# Randomised Broadcasting: Memory vs. Randomness*

Petra Berenbrink
School of Computing Science
Simon Fraser University
Burnaby B.C. V5A 1S6, Canada
petra@cs.sfu.ca

Robert Elsässer
Institute for Computer Science
University of Paderborn
33102 Paderborn, Germany
elsa@upb.de

Thomas Sauerwald
Max-Planck-Institut für Informatik
66123 Saarbrücken, Germany
sauerwal@mpi-inf.mpg.de

**Abstract**

In this paper we analyse broadcasting in $d$-regular networks with good expansion properties. For the underlying communication, we consider modifications of the so called random phone call model. In the standard version of this model, each node is allowed in every step to open a channel to a randomly chosen neighbour, and the channels can be used for bi-directional communication. Then, broadcasting on the graphs mentioned above can be performed in time $O(\log n)$, where $n$ is the size of the network. However, every broadcast algorithm with runtime $O(\log n)$ needs on average $\Omega(\log n/\log d)$ message transmissions per node.

In this paper we show that it is possible to save significantly on communications if the standard model is modified such that nodes can avoid opening channels to exactly the same neighbours in two consecutive steps. We consider the so called Rr model where we assume that every node has a cyclic list of all of its neighbours, ordered in a random way. Then, in step $i$ the node communicates with the $i$-th neighbour from that list. We provide an $O(\log n)$ time algorithm which produces in average $O(\sqrt{\log n})$ transmissions per node in networks with suitably defined expansion properties. Furthermore, we present a related lower bound of $\Omega(\sqrt{\log n/\log\log n})$ for the average number of message transmissions. These results show that by using memory it is possible to reduce the number of transmissions per node by almost a quadratic factor.

## 1 Introduction

We consider randomised broadcasting in (almost) regular graphs with good expansion properties. In the broadcasting problem, the goal is to spread a message from one vertex to all vertices of a network. Our interest in these graphs is motivated by overlay topologies in peer to peer (P2P) systems. Important topological properties of these networks include good connectivity, high expansion, and small diameter; all these properties are perfectly fulfilled

*An extended abstract of this paper has appeared in the Proceedings of LATIN 2010 [4]

by the graphs considered here. Our aim is to develop time-efficient broadcasting algorithms which produce a minimal number of message transmissions in the graphs described above. Since P2P systems are significant decentralised platforms for sharing data and computing resources, it is very important to provide efficient, simple, and robust broadcasting algorithms for these overlay networks. Minimising the number of transmissions is important in applications such as the maintenance of replicated databases in which broadcasts are necessary to deal with frequent updates in the system [18, 22].

In this paper we assume the so called *phone call model* (see [22]). In this model, each node $v$ may perform the following actions in every step:

1) create a new message to be broadcast

2) establish a communication channel between $v$ itself and one of its neighbours

3) transmit messages over incident channels opened by itself or by some of its $d$ neighbours

At the end of each step, the nodes close all open channels. Note that open channels can be used for bi-directional (push&pull) communications. In the case of `push` transmissions, calling nodes (i.e., the nodes that opened the channels) send their messages to their neighbours. In the case of `pull` transmissions, messages are transmitted from called nodes to the calling ones (we also say that the **called nodes** perform pull transmissions). Note that nodes do not have to send messages over open channels, they can *choose* to do so. For example, if node $v$ opens a channel to $w$, $w$ does not have to send a message to $v$. If many distinct messages are to be spread in the network, then the nodes can combine several broadcast messages to larger ones which can be sent over a channel in one time step.

In the standard phone call model it is assumed that nodes open a channel to a randomly chosen neighbour, and the nodes have to decide whether to transmit a specific message over a channel, without knowing if they opened a channel to the corresponding node in earlier steps [22]. In this paper we assume that every node has a cyclic list of all of its neighbours, ordered in a random way. In step $i$ the node opens a channel to the $i$ (modulo $d$)-th neighbour from that list. This model is called Rr *model* (Rr its standing for round robin) in the following. The Rr rule prevents a node to open a channel to a neighbour for a second time before it opened a channel to all of its neighbours. Hence, the rule helps nodes to communicate with more of their neighbours.

The question we address in this paper is whether remembering the communication partners from earlier rounds helps or not. We give a positive answer to this question and provide further evidence for the power of memory in randomised broadcasting (see [15]). More precisely, we present an algorithm, and show that w.r.t. the average number of transmissions per node this algorithm performs significantly better than any algorithm in the so called Random$[c]$-model introduced in [15] (i.e. we achieve an almost quadratic improvement). Random$[c]$ is similar to the standard random phone call model, however, every node may open channels to $c$ different randomly chosen neighbours simultaneously in each step. Our algorithm is *address oblivious*, i.e., the send decisions of the nodes do not depend on the IDs of the nodes to which they open channels in the actual step. However, the nodes are allowed to remember with which nodes they communicated in the steps before [22].

## 1.1 Related Work

There is a huge amount of work considering epidemic type (broadcasting) algorithms on graph models for P2P overlays. Most of these papers deal with the empirical analysis of

these algorithms e.g. [23, 26]. Due to space constraints, we can only describe here the results which focus on the analytical study of push&pull algorithms.

**Runtime.** Most randomised broadcasting results analyse the runtime of the push algorithm. For complete graphs of size $n$, Frieze and Grimmett [19] present an algorithm that broadcasts a message in time $\log_2(n)+\ln(n)+o(\log n)$ with a probability of $1-o(1)$. Later, Pittel [27] shows that (with probability $1-o(1)$) it is possible to broadcast a message in time $\log_2(n)+\ln(n)+f(n)$ [27], where $f(n)$ can be any slow growing function. In [18], Feige et al. determine asymptotically optimal upper bounds for the runtime of the push algorithm on $G(n,p)$ graphs (i.e., traditional Erdös-Rényi random graphs [16, 17]), bounded degree graphs, and Hypercubes. In [14] Elsässer and Sauerwald analyze certain Cayley graphs on which the push algorithm performs (asymptotically) optimal. Boyd et al. consider the combined push&pull model in arbitrary graphs of size $n$, and show that the running time is asymptotically bounded by the mixing time of a corresponding Markov chain plus an $O(\log n)$ value [5]. In [10] Doerr et al. analyse the so called quasi-random rumor spreading in an adversarial version of the RR model where the order of the lists is assumed to be given by an adversary. However, the nodes choose a random position in their lists to start with communication. They show for hypercubes and $G(n,p)$ graphs that $O(\log n)$ push steps suffice to inform every node, w.h.p.[1]. These bounds are similar to the ones in traditional randomised broadcasting (push model). These results have been extended to further graph classes with good expansion properties [11]. Recently, Doerr et al. showed in [9] that by using the RANDOM[2]-model, one can improve the running time of broadcasting on the so called preferential attachment graph [1] of size $n$ by a factor of $\log\log n$.

**Number of transmissions.** Karp et al. [22] show that in complete graphs the pull approach is inferior to the push approach, until roughly $n/2$ nodes receive the message, and then the pull approach becomes superior. They present a `push` and `pull` algorithm, together with a termination mechanism, which reduces the number of total transmissions to $O(n\log\log n)$ (w.h.p.), and show that this result is asymptotically optimal. They also consider communication failures and analyse the performance of their method in cases where the connections are established using arbitrary probability distributions.

For sparser graphs it is not possible to get $O(n\log\log n)$ message transmissions together with a broadcast time of $O(\log n)$ in the standard phone call model. In [12] Elsässer considers random $G(n,p)$ graphs, and shows a lower bound of $\Omega(n\log n/\log(pn))$ message transmissions for broadcast algorithms with a runtime of $O(\log n)$. On the positive side, for $p>\log^2 n/n$ he develops an algorithm that broadcasts in time $O(\log n)$ using $O(n\cdot(\log\log n+\log n/\log(pn)))$ transmissions, w.h.p.

In [15] the authors consider a simple modification of the standard phone call model called RANDOM[$c$] defined above. For $G(n,p)$ graphs with $p>\log^2 n/n$, they show that this modification results in a reduction of the number of message transmissions down to $O(n\log\log n)$. In [2] the authors show similar results for random $d$-regular graphs with $d=O(\log n)$. A further extension to random power law graphs has been obtained in [13]. Recently, the authors also considered quasi-random rumor spreading in random graphs and hypercubes, and obtained asymptotically optimal results w.r.t. the running time and number of message transmissions [3]. However, by using the techniques of [3], it is not possible to obtain similar results for the more general case of Edge-Node expanders considered in this paper.

[1] W.h.p. or "with high probability" means with probability $1-o(n^{-1})$

## 1.2 Models and Results

In this paper, we consider the running time and number of message transmissions produced by randomised broadcasting in more general expander graphs. We assume that every node has an estimation of $n$ which is accurate to within a constant factor. We also assume that all nodes have access to a global clock, and that they work synchronously. In each step every node can create an arbitrary amount of messages to be broadcasted. The number of message transmissions for a certain message is defined as the the number of open channels traversed by the message during the execution of the algorithm. As in [22], we assume here that new pieces of information are generated frequently in the network, and then the cost of establishing communication channels amortises over all message transmissions. However, we only concentrate on the distribution and lifetime of a single message, and consider broadcasting in the following graphs.

**Edge-Node Expanders.** Let $G = (V, E)$ be a $d$-regular graph of size $n$. For $A \subset V$, let $E(A, \overline{A})$ denote the set of edges between $A$ and $\overline{A} = V \setminus A$, and let $N(A) = \{v \in \overline{A} \mid (u, v) \in E \text{ and } u \in A\}$. For a constant $\alpha$, $G$ is called *$\alpha$-Edge-Node expander* (or simply Edge-Node expander) if the following holds:

1. For any set $A \subset V$ with $|A| \leq n/2$ we have $|E(A, \overline{A})| \geq \alpha d \cdot |A|$.
2. For any set $A \subset V$ with $|A| \leq \phi n/d$ it holds that $|N(A)| \geq \alpha d \cdot |A|$, where $\phi$ is a (large) constant.

In this paper we show three results.

- We show that there exists a family of Edge-Node Expanders for which every broadcasting algorithm with runtime $O(\log n)$ in the RANDOM$[r]$ communication model (with constant $r$) needs $\Omega(n \log n / \log\log n)$ message transmissions, w.h.p. (see Theorem 1).
- We present an algorithm in the RR model that broadcasts a message in a regular Edge-Node Expander in time $O(\log n)$ by using $O(n\sqrt{\log n})$ transmissions, w.h.p. (Theorem 2). The result holds for graphs with degree $d \in \{f(n) \cdot (\log^{3/2} n), 2^{o(\sqrt{\log n})}\}$. $f : \mathbb{N} \to \mathbb{R}$ is a function such that $\lim_{n\to\infty} f(n) = \infty$.
- We show that there exists a family of Edge-Node Expanders for which every broadcast algorithm in the oblivious model with runtime $O(\log n)$ in the RR communication model needs
$$\Omega(n\sqrt{(\log n)/\log\log n})$$
message transmissions, w.h.p. (Theorem 3).

Note that it might be very well possible to relax the bounds on $d$. The lower bound on $d$ comes from Equation 1 in Claim 4 together with Claim 3. The algorithm in Section 2.1 requires that $d = 2^{o(\sqrt{\log n})}$, since otherwise the message complexity would be larger than $O(n\sqrt{\log n})$ (cf. Phase 3).

We believe that $d \in \{f(n) \cdot (\log^{3/2} n), 2^{o(\sqrt{\log n})}\}$ covers all interesting cases. For small rvalues of $d$ it is clear that in the RR model the nodes communicate with each of their neighbours, which makes the analysis simple. On the other hand, for very large values of $d$ the difference between the standard phone call model and the RR model become negligible since it is very unlikely that nodes communicate several times with the same neighbours.

In [15, 2] we showed that in random graphs one can save on the number of message transmissions if the nodes avoid communication with the neighbours chosen already in some

recent steps. In the analysis we used the randomised construction of these graphs, and integrated the dynamical behaviour of the RANDOM[$c$] model (i.e., the parallelised version of the model described above) into the random structure of the underlying topology. However, the methods derived in [15, 2] cannot be generalised to non-random graphs with similar expansion and connectivity properties, not even to pseudo-random graphs [24]. Therefore, the main question is whether the same result also holds in graphs with random graph like properties. To answer this question, we show that there exists Edge-Node expanders that require $\Omega(n \log n / \log\log n)$ message transmissions for constant $c$ (Theorem 1).

To show Theorem 2 we introduce a new combinatorial technique which only uses the structural properties of Edge-Node expanders to show that our algorithm algorithm completes broadcasting in time $O(\log n)$ and generates $O(n\sqrt{\log n})$ message transmissions (see Sections 2.1-2.2). Our lower bound of Theorem 3 on the number of message transmissions shows that our analysis is tight up to a $\sqrt{\log\log n}$ factor.

Note that the upper bound on the number of message transmissions in the RR model is significantly smaller than the lower bound in the RANDOM[$c$] model, which substantiates the importance of memory in randomised broadcasting. Notice that all (regular) graphs $G$ for which $\lambda_2 = d - O(\sqrt{d})$ ($\lambda_2$ is the second smallest eigenvalue of the Laplacian of $G$) obey the properties described above (cf. [7, 21, 28]). Examples for such graphs are so called Ramanujan graphs which include the class of random regular graphs (cf. [10, Section 4.1.4]).

# 2 Broadcasting on Edge-Node Expanders

In this section we first consider the following lower bound w.r.t. the performance of the RANDOM[$c$] model in Edge-Node Expanders.

**Theorem 1** *Assume $A$ is a (Monte-Carlo) broadcasting algorithm with runtime $O(\log n)$ in the* RANDOM[$r$] *communication model, where $r$ is a constant. There exists a family of Edge-Node Expanders for which $A$ needs $\Omega(n \log n / \log\log n)$ message transmissions, w.h.p.*

**Proof:** We assume that $c \log n$ is the runtime of the algorithm in the RANDOM[$r$] model, where $r$ is a constant. We also assume that the structure of the graph is not known to the nodes and can not be learned during the execution of the algorithm (i.e., no information about the structure is sent from a node to another one). The theorem can be shown by extending the proof of Theorem 2 of [12] to the Cartesian product of an Erdős-Rényi random graph with a $K_{r+1}$. Here we present a somewhat different approach. The graph $G = (V, E)$ consists of a (random) $\alpha$-Edge-Node Expander $G' = (V', E')$ (with some parameter $\phi > r+1$, cf. definition of Edge-Node Expanders) of size $n$ and degree $\log^2 n - r - 1$, together with $\ell = n/(\log^2 n - r)$ cliques $K_1, \ldots K_\ell$ of size $r + 1$. For $1 \leq i \leq \ell$, all nodes of $K_i$ are connected to the same set of $\log^2 n - r$ nodes in $V'$. However, two nodes $u \in K_i$ and $v \in K_j$ are connected to different sets of nodes in $V'$, for any $i \neq j$. Hence, every node in $G$ has degree $d = \log^2 n$. In the following $G'$ is called the *original graph*, with *original nodes* and *original edges*. The clique nodes are called *clique nodes*, and edges between two node of the same clique will be called *inner edges*. Edges connecting a clique node to $G'$ are called *external edges*.

In order to show that $G$ is still an $\alpha/(2(r + 2))$-Edge-Node Expander, consider some subset $Q \subset V$ of size at most $\phi n/(d - r - 1) \geq \phi/2 \cdot (n + \ell(r + 1))/d$. Assume that

$|Q \cap V'| > |Q|/2$. Then, since $G'$ is an $\alpha$-Edge-Node Expander with parameter $\phi$, we have

$$|N(Q)| \geq \alpha d|Q \cap V'| \geq \frac{\alpha d|Q|}{2}.$$

If $|Q \cap V'| \leq |Q|/2$, then we have at least $|Q|/2$ clique nodes in $Q$. Since each of these nodes has at least $d - r$ neighbours in $V'$, and at most $r + 1$ nodes may have the same set of neighbours in $V'$, we obtain

$$|N(Q)| \geq (d - r)\frac{|Q|}{2(r + 1)} - \frac{|Q|}{2} \geq \frac{\alpha d|Q|}{2(r + 2)}.$$

The edge property from the definition of Edge-Node Expanders is shown in a similar way. Thus, $G$ is an $\alpha/(2(r + 2))$-Edge-Node Expander (with parameter $\phi/2$).

We assume that there is an algorithm $A$ which has running time $c \log n$ and that generates at most $\epsilon \cdot n \log n/ \log \log n$ message transmissions, where $\epsilon$ is a small constant. To prove this theorem we show in the following that, with a good probability, $A$ will not be able to inform all the nodes from every clique of the graph. Since $A$ sends at most $\epsilon \cdot n \log n/ \log \log n$ messages, a simple pigeonhole argument shows that there exists at least $l/2$ cliques $K_{i_1}, \ldots, K_{i_{l/2}}$ with the following property. Every node in $K_{i_1}, \ldots, K_{i_{l/2}}$ has at least $\gamma \cdot (\log^2 n - r)$ neighbours in $G'$ that send at most $2\epsilon \cdot \log n/((1 - \gamma) \log \log n)$ messages each (`push` or `pull`), where $\gamma < 1$ is a constant close to 1.

Now let us fix a node $v \in K_{i_j}$. We divide the steps into *dangerous* and *safe* steps. In a dangerous step, there are more than $(1 - \gamma + (1 - \gamma)^2) \cdot d - r$ nodes in $N(v, V')$ which perform a `pull` transmission. The steps which are not dangerous are called safe. Our goal is now to show that there is a large probability that in dangerous steps $v$ always communicates with its neighbours in $K_{i_j}$, whereas in safe steps $v$ always communicates with neighbours not performing `pull` in that step. First, we show an upper bound on the number of dangerous steps. Since a fraction $\gamma$ of the nodes from $N(K_{i_j}, V')$ perform at most $2\epsilon \log n/((1 - \gamma) \cdot \log \log n)$ message transmissions, the number of dangerous steps is bounded by

$$\frac{\gamma(d - r) \cdot 2\epsilon \log n/((1 - \gamma) \cdot \log \log n)}{(1 - \gamma)^2 d - r} < \frac{2\gamma + 1}{(1 - \gamma)^3} \cdot \epsilon \frac{\log n}{\log \log n}.$$

The probability that $v$ uses internal edges in all dangerous steps is at least

$$p' = d^{-r \cdot (2\gamma+1)\epsilon \log n/((1-\gamma)^3 \log \log n)}.$$

The probability that $v$ chooses in all safe steps only nenighbours which decided not to answer the `pull` request is at least

$$p'' = \left(\frac{(\gamma - (1 - \gamma)^2)d - r}{d}\right)^{r \cdot c \log n}.$$

Hence, the probability that all nodes of $K_{i_j}$ do not receive the message via `pull` is at least

$$P_1 = (p' \cdot p'')^{r+1} = d^{-r \cdot (2\gamma+1)\epsilon \log n/((1-\gamma)^3 \gamma \log \log n)} \left(\frac{(\gamma - (1 - \gamma)^2)d - r}{d}\right)^{r \cdot c \log n}$$

which is larger than $n^{-1/8}$ whenever $\epsilon$ is small enough and $\gamma$ is large enough.

Now we concentrate on push transmissions. As before, there are at most $\frac{2\gamma+1}{(1-\gamma)^3} \cdot \epsilon \frac{\log n}{\log\log n}$ dangerous steps. Node $v$ does not get the message via `push` in one of the dangerous steps with probability at least

$$\left(1-\frac{r}{d}\right)^{d\cdot\frac{2\gamma+1}{(1-\gamma)^3}\epsilon\frac{\log n}{\log\log n}} \geq e^{-\frac{2\gamma+1}{(1-\gamma)^3}\cdot\epsilon r\frac{\log n}{\log\log n}}.$$

Similarly, $v$ does not get the message by `push` in one of the safe steps with a probability of at least

$$\left(1-\frac{r}{d}\right)^{((1-\gamma+(1-\gamma)^2)d-r)c\log n}.$$

The probability that none of the nodes of $K_{i_j}$ receives the message via `push` is at least

$$P_2 = \left(e^{-\frac{2\gamma+1}{(1-\gamma)^3}\cdot\epsilon r\frac{\log n}{\log\log n}}\left(1-\frac{r}{d}\right)^{((1-\gamma+(1-\gamma)^2)d-r)c\log n}\right)^{r+1} \geq n^{-1/8},$$

whenever $\epsilon$ is small enough and $\gamma$ is large enough. Since $l \gg n^{1/4}$ and $P_1 \cdot P_2 \geq n^{-1/4}$, the theorem follows. □

## 2.1 The Algorithm

The following procedure describes one step of the algorithm. At the beginning of each step each node opens a channel, decides which messages to forward via `push` and `pull` (as described by the algorithm below). Then all nodes close all open channels. Recall that `push` sends from the calling node to the called node, and `pull` sends the message from the called node to the calling one. Nodes do not have to send messages over a channel opened by themselves (`push`) and they can also ignore nodes calling on them (`pull`).

To chose the neighbour to which the channel is directed we assume that every node $v$ stores a cyclic list $\ell_v$ with a random permutation of all its neighbours. Let $\ell_v(t)$ be the $t$-th entry in the list. Then we assume that $v$ opens a communication channel to $\ell_v(t)$ in step $t$ (we omit the division by $d$ for $t > d$). For $t' > t$ we define $L_v[t,t'] = \{\ell_v(t), \ell_v(t+1), \ldots \ell_v(t')\}$ as the set of nodes to which $v$ opened a channel in steps $t, \ldots t'$.

The following algorithm describes the behavior of the nodes w.r.t. one specific message $m$. We assume that every message that is forwarded contains the *age* of the message, which is defined as the difference between the current time step and the time step in which the message has been generated. The algorithm can also be applied in dynamic settings where nodes can generate new messages in every step. In that case the algorithm will be run for every message on every node, meaning that the algorithm has to decide which messages to send and which ones not, depending on the age of the message.

The age of a message determines which of the following phases of the algorithm applies to it. It also determines when the algorithm terminates for a message. If the age of a message is so large that none of the phases applies for the message, the algorithm stops forwarding the message. Note that, with a small probability, the message did notch reach all nodes of the graph. Hence, the algorithms terminates for a message after a fixed number of time steps but it does not deliver every message with probability one. Also note that we assume that nodes can combine all the messages that they would like to forward over a fixed edge to one large message. In our cost model, the cost of broadcasting a message is the total

number of transmissions in which the message is contained. One can regard our cost model as counting the total amount of data volume that is transmitted for every message.

We assume that $\rho > 40/\alpha^2$ is a (large) constant. A node is called *informed* if it got a copy of that message. We also assume $m$ is generated at time step 0 (i.e., at step $t$ the age of the message equals $t$), and that $t$ is the actual time. Recall that $\alpha$ is the expansion value of the graph.

**Phase 0:** $[\mathbf{age} \leq \lceil \rho \log n \rceil]$ The node $v$ which generates the message uses the channel to $\ell_v(t)$ for a `push` transmission in each step $t$ of this phase. No other node transmits the message in this phase.

**Phase 1:** $[\lceil \rho \log n \rceil + 1 \leq \mathbf{age} \leq 2 \cdot \lceil \rho \log n \rceil + \lceil 80/\alpha^2 \rceil]$ Nodes $v$ that received the message in Phase 0 use the first $\lceil 80/\alpha^2 \rceil$ steps of this phase to perform a `push` transmission to $\ell_v(t)$. If a node receives a message for the *first* time at time step $t' \in \{\lceil \rho \log n \rceil + 1, \ldots, 2 \cdot \lceil \rho \log n \rceil\}$, then the node will use the next $\lceil 80/\alpha^2 \rceil$ steps to perform a `push` transmission to $\ell_v(t)$.

**Phase 2:** $[2 \cdot \lceil \rho \log n \rceil + \lceil 80/\alpha^2 \rceil + 1 \leq \mathbf{age} \leq 2 \cdot \lceil \rho \log n \rceil + \lceil \rho \log d \rceil]$ *Every informed* node performs a `push` transmission to $\ell_v(t)$.

**Phase 3:** $[2 \cdot \lceil \rho \log n \rceil + \lceil \rho \log d \rceil + 1 \leq \mathbf{age} \leq 2 \cdot \lceil \rho \log n \rceil + 2 \cdot \lceil \rho \log d \rceil]$ *Every informed* node performs a `pull` transmission to every calling node in each step of this phase.

**Phase 4:** $[2 \cdot \lceil \rho \log n \rceil + 2 \cdot \lceil \rho \log d \rceil + 1 \leq \mathbf{age} \leq 2 \cdot \lceil \rho \log n \rceil + 2 \cdot \lceil \rho \log d \rceil + \lceil \rho \sqrt{\log n} \rceil]$ *Every informed* node performs a `pull` transmission to every calling node in each step of this phase.

**Phase 5:** $[2 \cdot \lceil \rho \log n \rceil + 2 \cdot \lceil \rho \log d \rceil + \lceil \rho \sqrt{\log n} \rceil + 1 \leq \mathbf{age} \leq 3 \cdot \lceil \rho \log n \rceil]$ Every node that receives the message *in Phase 4 or 5* performs a `pull` transmission to every calling node in each step of this phase.

The other informed nodes flip a coin and performs a `pull` transmission to every calling node with probability $1/\sqrt{\log n}$.

**Phase 6:** $[3 \cdot \lceil \rho \log n \rceil + 1 \leq \mathbf{age} \leq 3 \cdot \lceil \rho \log n \rceil + \lceil \rho \sqrt{\log n} \rceil]$ *Every informed* node performs a `pull` transmission to every calling node in each step of this phase.

In the first 3 phases the above algorithm performs `push` transmissions, in the remaining 4 phases it performs `pull` transmissions. Note that here is no algorithmic difference between Phase 3 and Phase 4. We introduce these two phases since they will be analysed separately. If a node decides to answer a `pull` request then it will answer all `pull` requests during that step. We say a node is *active* in a phase if it performs transmissions in that phase. The idea of the algorithm is as follows.

**`push` Phases.**

- In Phase 0 the node which generated the message performs a `push` transmission in every step. At the end of the phase $O(\log n)$ nodes are informed, w.h.p.
- In Phase 1 every informed node performs *a constant number* of `push` transmissions. After that we have w.h.p. $n/d$ informed nodes. The restriction to a constant number of transmissions per node helps to reduce the transmission number.

- The purpose of the third phase is to inform $n/2$ nodes. In this phase every informed node performs a `push` transmission in every step of the phase. Note that this phase consists only of $O(\log d)$ many steps.

The first 3 Phases are analysed in Observation 1, Lemma 1, and Lemma 2.

**`pull` Phases.**

- In every step of Phase 3 every informed node uses all incoming channels for `pull` transmission. At the end of the Phase 3 we have $n - n/d^3$ informed nodes, w.h.p (Lemma 3).
- In every step of Phase 4 every informed node uses all incoming channels for `pull` transmission.

  In Phase 5 the nodes that were informed during the last two phases become active. All remaining nodes will become active with a probability of $1/\sqrt{\log n}$ per step. (This helps to keep the number of transmissions low.) Every active nodes use all incoming channels for a `pull` transmission in every step.

  Phase 4 and Phase 5 are responsible to inform w.h.p. all uninformed nodes that have, in turn, many uninformed neighbours at the beginning of Phase 4. These two phases are analysed in Lemma 4.
- The remaining nodes are informed in Phase 6 where every informed node uses every incoming channel for a `pull` transmission in every step. This is shown in Lemma 5.

## 2.2 Analysis of the Algorithm

The analysis of the algorithm is more or less divided into the same phases as the algorithm. First we show (Lemma 1 and Lemma 2) that the algorithm informs w.h.p. at least $n/2$ nodes during the first $O(\log n)$ steps of Phase 1 to Phase 3 , using $O(n)$ message transmissions. More precisely, we show that (w.h.p.) in a constant number of steps the number of informed nodes increases by a constant factor, as long as the number of informed nodes is less than $n/2$.

As soon as the number of informed nodes is larger than $n/2$ the analysis becomes much more complicated. If we were only interested in the running time of our algorithm, then we could apply a backward analysis as in e.g. [15] to show that the algorithm completes broadcasting in $O(\log n)$ steps, w.h.p. However, this would result in a bound of $\Theta(n \log n)$ on the number of message transmissions. Since our goal is to significantly reduce the number of message transmission per node we need new analytical techniques for this case. Thus, we first analyse the distribution of edges in the set of uninformed nodes as well as the distribution of the so called cut edges separating informed and uninformed nodes from each other. To obtain the desired result, we design a new combinatorial technique that combines the information flow from informed to uninformed vertices with the distribution of the cut edges.

In our proofs we assume for simplicity that $\rho \log n$, $\rho \log d$, and $\rho\sqrt{\log n}$ are all integers. We assume that $d \geq f(n) \cdot \log^{3/2} n$ with $f : \mathbb{N} \to \mathbb{R}$ being a function such that $\lim_{n\to\infty} f(n) = \infty$. Note that whenever $d = O(\sqrt{\log n})$, the bound on the number of message transmissions can be trivially fulfilled, since each node can communicate with all its neighbors in $O(\sqrt{\log n})$ steps. However, the analysis cannot be extended to values between $\sqrt{\log n}$ and $\log^{3/2} n$ as

already mentioned in Section 1.2. Whenever we analyse a phase of our algorithm we assume that all earlier phases were successfull in the sense that they informed the right number of nodes. To get the failure probability of Algorithm RR one has to add up the failure probability of all 6 phases. We will use the following definitions.

- $I(t)$ is the set of informed nodes at the beginning of step $t$.
- Let $I^+(t) = I(t+1) \setminus I(t)$, that is, the nodes that get informed in step $t$.
- Let $\tau = t, t+1, \ldots t'$ be some consecutive steps of our algorithm. Then $I^+(\tau)$ is the set of nodes that get informed in steps $t, t+1, \ldots t'$ from one of the nodes of $I(t)$.
- $H(t)$ is the set of uninformed nodes $V \setminus I(t)$ at time $t$.
- $E(S, \overline{S})$ is the set of edges between $S$ and $\overline{S}$.
- $N(S, S')$ is the set of neighbours of $S$ that are in $S'$. Accordingly, $N(u, S')$ is the set of neighbours of $u \in V$ in $S'$.

### 2.2.1 Phase 0

Since $d \geq f(n) \cdot (\log n)^{3/2}$ it is easy to see that in Phase 0 the node, on which the message is generated, informs $\rho \log n$ different neighbours, which results in the following observation.

**Observation 1** *At the end of Phase 0 there are $\rho \log n$ informed nodes.*

### 2.2.2 Phase 1

Claim 1 below is used to analyse Phase 1. We divide Phase 1 into $k = (\rho \log n)/\ell + 2$ *time intervals* $\tau_1, \ldots, \tau_k$ of length $\ell = 40/\alpha^2$ each. Note that due to the definition of Phase 1, every node that is informed in time interval $\tau_i$ will perform `push` transmissions during the whole time interval $\tau_{i+1}$.

**Claim 1** *Let $\tau_1, \ldots, \tau_i, \ldots$ be the time intervals of Phase 1 and let $t_i$ be the beginning of $\tau_i$. Assume that there may be $o(|I(t_i)|)$ nodes in $I^+(\tau_{i-1})$ which do not send the message in the whole time interval $\tau_i$. Assume further that*

$$|I(t_i)| \leq \frac{n}{d} \quad \textit{and} \quad |I(t_i)| \geq \frac{8}{\alpha} \cdot |I(t_{i-1})|.$$

*Then with a probability of $1 - n^{-3}$ we have*

$$|I^+(\tau_i)| \geq \frac{8}{\alpha} \cdot |I(t_i)|.$$

**Proof:** To show this result we assume (as a worst case assumption) that the lists of the nodes are ordered by an adversary. Each node chooses at the beginning a random list position. To create the list, the adversary may have total knowledge about the topology of the network, but she cannot foresee any node's random choice w.r.t. the position selected at the beginning.

For the proof we can assume that there are at most $o(I(t_i) = o(|I^+(\tau_{i-1})|)$ nodes in $I^+(\tau_{i-1})$ which do not transmit in time interval $\tau_i$. We denote this subset of nodes by

$I^-(\tau_{i-1})$. Using the expansion properties of the graph and $H(t_i) \subset H(t_{i-1})$ we obtain that $|N(I^+(\tau_{i-1}) \setminus I^-(\tau_{i-1}), H(t_i))|$ is at least

$$\begin{aligned}
&|N(I(t_i), H(t_i))| - |N(I(t_{i-1}), H(t_i))| - |N(I^-(\tau_{i-1}), H(t_i))| \\
&\geq |N(I(t_i), H(t_i))| - |N(I(t_{i-1}), H(t_{i-1}))| - |N(I^-(\tau_{i-1}), H(t_i))| \\
&\geq \alpha \cdot d \cdot |I(t_i)| - d \cdot |I(t_{i-1})| - d \cdot o(|I(t_i)|) \\
&\geq \alpha \cdot d \cdot |I(t_i)| - d \cdot \frac{\alpha}{8} \cdot |I(t_i)| - o(d|I(t_i)|) \\
&= \frac{7}{8} \cdot \alpha \cdot d \cdot |I(t_i)|(1 - o(1)).
\end{aligned}$$

The last inequality holds due to the second precondition of the claim. Every node $v \in I^+(\tau_{i-1}) \setminus I^-(\tau_{i-1})$ performs $\ell = 40/\alpha^2$ push transmissions in time interval $\tau_i$. Now fix node $u \in N(I^+(\tau_{i-1}) \setminus I^-(\tau_{i-1}))$ and assume the node has $r$ neighbours in $I^+(\tau_{i-1}) \setminus I^-(\tau_{i-1})$. Then

$$\mathbf{Pr}\left[\, u \in I^+(\tau_i) \,\right] = 1 - \left(1 - \frac{40}{\alpha^2 d}\right)^r \geq 1 - \left(1 - \frac{40}{\alpha^2 d}\right) = \frac{40}{\alpha^2 d}.$$

By linearity of expectations we get

$$\mathbf{E}\left[\, |I^+(\tau_i)| \,\right] \geq \left(\frac{7}{8} \cdot \alpha \cdot d \cdot |I(t_i)|(1 - o(1))\right) \cdot \left(\frac{40}{\alpha^2 d}\right) \geq \frac{35}{\alpha} \cdot |I(t_i)|(1 - o(1)).$$

For $v \in I^+(\tau_{i-1})$, let $S_v = \{s_v^1, \ldots s_v^\ell\}$ be the random variables determining the choices of $v$, i.e. determining the nodes to which $v$ opens a channel in the interval $\tau_i$, and let

$$S = \bigcup_{v \in I^+(\tau_{i-1}) \setminus I^-(\tau_{i-1})} S_v.$$

Note that the choices in $S_v$ and $S_w$ are independent from each other for $v \neq w$. Since every $v \in I^+(\tau_{i-1})$ can only inform at most $40/\alpha^2$ in time interval $\tau_i$, $I^+(\tau_i)$ satisfies the $40/\alpha^2$-Lipschitz condition and the method of independent bounded differences [25] gives

$$\mathbf{Pr}\left[\, I^+(\tau_i) \leq \mathbf{E}\left[\, I^+(\tau_i) \,\right] - \lambda \,\right] \leq \exp\left(-\frac{\lambda^2}{2|I^+(\tau_i)|(40/\alpha^2)^2}\right).$$

With $\lambda = 27|I(t_i)|/\alpha$ we can conclude that

$$\mathbf{Pr}\left[\, I^+(\tau_i) \leq \frac{8}{\alpha} \cdot |I(t_i)| \,\right] \leq \exp\left(-\mathcal{O}(\log n)\right) \leq n^{-3},$$

since $|I(t_i)| \geq \rho \cdot \log n$ with a sufficiently large $\rho$. □

**Lemma 1** *With a probability of $1 - n^{-2}$ at least $n/d$ nodes are informed at the end of Phase 1.*

**Proof:** Let $t$ be the beginning of Phase 1. To show the result we will prove that $|I^+(t + \rho \log n))| \geq n/d$ with a probability of at least $1 - n^{-2}$.

Recall that we divided Phase 1 into $k = (\rho \log n)/\ell + 2$ time intervals $\tau_1, \ldots, \tau_k$ of length $\ell = 40/\alpha^2$ each. We assume that $t_i$ is the first step of time interval $\tau_i$. Let $I^+(\tau_i)$ be the random variable that counts the number of nodes that get informed in time interval $\tau_i$. Note that all nodes in $I^+(\tau_i)$ will perform push transmissions in every step of interval $\tau_{i+1}$ (for $i < k-1$). Since $\rho > 4$ we can assume that we have already $4 \log n$ informed nodes at the beginning of Phase 1. None of these nodes has transmitted the message yet. The proof of the lemma is based on Claim 1. The claim shows that, with probability at least $1 - n^{-3}$,

$$|I(t_{i+1})| \geq (8/\alpha) \cdot |I(t_i)|.$$

By repeatedly applying Claim 1 we conclude that if $\rho$ is large enough, then after $k$ many time intervals at least $n/d$ nodes are informed with a probability of at least $1 - n^{-2}$. $\square$

### 2.2.3 Phase 2

Next we consider Phase 2 where informed nodes perform `push` for roughly $\rho \log d$ many steps. Note that the statement of this lemma holds conditioned on the event that Phase 1 was successful.

**Lemma 2** *With a probability of $1 - n^{-2}$ at least $(n/2)$ nodes are informed at the end of Phase 2.*

**Proof:** Notice that the RR-model is more powerful than the RANDOM[1]-model, since an edge used for communication once cannot be used again to inform a node. Thus, we assume in this proof that the communication relies on the RANDOM[1]-model. Let $0 < \delta < 1$ be a suitably chosen constant and let

$$\Phi_{I(t)} := \frac{|E(I(t), H(t)|}{|I(t)| \cdot d}$$

be the edge expansion of $I(t)$. Due to our definition of Edge-Node Expanders we have $\Phi_{I(t)} \geq \alpha$ for $|I(t)| \leq n/2$. In order to prove the lemma we will first show that w.h.p.

$$|I^+(t)| \geq |I(t)| \cdot (1 - \delta) \cdot \Phi_{I(t)} \cdot \frac{1}{2}.$$

Let $\{v_1, \ldots, v_{H(t)}\}$ be the set of uninformed nodes in step $t$. Then,

$$\sum_{i=1}^{|H(t)|} |N(v_i, I(t))| = |E(I(t), H(t))| = \Phi_{I(t)} \cdot d \cdot |I(t)|,$$

and for $v \in H(t)$ it holds that

$$\mathbf{Pr}\left[\, v \in I^+(t) \,\right] = 1 - ((d-1)/d)^{|N(v, I(t))|}.$$

By linearity of expectations we have

$$\mathbf{E}\left[\, |I^+(t)| \,\right] = \sum_{v \in H(t)} \mathbf{Pr}\left[\, v \in I^+(t) \,\right].$$

It follows that

$$\mathbf{E}\left[\,|I^+(t)|\,\right] = \sum_{v\in H(t)} 1-\Big(\frac{d-1}{d}\Big)^{|N(v,I(t))|} \geq \sum_{v\in H(t)} 1-e^{\left(-\frac{|N(v,I(t))|}{d}\right)}.$$

Using the fact that $\exp(-x) \leq 1-\frac{x}{2}$ for $-1.5 < x < 0$ we obtain

$$\mathbf{E}\left[\,|I^+(t)|\,\right] \geq \sum_{v\in H(t)} \frac{|N(v,I(t))|}{2d} = \frac{\Phi_{I(t)}\cdot|I(t)|}{2}.$$

We can apply the method of independent bounded differences as in the proof of Claim 1, and obtain that

$$\mathbf{Pr}\left[\,|I^+(t)| \leq (1-\delta)\cdot\Phi_{I(t)}\cdot|I(t)|\cdot\big(1-\frac{1}{e}\big)\,\right] \leq \exp(-\Omega(\Phi_{I(t)}\cdot|I(t)|)).$$

So far we have shown that in one step of Phase 2 we have w.h.p.

$$|I^+(t)| \geq (1-\delta)\cdot\alpha\cdot(1-e^{-1})\cdot|I(t)|.$$

At the beginning of Phase 2 we have $I(t) \geq n/d$. Hence, after $\rho\log d$ additional steps we have $|I(t)| \geq n/2$ for $\rho$ large enough. □

### 2.2.4 Phase 3

**Lemma 3** *With a probability of $1-n^{-2}$ at least $n-n/d^3$ nodes are informed at the end of Phase 3.*

**Proof:** To prove the lemma, we use similar techniques as in the proof of Lemma 2. Let $t$ be the beginning of Phase 3. We assume that $|I(t)| > n/2$. We know that

$$|E(H(t'),I(t'))| \geq \alpha d\cdot|H(t')|,$$

where $t' \in \{t,\ldots,t+\rho\log d\}$. Let $H'(t')$ be the set of nodes $v\in H(t')$ with

$$|N(v,I(t'))| \geq \frac{\alpha d}{2}.$$

We can assume that there are no more than $(1-\alpha/2)\cdot|H(t')|$ nodes in $H(t')\setminus H'(t)$, since otherwise

$$|E(H(t'),I(t'))| \leq \Big(1-\frac{\alpha}{2}\Big)\cdot\frac{\alpha d}{2}\cdot|H(t')| + \frac{\alpha}{2}\cdot|H(t')|d < \alpha d\cdot|H(t')|$$

which contradicts our assumption that the graph is an $\alpha$-expander.

Fix $t'\in\{t,\ldots,t+\rho\log d\}$. Since every node $v\in H'(t')$ opens a channel in step $t'$, each of these nodes receives the message with a probability of at least $\alpha/2$, independently of the other nodes in $H'(t')$. Thus, applying simple Chernoff bounds [6], we can conclude that with probability $1-o(n^{-2})$

$$I^+(t') \geq |H'(t')|\cdot\alpha\cdot\frac{1-o(1)}{2}.$$

Since

$$|H'(t')| \geq |H(t')| \cdot \frac{\alpha}{2}$$

we have

$$|I^+(t')| \geq |H(t')| \cdot \alpha^2 \cdot \frac{1-o(1)}{4}.$$

Hence, $|H(t+\rho \log d)| \leq n/d^3$ for some properly chosen constant $\rho$. □

### 2.2.5 Phase 4 and Phase 5

Now we focus on Phases 4 and 5. Assume that $t$ is the beginning of Phase 4 and that there are at least $n - n/d^3$ informed nodes at that time. Recall that

- in Phase 4 (age $t$ to $t + \rho\sqrt{\log n}$) all informed nodes perform pull transmissions.
- In Phase 5 (age $t + \rho\sqrt{(\log n)} + 1$ to $3\rho \log n$)
  - every node that was informed in Phase 1-3 performs pull transmissions with probability $1/\sqrt{\log n}$,
  - and every node that was informed in Phase 4 or Phase 5 performs pull transmissions with probability 1.

**Lemma 4** *Let $t$ be the beginning of Phase 4. A node $v \in H(t)$ with $|N(v, I(t))| \leq d/2$ receives the message with probability $1 - n^{-2}$ by the end of Phase 5.*

**Proof:** To prove the lemma we use a so called *backward analysis*, for which we need some new definitions. We consider a node $v$ and the end of Phase 5, and in terms of Claim 5, this node is assumed to be colored blue. The node $\ell_v(\kappa)$ is called the $\kappa$-*active* neighbour of $v$, and the nodes $L_v[\kappa_1, \kappa_2]$ are called $(\kappa_1, \kappa_2)$-*active* neighbours of $v$. A node $w$ is called $\kappa$-*predecessor* of $v$ if there exists $k \leq \kappa$, some nodes $w_1, \ldots, w_k$ and time steps $t_0 < t_1 < \cdots < t_k \leq \kappa$ such that $w$ is the $t_0$-active neighbour of $w_1$, node $w_i$ is the $t_i$-active neighbour of $w_{i+1}$ $(1 \leq i < k)$, and $w_k$ the $t_k$-active neighbour of $v$. This means that $v$ is connected to $w$ by a path consisting of edges that were active in the time interval $[t_0, \kappa]$.

Note that, if $w$ is a $\kappa$-predecessor of $v$ and $w$ is informed at step $t_0$, then $v$ becomes informed at time $\kappa$. For different choices of $k$ and $t_0, t_1, \ldots, t_k$ one might regard $v$ as being a node of a tree consisting of $\kappa$-predecessors. If one of the nodes in the tree is informed in the time step that corresponds to the node, $v$ will get the message via the corresponding path in the tree. Therefore, the analysis is performed backwards in time by considering the predecessors, step by step. First we consider the predecessors to which $v$ opens a channel close to the end of Phase 5. Then we consider the nodes to which these neighbors of $v$ open channels in some steps earlier in time, etc... Finally we show that at some point in time (which occurs after the end of Phase 3), the number of predecessors must exceed the total number of uninformed nodes at the end of Phase 3. This would imply that the message reaches node $v$ by the end of Phase 5.

We define

$$T_0 = t + \rho\sqrt{\log n},\ T_1 = t + (\rho \cdot \log n)/4,\ T_2 = t + (\rho \cdot \log n)/2, \text{ and } T_3 = t + \rho \log n.$$

In the following we will consider time intervals $[T_i, T_{i+1}]$ with $i \in 0, 1, 2$. We assume that $[T_i, T_{i+1}]$ down not contain step $T_i$.

First show the following claim stating that every uninformed node with many uninformed neighbours has at least $(\rho \cdot \log n)/8$ many $(T_2, T_3)$-active neighbours in $H(t)$.

**Claim 2** *Let $v$ be a node in $H(t)$ with $|N(v, I(t))| \leq d/2$. With probability $1 - n^{-4}$*

$$H(t) \cap L_v[T_2, T_3] \geq \frac{\rho \log n}{8}.$$

**Proof:** Fix $\kappa \in [T_2, T_3]$. Let $p'$ be the probability that $\ell_v(\kappa) \in H(t)$. By time step $\kappa$ vertex $v$ used at most $\rho \log d + \rho \log n \leq 2\rho \log n$ many edges for pull transmissions ($\rho \log d$ edges in Phase 3 and at most $\rho \log n$ edges in Phase 4 and Phase 5. Note that the algorithm only performs push transmissions in the first three phases). Since $v$ has at least $d/2$ neighbors in $H(t)$ we have

$$p' \geq (d/2 - 2\rho \log n)/d.$$

With $d = f(n) \cdot (\log n)^{3/2}$ we get $p \geq 3/8$. The expected number of nodes in $I(t) \cap L_v[T_2, T_3]$ is at most $5\rho \log n/16$ and we can apply Chernoff bounds [6, 20] to conclude that whenever $\rho$ is large enough

$$\mathbf{Pr}\left[\, I(t) \cap L_v[T_2, T_3] \geq (7\rho \log n)/8 \,\right] \leq n^{-4}.$$

This finishes the proof of Claim 2. □

Applying the claim, we can assume for the rest of the proof that $v$ has at least $(\rho \log n)/8$ many $(T_2, T_3)$-active neighbours in $H(t)$. We say a node $v \in H(t)$ is $I(t)$-*good* if it has at least $d/2$ of its neighbours in $I(t)$ (meaning $|N(v, I(t))| \geq d/2$). Otherwise the node is called $I(t)$-*bad*.

In the following we show that every node that is $I(t)$-bad ($|N(v, I(t))| \leq d/2$) will receive the message from a node in $I(t)$ in Phase 4 either directly via one of its $I(t)$-good neighbours, or via a longer path to a node in $I(t)$ consisting of nodes which are in $H(t)$. Note that this shows Lemma 4 since it only states that nodes $v \in H(t)$ with $|N(v, I(t))| \leq d/2$ will be informed by the end of Phase 5.

To show that every node that is $I(t)$-bad will receive the message in Phase 4. We consider two cases.

**Case 1: In $L_v[T_2, T_3] \cap H(t)$ are at least $\rho \cdot \sqrt{\log n}$ $I(t)$-good nodes.** Let $U$ be the set of $I(t)$-good neighbours of $L_v[T_2, T_3]$. Note that $v$ receives the message in $[T_2, T_3]$ if there exists a node $u \in U$ that received the message in Phase 4. This holds since all nodes which receive the message for the first time in Phase 4 or Phase 5 respond to every pull request. The probability that a node $w \in U$ is still uninformed at the end of Phase 4 (step $t + \rho\sqrt{\log n}$) is at most $(3/8)^{-\rho\sqrt{\log n}}$. This holds since for every $\kappa \in [t, t + \rho\sqrt{\log n}]$ the probability that $\ell_w(\kappa) \in I(t)$ is at least

$$(d/2 - \rho \log d - \rho\sqrt{\log n})/d \geq 5/8$$

($\rho \log d$ edges in Phase 3 and at most $\rho\sqrt{\log n}$ edges in Phase 4 might have already been used for pull requests, and $d \geq f(n) \cdot (\log n)^{3/2}$). For $\rho > 4$, all nodes of $U$ are still uninformed at time $t + \rho\sqrt{\log n}$ with a probability of at most $(3/8)^{-\rho^2(\sqrt{\log n})^2} = n^{-4}$.

**Case 2: In $L_v[T_2, T_3] \cap H(t)$ are fewer than $\rho \cdot \sqrt{\log n}$ $I(t)$-good nodes.** Due to our assumption that $v$ has at least $\rho \log n/8$ many $(T_2, T_3)$-active neighbours in $H(t)$, at least $(\rho/8) \cdot \log n - \rho\sqrt{\log n}$ of $v$'s $(T_2, T_3)$-active neighbours (in $H(t)$) are $I(t)$-bad (see Claim 2). Let $U = w_1, \ldots w_k$ be an arbitrary subset of size $k = \sqrt{f(n) \log n}$ of the $I(t)$-bad neighbours of $v$. Now we step back in time and consider the time interval $[T_1, T_2]$.

**Claim 3** *Let*

$$U' = \bigcup_{w \in U} L_w(T_1, T_2).$$

*Then with a probability of* $1 - o(n^{-3})$

$$|U' \cap H(t)| = \Omega(\sqrt{f(n)}(\log n)^{3/2}).$$

**Proof:** To bound the size of $U' \cap H(t)$ we consider one node of $U$ after the other. We define $U'_0 = \emptyset$ and for $1 \leq i \leq k$

$$U'_i = \begin{cases} U'_{i-1} & \text{if } |U'_{i-1}| \geq \sqrt{f(n)} \cdot (\log n)^{3/2} \\ \bigcup_{j=1}^{i} L_{w_j}[T_1, T_2] \cap H(t) & \text{otherwise.} \end{cases}$$

Finally, we define $U' = U_k$. We calculate the probability that the construction ends with $|U'| < \sqrt{f(n)} \cdot (\log n)^{3/2}$.

Assume $|U'_{i-1}| < \sqrt{f(n)} \cdot (\log n)^{3/2}$ for some $i$. Then, for $w_i \in U$ and $\ell_j \in L_{w_i}[T_1, T_2]$ we have

$$\begin{aligned} \mathbf{Pr}\left[\, \ell_j \in U'_{i-1} \,\right] &\leq \frac{\sqrt{f(n)} \cdot (\log n)^{3/2} + \cdot\rho \cdot \log n/2}{d - \rho \log n} \\ &\leq \frac{\sqrt{f(n)} \cdot (\log n)^{3/2} + (\rho \cdot \log n)/2}{f(n) \cdot (\log n)^{3/2} - \rho \log n} \leq \frac{2}{\sqrt{f(n)}}. \end{aligned}$$

Here, the term $\sqrt{f(n)} \cdot (\log n)^{3/2}$ stands for the maximum size of $U'_{i-1}$, and the additional term $\rho \cdot \log n/2$ for $|L_{w_i}[T_0, T_2]|$. The term $\rho \log n$ in the denominator represents an upper bound on the number of neighbours chosen already in Phases 3, 4, and 5. Notice that if $d$ was $O(\log^{3/2} n)$, then we could not obtain a proper value for the probability above. The value we obtained will be applied in Claim 4 below.

By time step $\kappa \in [T_1, T_2]$ $w_i$ used at most

$$\rho \log d + (\rho \log n)/2 < \rho \log n$$

many edges for pull transmissions. Since we assumed that $d \geq f(n) \cdot (\log n)^{3/2}$, we can argue that

$$\mathbf{Pr}\left[\, \ell_j \in I(t) \,\right] \leq \frac{d/2}{d - (\rho \log n)} \leq \frac{5}{8}.$$

Hence,

$$\mathbf{Pr}\left[\, \ell_j \in I(t) \cup U'_{i-1} \,\right] \leq \frac{2}{\sqrt{f(n)}} + \frac{5}{8} \leq \frac{3}{4},$$

regardless of the sets $U_{i-1}$ and $I(t)$. The expected number of nodes in

$$L_{w_i}[T_1, T_2] \cap (I(t) \cup U'_{i-1})$$

is at most

$$3 \cdot (T_2 - T_1)/4 \geq (3\rho \log n)/16.$$

We can apply Chernoff bounds [6] to conclude that with probability $1 - O(n^{-4}/\sqrt{f(n) \log n})$ we have

$$L_{w_i}[T_1, T_2] \cap (I(t) \cup U'_{i-1}) \leq 7\rho \log n/32.$$

Thus,

$$|U'_i| \geq |U'_{i-1}| + (\rho \log n)/4 - 7\rho \log n/32 \geq |U'_{i-1}| + (\rho \log n)/32.$$

Since $k = \sqrt{f(n) \cdot \log n}$, with a probability of $1 - o(n^{-3})$ we have

$$|U' \cap H(t)| \geq k \cdot (\rho \log n)/32 \geq (\rho \sqrt{f(n)} \cdot (\log n)^{3/2})/32.$$

This finishes the proof of Claim 3.

□

Now we are back to proving Case 2. Next we show that the set of predecessors of any node $w \in U' \cap H(t)$ grows the further we go backwards from $T_1$ to $T_0$. We define $\ell = 40/\alpha$ and divide the time interval $[T_0, T_1]$ into $(T_1 - T_0)/\ell$ rounds of equal length. For $0 \leq i \leq k' - 1$ we define

$$\tilde{T}_i = [T_1 - i \cdot \ell, T_1 - (i+1) \cdot \ell].$$

Let

$$U_0^H = \bigcup_{w \in U' \cap H(t)} L_w[\tilde{T}_0] \cap H(t) \text{ and } U_0^I = \bigcup_{w \in U' \cap H(t)} L_w[\tilde{T}_0] \cap I(t)$$

be the corresponding ($\tilde{T}_0$-active) set of uninformed and informed neighbours of $w \in U' \cap H(t)$, respectively. For $1 \leq i \leq k' - 1$ we define

$$U_i^H = \bigcup_{w \in U_{i-1}^H} L_w[\tilde{T}_i] \cap H(t) \text{ and } U_i^I = \bigcup_{w \in U_{i-1}^H} L_w[\tilde{T}_i] \cap I(t).$$

For every node $\overline{w}_i \in U_i^I$ there is a path $P = (\overline{w_i}, \ldots, \overline{w}_0, w', w)$ to a node $w \in U$, where $\overline{w}_{i-1}, \ldots, \overline{w}_0$, $w', w \in H(t)$, and $\overline{w}_j$ $(i \leq j \leq 0)$ is an active neighbour of its predecessor in $\tilde{T}_{j+1}$. Hence, together $U_i^H$ and $U_i^I$ can be regarded as the $i$-th level of a tree routed in $w$ and, consequently (since $w$ is an active neighbour of $v$), also in $v$. Nodes in $H(t)$ are inner nodes of the tree, and the leaves are nodes in $I(t)$ or nodes in $H(t)$ on level $k' - 1$. Note that nodes can occur several times on several different levels of the tree.

In the following we argue that the amount of different leaves (informed nodes) in the tree is w.h.p. at least $\rho \cdot (\log n)^{3/2}$. Then we will show that the informed leaves are sufficient to inform node $v$. Let

$$U_{0 \to i}^H = \left[ \bigcup_{j=0}^{i} U_j^H \right] \text{ and } U_{0 \to i}^I = \left[ \bigcup_{j=0}^{i} U_j^I \right].$$

The following claim shows that w.h.p. there exists a time interval $i$ with $|U_{0 \to i}^I| \geq \rho \cdot (\log n)^{3/2}$.

**Claim 4** *With a probability of $1 - n^{-3}$ there exists $i \in 1, \ldots, (T_1 - T_0)/\ell - 1$ with $U_{0 \to i}^I > \rho \cdot (\log n)^{3/2}$.*

**Proof:** Define

$$\Delta U_{0 \to j}^H = U_{0 \to j}^H \setminus U_{0 \to j-1}^H, \text{ and } \Delta U_{0 \to j}^I = U_{0 \to j}^I \setminus U_{0 \to j-1}^I.$$

According to Claim 5 (see below), as long as

$$|U_j^I| \leq \rho (\log n)^{3/2},$$

and

$$|\Delta U_j^H| = \Omega(\sqrt{f(n)}(\log n)^{3/2}),$$

it holds with a probability of $1 - n^{-3}$ that

$$\Delta U_{1\to j+1}^I \cup \Delta U_{1\to j+1}^H \geq \left(\frac{8}{\alpha}\right) \cdot |U_j^H|.$$

Here, $\Delta U_{1\to j+1}^I \cup \Delta U_{1\to j+1}^H$ represents the set of newly colored blue nodes in terms of Claim 5. With a probability of $1 - n^{-3}$ we have $|H(t)| < n/d^3$. Due to the exponential growth of the set $\Delta U_{1\to j+1}^I \cup \Delta U_{1\to j+1}^H$ there must be some $i = O(\log n)$ such that with probability $1 - n^{-3}$

$$U_{0\to i}^I > \rho \cdot (\log n)^{3/2}.$$

□

Now let $i \leq k' - 1$ be the value with $U_{0\to i}^I > \rho \cdot (\log n)^{3/2}$. Using the claim we can easily argue that $v$ will get the message over a path starting at one of the nodes in $U_{0\to i}^I$. Recall that every node $u \in U_{0\to i}^I$ answers pull requests with a probability of $1/\sqrt{\log n}$. Since for such a $u$ there is a path $P = (u, \overline{w_s}, \ldots, \overline{w}_0, w', w, v)$ consisting of nodes in $H(t)$, all the nodes on this path answer every pull request. Hence, if there exists a node $u \in U_{0\to i}^I$ which answers the pull request at the time $u$ was the active neighbour of $\overline{w}_s$, $v$ will be informed. This happens with a probability of

$$1 - \left(1 - \frac{1}{\sqrt{\log n}}\right)^{\rho \cdot (\log n)^{3/2}} \geq 1 - e^{-\rho \log n} \geq 1 - n^{-3}, \tag{1}$$

Note that in the above equation we need that $U_{0\to i}^I > \rho \cdot (\log n)^{3/2}$.

□

In Claim 5 we follow the "spread" of uninformed nodes back in time. We start with an arbitrary but fixed node $v$ and time step $t$. We assume that $v$ is colored *blue*. If a node $w'$ is blue at $t'' \leq t$, and $w'$ opens a communication channel to $w''$ in step $t'' - 1$, then $w''$ becomes blue in step $t'' - 1$.

Let $\tau_1', \tau_2', \ldots, \tau_k' \leq t$ be $k$ consecutive time intervals of length $40/\alpha^2$, which go back in time and let $t_i$ be the end of $\tau_i$. We assume that $\tau_i'$ starts directly after $\tau_{i+1}'$ (on the real time axis). Furthermore, we assume that at the end of $\tau_1'$ there are at least $\sqrt{f(n)}\log^{3/2} n$ blue nodes. Let $I'(t)$ be the set of blue nodes at time $t$, and let $I'^{+}(t)$ and $I'^{+}(\tau_i)$ be defined accordingly (note that $|I'(t)| < |I'(t')|$ for $t > t'$ since we consider the spreading process backwards in time).

**Claim 5** *Assume that are $o(|I'(t_i)|)$ nodes in $I'^{+}(\tau_{i-1})$ which do not colour neighbours blue in the whole time interval $\tau_i'$. Assume further that*

$$|I'(t_i)| \leq \frac{n}{d} \quad \textit{and} \quad |I'(t_i)| \geq \frac{8}{\alpha} \cdot |I'(t_{i-1})|.$$

*Then with a probability of $1 - n^{-3}$ we have*

$$|I'^{+}(\tau_i)| \geq \frac{8}{\alpha} \cdot |I'(t_i)|.$$

The proof of the claim is the same as the proof of Claim 1, we only have to replace $I$ with $I'$ and $\tau$ with $\tau'$.

### 2.2.6 Phase 6

**Lemma 5** *At the end of Phase 6 every node is informed with probability* $1 - n^{-2}$.

**Proof:** We have shown that every node which has fewer than $d/2$ neighbours in $I(2(\rho \log n + \rho \log d))$ (beginning of Phase 4), is informed at the end of Phase 5. Thus, the only nodes which are still uninformed have at least $d/2$ neighbours in $I(3\rho \log n)$.

We know that in time interval

$$[3\rho \log n + \rho\sqrt{\log n}/2 + 1, 3\rho \log n + \rho\sqrt{\log n}]$$

every uniformed node $v$ contacts $\rho\sqrt{\log n}/2$ nodes. If one of these nodes was informed at time $3\rho \log n$, then $v$ becomes informed as well.

If none of these nodes is informed at time $3\rho \log n$, then all of them have at least $d/2$ neighbours in $I(3\rho \log n)$. Consider now a fixed node $w$ among these nodes. $w$ remains uninformed in time interval $[3\rho \log n + 1, 3\rho \log n + \rho\sqrt{\log n}/2]$ with probability at most $2^{-\rho\sqrt{\log n}/2}$, independently of the others. Thus, the probability that all of these nodes are still uninformed at time $3\rho \log n + \rho\sqrt{\log n}/2$ is

$$2^{-\rho^2 \log n/4} \leq n^{-4}$$

for any $\rho > 4$. This implies that $v$ becomes informed in Phase 6 with probability $1 - o(n^{-3})$. Then, the lemma follows by applying the union bound over all nodes which were uninformed at the end of Phase 5. □

By summarizing the results of the lemmas of this section we obtain the following theorem.

**Theorem 2** *Assume that* $f : \mathbb{N} \to \mathbb{R}$ *is a function such that* $\lim_{n\to\infty} f(n) = \infty$ *and* $d \in \{f(n) \cdot (\log n)^{3/2}, 2^{o(\sqrt{\log n})}\}$.

*For every Edge-Node Expander $G$ with degree $d$ our algorithm broadcasts a message in $G$ in time $O(\log n)$ by using $O(n\sqrt{\log n})$ transmissions, w.h.p.*

## 2.3 Lower Bound

In this section we show that the result of Theorem 2 is tight up to a $\sqrt{\log \log n}$ factor. For the following bound we also assume the *oblivious communication model*. In this model, a node's decision whether to transmit in a fixed step can depend on the age of the message and on any information the node might have aquired before the current step. However, the node's decision is not influenced by the ID's of the nodes at the other end of a currently open channel. We can assume that the nodes decide if they want to transmit in a fixed step or not *before* the channels become opened. Furthermore, the algorithms is not allowed to depend on the structure of the graph, i.e., we may assume that the graph is constructed by some random process, and the algorithm is set before the graph is constructed.

**Theorem 3** *Assume $A$ is a (Monte-Carlo) broadcast algorithm in the oblivious model with runtime $O(\log n)$ in the* RR *communication model. There exists a family of Edge-Node Expanders for which $A$ needs*

$$\Omega\left(n \cdot \sqrt{\frac{\log n}{\log \log n}}\right)$$

*message transmissions, w.h.p.*

**Proof:** We assume that $c \log n$ is the runtime of the algorithm. Let $G' = (V', E')$ be an $\alpha$-Edge-Node expander (i.e., a corresponding random graph) of size $n$ and with degree $\log^2 n - 1$. Define

$$r = \sqrt{\frac{\log n}{\log \log n}}.$$

For

$$\ell = \frac{n}{(r+1) \cdot (\log^2 n - r)}$$

and $i \in \{1, \ldots, l\}$, let $K_i$ be a clique (complete graph) of size $r$. To construct $G$ we connect the nodes of the cliques to the nodes of $G'$ such that every $K_i$ $(1 \leq i \leq \ell)$ has $\log^2 n - r + 1$ different neighbours in $G'$, and that each node of $G'$ has exactly one neighbour in $K_1, \ldots K_\ell$ chosen randomly. That is, each node in $G'$ has at any time step the same probability to be connected to a node in a completely uninformed clique unless a (message) communication between these nodes has been performed before. The edges between nodes in $G'$ are independent of $G'$-clique edges. Due to our construction, the degree of every node is $d = \log^2 n$, and the graph is still an Edge-Node Expander. In the following $G'$ and will be called the *original graph*, with *original nodes* and *original edges*. The nodes in the cliques are called *clique nodes*, and edges between two node of the same clique will be called *inner edges*. Edges connecting a clique node to $G'$ are called *external edges*.

Now we assume that there is an algorithm $A$, which has running time $c \log n$ and produces at most $\epsilon \cdot nr$ message transmissions with $\epsilon < 1/17$. We will show that, with a good probability, $A$ will not be able to inform all clique nodes of the graph. Since $A$ sends at most $\epsilon \cdot nr$ message, a simple pigeonhole argument shows that there exists at least $l/2$ cliques where at least half of the nodes will have a neighbourhood in $G'$ which (altogether) sends at most $4\epsilon \cdot dr$ messages. That is, there are $l/2$ complete graphs $K_{i_1}, \ldots, K_{i_{l/2}}$ such that in any $K_{i_j}$ at least $\frac{1}{2} \cdot r$ nodes belong to a subset $K'_{i_j}$ with the following properties:

- $|K'_{i_j}| \geq \frac{|K_{i_j}|}{2} = \frac{r}{2}$, and
- For each $v \in K'_{i_j}$ the nodes in $N(v, V')$ produce altogether at most $4\epsilon \cdot dr$ message transmissions (`push` or `pull`).

In the rest of the proof the nodes belonging to one of the $K'_{i_j}$ sets are called *good*, and the nodes of $K_{i_j} \setminus K'_{i_j}$ are called *bad*. Our goal is now to show that not all good nodes will get the message with a good probability. Let us fix a good node $v \in K_{i_j}$. We divide the steps into *dangerous* and *safe* steps now. In a dangerous step at least $d - i - 2r$ nodes of $N(v, V')$ send a `pull` message, and in a safe step fewer than $d - i - 2r$ nodes of $N(v, V')$ send a `pull` message. Our goal is now to show that there is a good probability that $v$ will communicate with good clique nodes in dangerous steps, and with non-sending nodes in $V'$ in safe steps. Since we assume the RR communication model, nodes can not (in $c \log n$ steps) communicate with the same node twice. Since $v$ has only $r/2$ good clique neighbours, we will try to "save" these nodes for dangerous steps and prevent $v$ from choosing inner edges during safe steps. In the following $x_i$ will be the set of nodes that $v$ should "avoid" in step $i$. For $0 \leq i \leq c \log n - 1$ we define $x_i$ as follows:

- Dangerous step $i$: $x_i$ is the number of nodes in $N(v, V')$ that perform `pull` in step $i$, together with the bad nodes of $K_{i_j}$, minus the nodes of these sets which have been called by $v$ before step $i$.

- Safe step $i$: $x_i$ is the number of nodes in $N(v, V')$ that perform `pull` in step $i$, together with the nodes of $K_{i_j}$, minus the nodes of these sets which have been called by $v$ before step $i$.

In a safe step we have

$$x_i \leq (d - i - 2r) + r = d - i - r.$$

Of course, in dangerous steps

$$x_i \leq (d - 2r) + (|K_{i_j}| - |K'_{i_j}|) \leq d - 1.5r.$$

To get a better bound on $x_i$ in dangerous steps we calculate the number of dangerous steps first. The total number of message transmissions performed by the nodes of $N(v, V')$ is at most $4\epsilon \cdot dr$. With $\epsilon < 1/17$, the total number of dangerous steps is at most

$$\frac{4\epsilon \cdot dr}{d - 1 - 2r} \leq \frac{17}{4}\epsilon \cdot r - 1 \leq \frac{r}{4} - 1.$$

Define $x_i(\text{good})$ as the number of good neighbours to which $v$ already opened a channel by step $i$. Since $v$ is only "allowed" to open a channel to good neighbours in dangerous steps $x_i(\text{good}) \leq r/4 - 1$. Now we can get that

$$x_i + i \leq d - r + 1 - (i - r/4 + 1) + r/2 + i = d - r/4.$$

Now we are ready to lower bound the probability that $v$ is NOT informed after $c \log n$ time steps. As a worst case assumption we assume in the following that $|K'_{i_j}| = |K_{i_j}|/2$. Let us consider pull transmissions first and define $P_1$ as the probability that $v$ does not get the message via a pull transmission. To bound this probability we will use the probability that, for $0 \leq i \leq c \log n - 1$, $v$ does not get the message via a pull transmission in step $i$ from one of the nodes in the set $x_i$. This means that $v$ is not allowed to open a channel to any node in $x_i$ in step $i$. Hence,

$$P_1 \geq \left(1 - \frac{x_0}{d}\right) \cdot \left(1 - \frac{x_1}{d-1}\right) \cdots \left(1 - \frac{x_{c \log n - 1}}{d - c \log n + 1}\right).$$

In order to derive a proper lower bound on $P_1$, we consider the following case analysis.

**Case 1:** $x_i < d/4$**.** Then $i$ is a safe step. Since $c \log n \ll d/4$ we have

$$1 - \frac{x_i}{d-i} \geq 1 - \frac{2x_i}{d} \geq \left(1 - \frac{1}{d}\right)^{4x_i}.$$

**Case 2:** $x_i \geq d/4$**.** In this case step $i$ can be either dangerous or safe. Since $x_i + i \leq d - r/4$

$$1 - \frac{x_i}{d-i} \geq 1 - \frac{x_i + i}{d} \geq 1 - \frac{d - r/4}{d} = \frac{r}{4d}.$$

Next we calculate the number of steps in which Case 2 applies. The total number of message transmissions performed by the nodes of $N(v, V')$ is at most $4\epsilon \cdot dr$. From the $x_i$

nodes there are at least $x_i - r$ nodes in $N(v, V')$. For $n$ large enough the number of steps with $x_i \geq d/4$ is at most

$$\frac{4\epsilon \cdot dr}{d/4 - r} \leq 17\epsilon \cdot r.$$

Let $I_1$ be the set of indices $0 \leq i \leq c \log n - 1$ with $x_i < d/4$ and let $I_2$ be the set of indices $i$ with $x_i \geq d/4$. Then we get with $r = \sqrt{\log n / \log\log n}$, $d = \log^2 n$, and

$$\sum_{i \in I_1} x_i \leq \sum_{i=0}^{c \log n - 1} x_i \leq 4\epsilon \cdot dr + c \log n \cdot r \leq 5\epsilon \cdot dr,$$

$$\begin{aligned} P_1 &\geq \prod_{i \in I_1} \left(1 - \frac{1}{d}\right)^{4x_i} \cdot \prod_{i \in I_2} \left(\frac{r}{4d}\right) \\ &\geq \left(1 - \frac{1}{d}\right)^{4 \sum_{i \in I_1} x_i} \cdot \left(\frac{\sqrt{\log n / \log\log n}}{4d}\right)^{17\epsilon \sqrt{\frac{\log n}{\log\log n}}} \\ &\geq \left(1 - \frac{1}{d}\right)^{20\epsilon d \sqrt{\frac{\log n}{\log\log n}}} \cdot \left(\frac{1}{4d}\right)^{17\epsilon \sqrt{\frac{\log n}{\log\log n}}} \geq \left(\frac{1}{d}\right)^{20\epsilon \sqrt{\frac{\log n}{\log\log n}}}. \end{aligned}$$

Now we bound the probability $P_2$ that $v$ gets the message via a push transmission. Let $x'_0, \ldots, x'_{c \log n - 1}$ denote the number of neighbours of $v$ in $N(v, V') \cup K_{i_j} \setminus K'_{i_j}$ which push the message (to any of their neighbours) in steps $1, \ldots, c \log n - 1$, respectively. We know that

$$\sum_{i=0}^{c \log n - 1} x'_i \leq 4\epsilon \cdot dr + r/2 \cdot c \log n = 5\epsilon \cdot dr.$$

The probability that $v$ does not receive the message by `push` in step $i$ is $(1 - 1/d)^{x'_i}$. Hence,

$$P_2 \geq \prod_{i=0}^{c \log n - 1} \left(1 - \frac{1}{d}\right)^{x'_i} = \left(1 - \frac{1}{d}\right)^{5\epsilon \cdot dr} \geq \left(\frac{1}{e}\right)^{-5\epsilon \sqrt{\frac{\log n}{\log\log n}}}.$$

Putting everything together, we obtain that all nodes of $K'_{i_j}$ remain uninformed with probability

$$(P_1 \cdot P_2)^{\frac{1}{2} \cdot \sqrt{\frac{\log n}{\log\log n}}} \geq \left(\frac{1}{d}\right)^{10\epsilon \frac{\log n}{\log\log n}} \cdot \left(\frac{1}{e}\right)^{5\epsilon \frac{\log n}{\log\log n}} > d^{-11\epsilon \frac{\log n}{\log\log n}}.$$

If we now choose $\epsilon$ small enough (i.e. $\epsilon < 1/45$), then all nodes of $K'_{i_j}$ remain uninformed with probability at least $1/\sqrt{n}$. Since there are $\ell/2 = n/((r+1) \cdot (\log^2 n - r)/2)$ complete graphs containing $\sqrt{\log n / \log\log n}/2$ good nodes, the theorem follows. □

# 3 Conclusion

In this paper we presented upper and lower bounds for broadcasting in Edge-Node Expanders. The results of this paper (together with the results of [15]) show that choosing *different* neighbours is very important to save on broadcast communication. In this sense, model RR can be regarded as more advantageous than RANDOM[$c$], which provides further evidence for the power of memory in randomised broadcasting.

# References


[1] A.-L. Barabási and R. Albert. Emergence of Scaling in Random Networks *Science*, 286, 1999.

[2] P. Berenbrink, R. Elsässer, T. Friedetzky. Efficient Randomised Broadcasting in Random Regular Networks with Applications in Peer-to-Peer Systems. In *Proc. of PODC'08*, pages 155–164, 2008.

[3] P. Berenbrink, R. Elsässer, T. Sauerwald. Communication Complexity of Quasirandom Rumor Spreading. In *Proc. of ESA'10*, pages 134–145, 2010.

[4] P. Berenbrink, R. Elsässer, T. Sauerwald. Randomised Broadcasting: Memory vs. Randomness. In *Proc. of LATIN'10*, pages 306–319, 2010.

[5] S. Boyd, A. Ghosh, B. Prabhakar, and D. Shah, Randomized Gossip Algorithms. *IEEE Transactions on Information Theory and IEEE/ACM Transactions on Networking*, 52:2508–2530, 2006.

[6] H. Chernoff. A measure of asymptotic efficiency for tests of a hypothesis based on the sum of observations. *The Annals of Mathematical Statistics*, 23:493–507, 1952.

[7] F.R.K. Chung. *Spectral Graph Theory*. American Mathematical Society, 1985.

[8] 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 PODC'87*, pages 1–12, 1987.

[9] B. Doerr, M. Fouz, T. Friedrich. Social networks spread rumors in sublogarithmic time, In *Proc. of STOC'11*, pages 21–30, 2011.

[10] B. Doerr, T. Friedrich, T. Sauerwald. Quasirandom Rumor Spreading In *Proc. of SODA'08*, pages 773–781, 2008. (full version available at: `http://arxiv.org/abs/1012.5351`)

[11] B. Doerr, T. Friedrich, T. Sauerwald. Quasirandom rumor spreading: expanders, push vs. pull, and robustness *Proc. of ICALP'09, track A*, pages 366–377, 2009.

[12] R. Elsässer. On the communication complexity of randomized broadcasting in random-like graphs. In *Proc. of SPAA'06*, pages 148–157, 2006.

[13] R. Elsässer and A. Ogierman. Efficient Broadcasting in Random Power Law Networks. In *Proc. of WG'10*, pages 279–291, 2010.

[14] R. Elsässer and T. Sauerwald. Broadcasting vs. mixing and information dissemination on Cayley graphs. In *Proc. of STACS'07*, pages 163–174, 2007.

[15] R. Elsässer and T. Sauerwald. The power of memory in randomized broadcasting. In *Proc. of SODA'08*, pages 218–227, 2008.

[16] P. Erdős and A. Rényi. On random graphs I. *Publ. Math. Debrecen*, 6:290–297, 1959.

[17] P. Erdős and A. Rényi. On the evolution of random graphs. *Publ. Math. Inst. Hungar. Acad. Sci.*, 5:17–61, 1960.

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

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

[20] T. Hagerup and C. Rüb. A guided tour of Chernoff bounds. *Information Processing Letters*, 36(6):305–308, 1990.

[21] N. Kahale. Eigenvalues and expansion of regular graphs. *Journal of the ACM*, 42:1091–1106, 1995.

[22] R. Karp, C. Schindelhauer, S. Shenker, and B. Vöcking. Randomized rumor spreading. In *Proc. of FOCS'00*, pages 565–574, 2000.

[23] A.-M. Kermarrec, L. Massouli, and A. J. Ganesh. Probabilistic reliable dissemination in large-scale systems. *IEEE Transactions on Parallel and Distributed Systems*, 14(3):248–258, 2003.

[24] M. Krivelevich and B. Sudakov. Pseudo-random graphs. *More Sets, Graphs and Numbers, Bolyai Society Mathematical Studies 15*, pages 199–262, 2006.

[25] C. McDiarmid. On the method of bounded differences. In *Surveys in Combinatorics*, London Math. Soc. Lectures Notes 141, Cambridge Univ. Press, Cambridge 1989, 148-188.

[26] R. Melamed and I. Keidar. Araneola: A scalable reliable multicast system for dynamic environments. In *Proc. of NCA'04*, pages 5–14, 2004.

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

[28] M. Tanner. Explicit concentrators from generalized n-gons. *SIAM J. Algebraic and Discrete Methods*, 5:287–293, 1984.