

The concept of lazy phases was already used in [7]. They greatly simplify the analysis, but they often leave a constant fraction of the informed vertices inactive. In order to obtain bounds for the broadcasting time that are correct up to the leading constant, we need to perform a much tighter analysis. The new contribution of this work is to show how phases with much less delaying can be analyzed. That is, we allow and exploit the fact that vertices, which become informed within the current phase, start informing other vertices already from the very next round on. We call such a phase a *busy* one. The analysis of busy phases is the heart of this work, and differs substantially from the strategy used in [7].

Overview of the Analysis The proof of Theorem 4 is divided into four main steps described in the lemmas below. To improve readability, we often allow non-integer values in contexts where integers are called for (numbers of rounds, vertices, etc.). In all such cases, however, it is easy to see that rounding up or down appropriately yields a correct proof.

The following lemma says that after only a few rounds we have logarithmically many newly informed vertices. We will use lazy phases in the proof.

Lemma 6. *Let $\delta > 0$. There exists a delayed model on the complete graph on n vertices such that after $2\delta \ln n$ rounds there will be $(\delta \ln n)^2$ newly informed vertices and $1+\delta \ln n+(\delta \ln n)^2$ overall informed vertices with probability $1 - \frac{(\delta \ln n)^4}{n-1}$.*

The main part of our proof is to show the following lemma, which states that after roughly $\log_2 n$ phases, a constant fraction of the vertices will be informed. The proof uses busy phases.

Lemma 7. *Let $\delta \in (0, \frac{1}{4})$ and assume that in some delayed model on the complete graph on n vertices we have $(\delta \ln n)^2$ newly informed vertices and $1 + \delta \ln n + (\delta \ln n)^2$ overall informed vertices. Then there exist $\nu = \nu(\delta), \zeta = \zeta(\delta)$ with $0 < \nu \leqslant \zeta$ and $\lim_{\delta \to 0} \zeta = 0$ and a delayed model such that after $(1 + 3\delta) \log_2 n$ rounds, we will have at least νn newly informed vertices and at most ζn informed vertices with probability $1 - n^{-\omega(1)}$.*

Once a constant fraction of the vertices knows the rumor, we may again go on with lazy phases to show that after a constant number of rounds, almost all vertices will be informed.

Lemma 8. *Let $0 < \nu \leqslant \zeta < \frac{1}{4}$ and assume that in some delayed model on the complete graph on n vertices we have at least νn newly informed vertices and at most ζn informed vertices. Then there exist constants $K = K(\nu, \zeta), c = c(\nu, \zeta)$ such that after performing a lazy phase of K rounds, we will have at least $(1 - 4\zeta)n$ newly informed vertices with probability $1 - e^{-cn}$.*

The fact of having all but a constant fraction of vertices newly informed will allow us to show that after roughly $\ln n$ additional rounds all vertices will be informed.

Lemma 9. *Let $\eta \in (0, 1)$ and assume that in some delayed model on the complete graph on n vertices we have at least $(1 - \eta)n$ newly informed vertices. Then within $(1 + \alpha) \ln n$ rounds, where $\alpha = \alpha(\eta) := \frac{2\eta}{1-\eta}$, we will have all vertices informed with probability $1 - \eta n^{-\eta}$.*

We now show how to derive the proof of Theorem 4, using Lemmas 6–9.

Proof of Theorem 4. Let $\varepsilon > 0$. Let $0 < \delta < \frac{1}{4}$ be such that if we choose $\zeta = \zeta(\delta)$ as in Lemma 7, $\eta = 4\zeta$ and $\alpha = \alpha(\eta)$ as in Lemma 9, we have $\alpha \leqslant \frac{\varepsilon}{4}$ and $3\delta \leqslant \varepsilon$.

We will construct a delayed model in which all vertices are informed within $(1 + \varepsilon)(\log_2 n + \ln n)$ steps with probability $1 - n^{-\eta}$. First, we apply Lemma 6. This guarantees that after at most $\frac{2}{3}\varepsilon \ln n$ rounds there are $(\delta \ln n)^2$ newly informed vertices and $1 + \delta \ln n + (\delta \ln n)^2$ overall informed vertices. Next, by applying Lemma 7 and Lemma 8 we infer that there are at least $(1 - 4\zeta)n$ newly informed vertices. For this we used at most $(1 + \varepsilon) \log_2 n + K$ more rounds, where K is the constant guaranteed to exist by Lemma 8. Finally, we apply Lemma 9 with $\eta := 4\zeta$, which guarantees the existence of a delayed model that will inform all remaining vertices after at most $(1 + \frac{\varepsilon}{4}) \ln n$ more rounds. The phases we described above consist in total of at most $\frac{2}{3}\varepsilon \ln n + (1+\varepsilon) \log_2 n + K + (1+\frac{\varepsilon}{4}) \ln n$ rounds, which is at most $(1 + \varepsilon)(\log_2 n + \ln n)$ for suitably large n.

It remains to estimate the probability for the above events. Via a straightforward union bound argument, the probability that a vertex remains uninformed is at most the sum of the failure probabilities of the four stages. The largest failure probability is $\eta n^{-\eta}$ in the last stage. Hence, the overall failure probability is bounded from above by $n^{-\eta}$. In other words, there exists a delayed model in which with probability $1 - n^{-\eta}$ all vertices are informed within $(1 + \varepsilon)(\log_2 n + \ln n)$ steps. By Lemma 5 this yields the theorem. $\square$

The proofs of Lemmas 6–9 can be found in Sections 4.1 to 4.4. Let us introduce some notation. For each time-step t, we will denote by I_t the set of informed vertices after that time-step and by N_t the set of newly informed vertices after that time-step. Recall that for a lazy phase consisting of rounds $t_1 + 1, \ldots, t_2$, N_{t_2} contains all vertices informed in rounds $t_1 + 1, \ldots, t_2$, whereas for a busy phase, N_{t_2} contains just the vertices that were informed in round t_2. Furthermore let $I_0 = N_0$ be the set containing only the initially informed vertex.

4.1 Proof of Lemma 6

Let $\delta > 0$ and $t_1 := \delta \ln n$. We will perform two lazy phases of t_1 rounds each. In the first one, the initially informed vertex informs exactly $\delta \ln n$ other vertices and we have

$|N_{t_1}| = \delta \ln n$ and $|I_{t_1}| = 1 + \delta \ln n$ with probability 1.

Let $t_2 := 2t_1$. The second lazy phase now comprises of the rounds $t_1 + 1, \ldots, t_2$. Throughout the whole phase only the vertices of N_{t_1} are assumed to contact other vertices. We shall show that with probability at least $1 - \frac{(\delta \ln n)^4}{n-1}$, only uninformed vertices are called and none of them more than once. Hence we will end up with $|N_{t_2}| = (\delta \ln n)^2$ newly informed vertices at the end of the phase. To this purpose, we will bound the probability that a given vertex v from N_{t_1} contacts a vertex, which has already been informed or is contacted also by another vertex of N_{t_1}. Apart from v itself, there are $\delta \ln n$ vertices already informed in the beginning of the phase. And there are at most $(\delta \ln n - 1)\delta \ln n$ vertices contacted by one of the other $\delta \ln n - 1$ vertices of N_{t_1}. So independent from the choices of the other vertices, there are at most $(\delta \ln n)^2$ vertices v has to "avoid" in the random segment of its list of length $\delta \ln n$. This means that the probability that v contacts a vertex, which has already been informed or is contacted also by another vertex of N_{t_1} is bounded from above by $\frac{(\delta \ln n)(\delta \ln n)^2}{n-1} = \frac{(\delta \ln n)^3}{n-1}$. Hence by a simple union bound, the probability that there exists such a vertex v in the set N_{t_1} is bounded from above by $\frac{(\delta \ln n)(\delta \ln n)^3}{n-1} = \frac{(\delta \ln n)^4}{n-1}$. In conclusion,

$$\Pr\left(|N_{t_2}| = (\delta \ln n)^2\right) \geqslant 1 - \frac{(\delta \ln n)^4}{n-1}.$$

4.2 Proof of Lemma 7

Before we proceed with the proof of this lemma we will introduce and investigate a delayed model that performs busy phases of constant length $k \geqslant 5$. Recall that in a busy phase, all vertices that become informed during the phase may start informing from the very next round on. Nevertheless, we will use some delaying in form of the "early bird assumption", which will be explained later. We stop after the phase where we have for the first time at least λn informed vertices in total, where

$$\lambda := 2^{-k}\zeta \quad \text{and} \quad \zeta := \frac{1 - 2^{-1/k}}{4k}. \tag{7}$$

We will first analyze a busy phase that consists of k rounds for any constant $k \geqslant 5$.

Analysis of a Single Busy Phase Let us assume that a busy phase of constant length $k \geqslant 5$ starts after time-step t, and additionally, that $|I_t| < \lambda n$, where λ is as in (7).

Let us make a few preliminary observations. First, for all $i \in \{1, \ldots, k\}$ we have

$$|N_{t+i}| \leqslant 2^{i-1}|N_t|, \tag{8}$$

as $|N_{t+1}| \leqslant |N_t|$ and after that step the number can at most double during each step. Thus, since $|N_t| \leqslant |I_t| < \lambda n$, we have

$$|I_{t+k}| = |I_t| + \sum_{i=1}^{k} |N_{t+i}| \leqslant |I_t| + \sum_{i=1}^{k} 2^{i-1}|N_t| \leqslant 2^k|I_t| < 2^k\lambda n = \zeta n. \tag{9}$$

The following proposition completes the analysis of a busy phase by giving corresponding lower bounds in the case that $|N_t| \geqslant (\delta \ln n)^2$ for a $\delta \in (0, \frac{1}{4})$, i.e., we are in a situation where the previous phase produced at least $(\delta \ln n)^2$ newly informed vertices.

Proposition 10. *Let $k \geqslant 5$ and $\delta \in (0, \frac{1}{4})$. Assume that in some delayed model on the complete graph on n vertices after t rounds we have $|N_t| \geqslant (\delta \ln n)^2$ and $|I_t| < \lambda n$, where λ is as in (7). Then there exists a delayed model such that with probability $1 - n^{-\omega(1)}$ the number of newly informed vertices after k rounds satisfies*

$$|N_{t+k}| \geqslant 2^{k-2}|N_t|, \tag{10}$$

and the total number of vertices informed at the end of the phase is

$$|I_{t+k}| \geqslant |I_t| + \frac{2^k - 1}{2}|N_t|. \tag{11}$$

With this proposition at hand, we now can prove Lemma 7.

Proof of Lemma 7. Let $0 < \delta < \frac{1}{4}$ and assume that at time-step t_2 in some delayed model on the complete graph on n vertices we have $|N_{t_2}| = (\delta \ln n)^2$ newly informed vertices and $|I_{t_2}| = 1 + \delta \ln n + (\delta \ln n)^2$ overall informed vertices. We perform busy phases of k rounds each, where

$$k := \frac{5}{4\delta},$$

and we stop after the phase where we have for the first time at least λn informed vertices in total, where λ as in (7). Let ℓ be the number of accomplished phases.

The preconditions of Proposition 10 are satisfied for all $t \in t_2 + \{0, \dots, \ell - 1\}k$ with probability $1 - n^{-\omega(1)}$, as $|I_t| < \lambda n$ by definition of ℓ and

$$|N_t| \geqslant (\delta \ln n)^2$$

follows inductively from $|N_{t_2}| \geqslant (\delta \ln n)^2$ and

$$|N_{t+k}| \overset{(10)}{\geqslant} 2^{k-2}|N_t| \geqslant |N_t|.$$

Having this, we can apply Proposition 10 to all $t \in t_2 + \{0, \dots, \ell - 1\}k$ and obtain that (10) and (11) hold with probability $1 - n^{-\omega(1)}$. By an iterative use of these equations, we get

$$\zeta n \overset{(9)}{>} |I_{t_2+\ell k}| \overset{(11)}{\geqslant} |I_{t_2}| + \frac{2^k - 1}{2} \sum_{s=0}^{\ell-1} |N_{t_2+sk}| \overset{(10)}{\geqslant} \frac{2^k - 1}{2} \sum_{s=0}^{\ell-1} \left(2^{k-2}\right)^s |N_{t_2}| \geqslant 2^{(k-2)\ell}$$

with probability $1 - n^{-\omega(1)}$, and thus

$$\ell < \tfrac{1}{k-2} \log_2 \zeta n.$$

So overall we performed

$$k\ell < \tfrac{k}{k-2}\log_2 \zeta n < \tfrac{k}{k-2}\log_2 n = \tfrac{5}{5-8\delta}\log_2 n < (1+3\delta)\log_2 n$$

time-steps with probability $1 - n^{-\omega(1)}$.

What remains is to show the lower bound for the number of newly informed vertices. For this we set $\nu := \frac{\lambda}{5}$. Let

$$t_3 := t_2 + \ell k.$$

By definition of ℓ this is the first time step where we have λn informed vertices in total, after an accomplished phase. This means that $\lambda n \leqslant |I_{t_3}| < \zeta n$. We will show that

$$|N_{t_3}| \geqslant \frac{1}{5}|I_{t_3}| \geqslant \nu n$$

holds with probability $1 - n^{-\omega(1)}$. To this aim we will prove by induction the more general statement that for all $s \in \{0, \ldots, \ell\}$ and $t := t_2 + sk$ one has

$$|N_t| \geqslant \frac{1}{5}|I_t|$$

with probability $1 - n^{-\omega(1)}$. For $s = 0$ the assertion holds as we have $|N_{t_2}| \geqslant (\delta \ln n)^2$ and thus $|I_{t_2}| = |N_{t_2}| + |I_{t_1}| = |N_{t_2}| + \delta \ln n + 1 \leqslant 5|N_{t_2}|$ if n is sufficiently large.

Let $s \in \{0, \ldots, \ell - 1\}$ and $t := t_2 + sk$ and assume $|N_t| \geqslant \frac{1}{5}|I_t|$. Then we have by (8)

$$|I_{t+k}| = |I_t| + \sum_{i=1}^{k} |N_{t+i}| \leqslant 5|N_t| + \sum_{i=1}^{k} |N_{t+i}| \leqslant 5|N_t| + \sum_{i=1}^{k} 2^{i-1}|N_t| = \left(2^k + 4\right)|N_t|$$

with probability $1 - n^{-\omega(1)}$. By (10) and as $k \geqslant 5$ this yields

$$|I_{t+k}| \leqslant \frac{2^k + 4}{2^{k-2}}|N_{t+k}| \leqslant 5|N_{t+k}|,$$

which shows that the statement is true for $s + 1$. □

The main work in the proof of Proposition 10 is a detailed analysis of how the information is propagated to uninformed vertices. To reduce the dependencies in this random experiment, we introduce a delaying concept as follows.

Early bird assumption. If some vertex v that was not informed at the beginning of a busy phase is contacted by two vertices x and y during that phase, and x was informed strictly before y, then we shall assume that v is informed at the time-step when x contacts v (even if y would have informed v much earlier). If there are several such x, we shall take into account the earliest contact time among them.

This assumption invokes the delaying concept: v does not start informing after being informed by y, but only after being contacted by x. Hence the resulting model is a delayed model, in the sense of Lemma 5.

One consequence of this assumption is the following. Suppose that a vertex x was informed within a busy phase with k' rounds remaining. Then, each vertex contacted by x during these k' rounds becomes actually informed by x except for the following cases: Either such a vertex y was already informed at the time when x does its random decision, or another vertex informed in the same round as x contacted y earlier.

For all $i \in \{0, \ldots, k\}$ we write N_{t+i} for the set of vertices that are newly informed in round $t+i$, *under the early bird assumption*. For the further investigation of this set, we introduce the following notation. For $j \in \{0, \ldots, i-1\}$ let N_{t+i}^{t+j} be the set of vertices that become informed by vertices of N_{t+j} in round $t+i$ under the early bird assumption.

Remark 11. *The early bird assumption ensures that N_{t+i}^{t+j} depends only on the random decisions of vertices informed up to round $t+j$. That means, at round $t+j$ we can explicitly say which vertices will be informed by vertices of N_{t+j} during the rest of the phase. Without the early bird assumption, the vertices informed by vertices of N_{t+j} might also depend on random decisions of vertices informed in some later round between $j+1$ and $k-1$. So they could not be described with the knowledge of the process up to round $t+j$.*

For all $j \in \{0, \ldots, k-1\}$ let N^{t+j} be the set of all vertices that were informed during the actual phase by vertices of N_{t+j} under the early bird assumption. This means in particular that

$$N^{t+j} = \bigcup_{i=j+1}^{k} N_{t+i}^{t+j}.$$

The formal, necessarily recursive definitions look as follows. N_{t+0} is just the set of newly informed vertices stemming from the previous phase. By the early bird assumption, for all $i \in \{1, \ldots, k\}$, the vertices in

$$N_{t+i}^{t+0} := \{v \in V \setminus I_t \mid t+i \text{ is the first time a vertex in } N_{t+0} \text{ contacts } v\}$$

are newly informed precisely in round $t+i$. Let

$$N_{t+1} := N_{t+1}^{t+0}, \quad I_{t+1} := I_t \cup N_{t+1}$$

and

$$N^{t+0} := \bigcup_{i=1}^{k} N_{t+i}^{t+0}.$$

Having defined N_{t+j+1}, I_{t+j+1} and N^{t+j} for some $j \in \{0, \ldots, k-2\}$, the early bird assumption yields for all $i \in \{j+2, \ldots, k\}$ that the vertices in

$$N_{t+i}^{t+j+1} := \{v \in V \setminus (I_t \cup N^t \cup \ldots \cup N^{t+j}) \mid t+i \text{ is the first time a vertex in } N_{t+j+1} \text{ contacts } v\}$$

are newly informed precisely in round $t+i$. Let

$$N_{t+j+1} := \bigcup_{j'=0}^{j} N_{t+j+1}^{t+j'}, \quad I_{t+j+1} := I_{t+j} \cup N_{t+j+1} \tag{12}$$

and

$$N^{t+j} := \bigcup_{i=j+1}^{k} N_{t+i}^{t+j}.$$

The main work of this section is the proof of the following proposition, where we estimate the sizes of the sets N_{t+i}^{t+j}. Here we make use of the early bird assumption.

Proposition 12. *For all $i \in \{0, \dots, k\}$ one has with probability $1 - n^{-\omega(1)}$ that*

$$\begin{aligned} |N_{t+i}| &\geqslant 2^{i-1-\frac{i}{k}}|N_t|, \text{ and} && (13)\\ \forall j \in \{0, \dots, i-1\}\colon\ |N_{t+i}^{t+j}| &\geqslant 2^{-\frac{1}{k}}|N_{t+j}|. && (14) \end{aligned}$$

With this, we can prove Proposition 10.

Proof of Proposition 10. By putting $i = k$ in (13) we directly get (10). It remains to show that (11) holds with probability $1 - n^{-\omega(1)}$. By (14)

$$|I_{t+k}| \overset{(12)}{=} |I_t| + \sum_{i=1}^{k}\sum_{j=0}^{i-1} |N_{t+i}^{t+j}| \overset{(14)}{\geqslant} |I_t| + \sum_{i=1}^{k}\sum_{j=0}^{i-1} 2^{-\frac{1}{k}}|N_{t+j}|,$$

and by (13) this is at least

$$\begin{aligned} |I_t| + 2^{-\frac{1}{k}} \sum_{i=1}^{k} \left(|N_t| + \sum_{j=1}^{i-1} 2^{j-1-\frac{j}{k}}|N_t| \right) &\geqslant |I_t| + \frac{1}{2}\sum_{i=1}^{k}\left(|N_t| + \sum_{j=1}^{i-1} 2^{j-1}|N_t|\right)\\ &= |I_t| + \frac{2^k-1}{2}|N_t|. \end{aligned}$$

□

To prove Proposition 12, we first estimate the expectation of $|N_{t+i}^{t+j}|$.

Proposition 13. *For all $i \in \{0, \dots, k\}$ and $j \in \{0, \dots, i-1\}$,*

$$\mathrm{E}\left(|N_{t+i}^{t+j}|\right) > (1-2\zeta)|N_{t+j}|.$$

Proof. Let $i \in \{0, \dots, k\}$ and $j \in \{0, \dots, i-1\}$. Then by (9), the number of vertices a priori excluded from being in N_{t+i}^{t+j} is

$$|I_t| + \sum_{\nu=0}^{j-1} |N^{t+\nu}| + \sum_{\mu=j+1}^{i-1} |N_{t+\mu}^{t+j}| \leqslant |I_{t+k}| \overset{(9)}{<} \zeta n. \tag{15}$$

Since $1 - x \leqslant e^{-x} \leqslant 1 - x + \frac{x^2}{2}$ holds for all $x > 0$, we get the following bound on the probability that one vertex gets contacted by any vertex from N_{t+j} in round $t+i$.

$$\begin{aligned} 1 - \left(1 - \frac{1}{n-1}\right)^{|N_{t+j}|} &\geqslant 1 - e^{-\frac{|N_{t+j}|}{n-1}} \geqslant \left(1 - \frac{|N_{t+j}|}{2(n-1)}\right)\frac{|N_{t+j}|}{n-1}\\ &\overset{(9)}{>} \left(1 - \frac{\zeta n}{2(n-1)}\right)\frac{|N_{t+j}|}{n-1} \geqslant (1-\zeta)\frac{|N_{t+j}|}{n-1}. \end{aligned} \tag{16}$$

Thus, conditional on the values of the outcomes of the rounds up to $t+i-1$,

$$\mathrm{E}\left(|N_{t+i}^{t+j}|\right) = \left(n - |I_t| - \sum_{\nu=0}^{j-1} |N^{t+\nu}| - \sum_{\mu=j+1}^{i-1} |N_{t+\mu}^{t+j}|\right)\left(1 - \left(1 - \frac{1}{n-1}\right)^{|N_{t+j}|}\right).$$

It readily follows from (15) and (16) that this is at least

$$n(1-\zeta)^2 \frac{|N_{t+j}|}{n-1} \;>\; (1-\zeta)^2 |N_{t+j}| \;>\; (1-2\zeta)|N_{t+j}|.$$

□

In the following proof of Proposition 12, we make finally use of the early bird assumption in form of Remark 11.

Proof of Proposition 12. We will use induction on i. The case $i = 0$ is clear, so let $i \in \{0, \ldots, k-1\}$ and assume

$$|N_{t+i'}| \geqslant 2^{i'-1}|N_t|2^{-\frac{i'}{k}} \quad \text{for all } i' \in \{0, \ldots, i\} \text{ and} \tag{17}$$

$$|N_{t+i}^{t+j}| \geqslant |N_{t+j}|2^{-\frac{1}{k}} \quad \text{for all } j \in \{0, \ldots, i-1\}.$$

We first show that for all $j \in \{0, \ldots, i\}$ one has

$$|N^{t+j}| \geqslant |N_{t+j}|(k - j - 1 + 2^{-\frac{1}{k}}) \tag{18}$$

with probability $1 - n^{-\omega(1)}$. To this aim we will give a bound on the expectation of $|N^{t+j}|$ and then use Azuma's inequality. As the N_{t+i}^{t+j}, $i \in \{j+1, \ldots, k\}$ are disjoint, we have

$$|N^{t+j}| = \sum_{i=j+1}^{k} |N_{t+i}^{t+j}|$$

by definition of N^{t+j}, and so one has by Proposition 13

$$\begin{aligned} \mathrm{E}\left(|N^{t+j}|\right) &= \sum_{i=j+1}^{k} \mathrm{E}\left(|N_{t+i}^{t+j}|\right) > (k-j)(1-2\zeta)|N_{t+j}| \geqslant (k-j-2k\zeta)|N_{t+j}| \\ &= \left(k - j - \tfrac{1-2^{-1/k}}{2}\right)|N_{t+j}|. \end{aligned} \tag{19}$$

Number the nodes of N_{t+j} from 1 to $|N_{t+j}|$. For all $v \in \{1, \ldots, |N_{t+j}|\}$ let X_v denote the random set of size $k-j$ of all vertices contacted by v. Note that $X_1, \ldots, X_{|N_{t+j}|}$ are mutually independent. Now we make use of the early bird assumption: By Remark 11, the value $|N^{t+j}|$ depends only on $X_1, \ldots, X_{|N_{t+j}|}$, namely we have

$$|N^{t+j}| = \left|\bigcup_{v=1}^{|N_{t+j}|} X_v \Big\backslash\, I_{t+j}\right| =: f(X_1, \ldots, X_{|N_{t+j}|}).$$

For all $v \in \{1, \ldots, |N_{t+j}|\}$ and for realizations $x_1, \ldots, x_{|N_{t+j}|}, x'_v$ of the random sets one has

$$|f(x_1, \ldots, x_{|N_{t+j}|}) - f(x_1, \ldots, x_{v-1}, x'_v, x_{v+1}, \ldots, x_{|N_{t+j}|})| \leqslant k - j.$$

For $j = 0$, $|N_{t+j}| \geqslant (\delta \ln n)^2 = 2^{-\frac{j}{k}}(\delta \ln n)^2$ trivially holds. By applying the induction hypothesis (17), we see that for $j \in \{1, \ldots, i\}$,

$$|N_{t+j}| \geqslant 2^{j-1-\frac{j}{k}}|N_t| \geqslant 2^{-\frac{j}{k}}|N_t| \geqslant 2^{-\frac{j}{k}}(\delta \ln n)^2.$$

With (19) and Lemma 2 we compute

$$\begin{aligned}
\Pr\left(|N^{t+j}| < |N_{t+j}|(k-j-1+2^{-\frac{1}{k}})\right) &\leqslant \Pr\left(|N^{t+j}| \leqslant \mathrm{E}\left(|N^{t+j}|\right) - \tfrac{1-2^{-1/k}}{2}|N_{t+j}|\right) \\
&\leqslant \Pr\left(\left|\mathrm{E}\left(|N^{t+j}|\right) - |N^{t+j}|\right| \geqslant \tfrac{1-2^{-1/k}}{2}|N_{t+j}|\right) \\
&\leqslant 2e^{-\frac{\left(1-2^{-1/k}\right)^2|N_{t+j}|^2}{2|N_{t+j}|(k-j)^2}} \leqslant 2e^{-\frac{\left(1-2^{-1/k}\right)^2(\delta \ln n)^2}{2^{1+j/k}(k-j)^2}} \\
&= 2n^{-\frac{\left(1-2^{-1/k}\right)^2\delta^2 \ln n}{2^{1+j/k}(k-j)^2}} = n^{-\omega(1)}.
\end{aligned}$$

As for all $j \in \{0, \ldots, i\}$ we have $|N^{t+j}_{t+i+1}| \leqslant |N_{t+j}|$, (18) yields for all $j \in \{0, \ldots, i\}$ that

$$|N^{t+j}_{t+i+1}| \geqslant 2^{-\frac{1}{k}}|N_{t+j}| \tag{20}$$

with probability $1 - n^{-\omega(1)}$. As the N^{t+j}_{t+i+1}, $j \in \{0, \ldots, i\}$ are disjoint, (20) and the induction hypothesis give us

$$\begin{aligned}
|N_{t+i+1}| &= \sum_{j=0}^{i} |N^{t+j}_{t+i+1}| \geqslant \sum_{j=0}^{i} 2^{-\frac{1}{k}}|N_{t+j}| = 2^{-\frac{1}{k}}\left(|N_t| + \sum_{j=1}^{i}|N_{t+j}|\right) \\
&\geqslant 2^{-\frac{1}{k}}\left(|N_t| + \sum_{j=1}^{i} 2^{j-1-\frac{j}{k}}|N_t|\right) \\
&\geqslant 2^{-\frac{i+1}{k}}\left(|N_t| + \sum_{j=1}^{i} 2^{j-1}|N_t|\right) = 2^{i-\frac{i+1}{k}}|N_t|
\end{aligned}$$

with probability $1 - n^{-\omega(1)}$. □

4.3 Proof of Lemma 8

Let $0 < \nu \leqslant \zeta < \frac{1}{4}$ and assume that in some delayed model on the complete graph on n vertices after t_3 rounds we have $\nu n \leqslant |N_{t_3}|$ and $|I_{t_3}| \leqslant \zeta n$. We perform one lazy phase of $K := \frac{1}{\nu} \ln \frac{1}{\zeta}$ steps. The probability that a given vertex from $V \setminus I_{t_3}$ does not get informed during this phase is

$$\left(1 - \frac{K}{n-1}\right)^{|N_{t_3}|} \leqslant e^{-\frac{K|N_{t_3}|}{n-1}} \leqslant e^{-K\nu} = \zeta.$$

So the probability that a given vertex from $V \setminus I_{t_3}$ becomes informed is at least $1-\zeta$, and with $t_4 := t_3 + K$ we obtain

$$\mathrm{E}\left(|N_{t_4}|\right) \geqslant (n - |I_{t_3}|)(1-\zeta) \geqslant (n-\zeta n)(1-\zeta) \geqslant (1-2\zeta)n. \tag{21}$$

Let $c := \frac{4\zeta}{K^2}$. To show that $|N_{t_4}| \geqslant (1-4\zeta)n$ with probability at least $1-e^{-cn}$ we will use Azuma's inequality. Number the nodes of N_{t_3} from 1 to $|N_{t_3}|$. For all $i \in \{1,\ldots,|N_{t_3}|\}$ let X_i denote the random set of size K of all vertices contacted by i. Then $X_1,\ldots,X_{|N_{t_3}|}$ are mutually independent and

$$|N_{t_4}| = \left| \bigcup_{i=1}^{|N_{t_3}|} X_i \setminus I_{t_3} \right| =: f(X_1,\ldots,X_{|N_{t_3}|}).$$

For all $i \in \{1,\ldots,|N_{t_3}|\}$ and for realizations $x_1,\ldots,x_{|N_{t_3}|}, x_i'$ of the random sets

$$|f(x_1,\ldots,x_{|N_{t_3}|}) - f(x_1,\ldots,x_{i-1},x_i',x_{i+1},\ldots,x_{|N_{t_3}|})| \leqslant K.$$

With Lemma 2 the proof completes as follows:

$$\begin{aligned}
\Pr\left(|N_{t_4}| < (1-4\zeta)n\right) &\overset{(21)}{\leqslant} \Pr\left(|N_{t_4}| \leqslant \mathrm{E}\left(|N_{t_4}|\right) - 2\zeta n\right) \\
&\leqslant \Pr\left(|\mathrm{E}\left(|N_{t_4}|\right) - |N_{t_4}|| \geqslant 2\zeta n\right) \\
&\leqslant 2e^{-\frac{8\zeta^2 n^2}{\zeta n K^2}} \\
&\leqslant e^{-cn}.
\end{aligned}$$

4.4 Proof of Lemma 9

Let $\eta > 0$ and assume that in the quasirandom model on the complete graph on n vertices after t_4 rounds we have $|N_{t_4}| \geqslant (1-\eta)n$. To inform the remaining few vertices, we use one lazy phase consisting of $(1+\alpha)\ln n$ rounds, where $\alpha = \alpha(\eta) := \frac{2\eta}{1-\eta}$. The probability that a given vertex from $V \setminus I_{t_4}$ does not get informed in this phase is

$$\begin{aligned}
\left(1 - \frac{(1+\alpha)\ln n}{n-1}\right)^{|N_{t_4}|} &\leqslant \left(1 - \frac{(1+\alpha)\ln n}{n-1}\right)^{(1-\eta)n} \\
\leqslant \ e^{-(1+\alpha)(1-\eta)\ln n} &= e^{-(1+\frac{2\eta}{1-\eta})(1-\eta)\ln n} = n^{-(1+\eta)}.
\end{aligned}$$

So the probability that all vertices get informed is

$$\begin{aligned}
&1 - \Pr\left(\text{there is a vertex in } V \setminus I_{t_4} \text{ which does not get informed}\right) \\
&\quad\geqslant\ 1 - |V \setminus I_{t_4}| \Pr\left(\text{a given vertex from } V \setminus I_{t_4} \text{ does not get informed}\right) \\
&\quad\geqslant\ 1 - |V \setminus I_{t_4}|\ n^{-(1+\eta)} \\
&\quad\geqslant\ 1 - \eta n\ n^{-(1+\eta)} \\
&\quad=\ 1 - \eta n^{-\eta}.
\end{aligned}$$

References

[1] A. Alavena, A. Demers, and J.E. Hopcroft. Correctness of a gossip based membership protocol. In *Proc. of the 16th Symposium on Principles of Distributed Computing (PODC)*, pages 292–301, New York, NY, USA, 2005. ACM.

[2] J. Cooper, B. Doerr, T. Friedrich, and J. Spencer. Deterministic random walks on regular trees. In *Proc. of the 19th ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 766–772, 2008.

[3] J. Cooper, B. Doerr, J. Spencer, and G. Tardos. Deterministic random walks on the integers. *European Journal of Combinatorics*, 28(8):2072–2090, 2007.

[4] J. Cooper and J. Spencer. Simulating a random walk with constant error. *Combinatorics, Probability and Computing*, 15(6):815–822, 2006.

[5] A. Demers, D. Greene, C. Hauser, W. Irish, J. Larson, S. Shenker, H. Sturgis, D. Swinehart, and D. Terry. Epidemic algorithms for replicated database maintenance. In *Proc. of the 6th ACM Symposium on Principles of Distributed Computing (PODC)*, pages 1–12, 1987.

[6] B. Doerr and T. Friedrich. Deterministic random walks on the two-dimensional grid. *Combinatorics, Probability and Computing*, 18(1-2):123–144, 2009.

[7] B. Doerr, T. Friedrich, and T. Sauerwald. Quasirandom rumor spreading. In *Proceedings of the 19th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 773–781, 2008.

[8] R. Elsässer and T. Sauerwald. Broadcasting vs. mixing and information dissemination on cayley graphs. In *Proceedings of the 24th International Symposium on Theoretical Aspects of Computer Science (STACS)*, pages 163–174, 2007.

[9] U. Feige, D. Peleg, P. Raghavan, and E. Upfal. Randomized broadcast in networks. *Random Structures and Algorithms*, 1(4):447–460, 1990.

[10] N. Fountoulakis and A. Huber. Quasirandom rumour spreading on the complete graph is as fast as randomized rumour spreading. Manuscript, 2009.

[11] A.M. Frieze and G.R. Grimmett. The shortest-path problem for graphs with random arc-lengths. *Discrete Applied Mathematics*, 10(1):57–77, 1985.

[12] H.W. Hethcore. Mathematics of infectious diseases. *SIAM Review 42*, pages 599–653, 2000.

[13] C. McDiarmid. On the method of bounded differences. In *Surveys in combinatorics, 1989 (Norwich, 1989)*, volume 141 of *London Mathematical Society Lecture Note Series*, pages 148–188. Cambridge Univ. Press, Cambridge, 1989.

[14] H. Niederreiter. *Random number generation and quasi-Monte Carlo methods.* SIAM, Philadelphia, PA, USA, 1992.

[15] B. Pittel. On spreading a rumor. *SIAM Journal on Applied Mathematics*, 47(1):213–223, 1987.

[16] V. B. Priezzhev, D. Dhar, A. Dhar, and S. Krishnamurthy. Eulerian walkers as a model of self-organized criticality. *Physical Review Letters*, 77(25):5079–5082, 1996.