

Quasirandom Rumor Spreading*

Benjamin Doerr[†]

Tobias Friedrich[†]

Thomas Sauerwald[‡]

Abstract

We propose and analyse a quasirandom analogue to the classical push model for disseminating information in networks (“randomized rumor spreading”).

In the classical model, in each round each informed node chooses a neighbor at random and informs it. Results of Frieze and Grimmett (Discrete Appl. Math. 1985) show that this simple protocol succeeds in spreading a rumor from one node of a complete graph to all others within $\mathcal{O}(\log n)$ rounds. For the network being a hypercube or a random graph $G(n, p)$ with $p \geq (1+\varepsilon)(\log n)/n$, also $\mathcal{O}(\log n)$ rounds suffice (Feige, Peleg, Raghavan, and Upfal, Random Struct. Algorithms 1990).

In the quasirandom model, we assume that each node has a (cyclic) list of its neighbors. Once informed, it starts at a random position of the list, but from then on informs its neighbors in the order of the list. Surprisingly, irrespective of the orders of the lists, the above mentioned bounds still hold. In addition, we also show a $\mathcal{O}(\log n)$ bound for sparsely connected random graphs $G(n, p)$ with $p = (\log n + f(n))/n$, where $f(n) \rightarrow \infty$ and $f(n) = \mathcal{O}(\log \log n)$. Here, the classical model needs $\Theta(\log^2(n))$ rounds.

Hence the quasirandom model achieves similar or better broadcasting times with a greatly reduced use of random bits.

1 Introduction

1.1 Randomized Broadcast in Networks. The study of information spreading in large networks has various fields of applications in distributed computing. One important example is the maintenance of replicated databases on name servers in a large network [7, 13]. There are updates injected at various nodes, and these updates must be propagated to all the nodes in the network. In each step, two neighboring nodes check whether their copies of the database agree and perform the updates, if necessary. In order to be able to let all copies of the database converge to the same content, efficient broadcasting algorithms have to be developed. Typically, these algorithms should be simple, resilient against failures and should work locally, i. e., the nodes

do not have any knowledge of the global topology.

One such broadcasting protocol is the so-called push model. Initially, only one node of a graph $G = (V, E)$ owns a piece of information (or equivalently, knows a rumor) which is spread iteratively to all other nodes: in each time-step $t = 1, 2, \dots$ every *informed* node chooses a neighbor uniformly at random, to which the piece of information is sent to. The crucial question is how many time-steps are required such that all nodes become informed (with high probability).

Pittel [20] proved that with a certain probability a piece of information is spread to all nodes by the push algorithm within $\log_2 n + \ln n + \mathcal{O}(1)$ steps in a complete graph K_n , tightening a former result by Frieze and Grimmett [15]. Feige, Peleg, Raghavan, and Upfal [13] were the first giving bounds which are valid for general graphs. Moreover, they gave asymptotically tight upper bounds on the runtime for hypercubes and random graphs.

1.2 Quasirandom Push Model. In this work, we propose a quasirandom analogue of the randomized model described above. The basic setup is as in the randomized push model, that is, in each time-step each informed node tries to inform one of its neighbors. However, the choices of these neighbors will not be independently at random. Instead, we assume that each node has a list of his neighbors and informs the neighbors in the order of the list.

It is easily seen that in this model without any randomness a bad choice of the lists can lead to a bad behavior of the protocol. Consider, e. g., the complete graph on n vertices labeled 1 to n and assume that each node informs its neighbors in increasing order. Then it takes $n - 1$ time-steps to spread a rumor from node n to all others (in time-step i , all informed vertices inform node i).

To avoid such behavior, we allow a little randomness. When a node receives the rumor for the first time, it chooses a random position on his list. In the sequel, it informs its neighbors starting with this position and then continuing in the order of the list. When the end of the list is reached, it continues at the beginning of the list.

We call this model *quasirandom push model*, as it

*This work was partially supported by German Science Foundation (DFG) Research Training Group GK-693 of the Paderborn Institute for Scientific Computation (PaSCo) and by the Integrated Project IST-15964 “Algorithmic Principles for Building Efficient Overlay Networks” (AEOLUS) of the European Union.

[†]Department 1: Algorithms and Complexity, Max-Planck-Institut für Informatik, Campus E1 4, 66123 Saarbrücken, Germany

[‡]Faculty of Computer Science, Electrical Engineering and Mathematics, Fürstenallee 11, 33102 Paderborn, Germany

aims at imitating properties of the classical push model with a much smaller degree of randomness. In our analysis, we adopt a worst-case view, that is, we prove bounds for the broadcast times independent of the particular lists. Hence in a practical application, the lists may be chosen to suit internal technical representations of the network.

1.3 Our Results. As previous work did for the random push model, we analyse how long it takes to spread a rumor from one node of a network to all other nodes. Surprisingly, the greatly reduced degree of randomness does not make broadcasting less efficient. For complete graphs, hypercubes and random graphs $\mathcal{G}(n, p)$, $p \geq (1 + \varepsilon)(\log n)/n$, we also obtain a bound of $\mathcal{O}(\log n)$ transmission rounds. These bounds hold for all starting vertices and all orders of the lists.

Our $\mathcal{O}(\log n)$ bound also holds for sparsely connected random graphs $G \in \mathcal{G}(n, p)$ if $p \geq c_n \log(n)/n$ and $(c_n - 1) \log(n) \rightarrow \infty$. This contrasts with the $\Omega(\log^2 n)$ bound shown by Feige et al. [13] for the case that $c_n = 1 + \mathcal{O}(\log \log n / \log n)$ and shows a further superiority of our model (in addition to the reduced need of random bits).

We also prove tight upper bounds of $\Delta \cdot \text{diam}(G)$ and $2n - 3$ for general graphs, which are again better than the corresponding bounds of [13] for the random model. All bounds at a glance are summarized in Table 1.

1.4 Related Work. This work is on quasirandom broadcasting in the push model. In this subsection, we mention a few results related to the push model and the concept of quasirandomness.

While the focus of the papers cited in the first subsection and of this one is only on the time needed to broadcast a rumor to all nodes of a network, one might also want to minimize the number of messages needed to do so. This aspect was regarded by Karp, Schindelhauer, Shenker, and Vöcking [16]. Amongst other results, they combined the push algorithm with the so-called pull algorithm and gave a distributed termination mechanism ensuring that on *complete graphs* only $\Theta(n \log \log n)$ messages are generated.

This analysis has been recently extended to certain random graphs [9]. In the random graph model considered there, every edge between two vertices exists independently with probability p . Note that these graphs serve as a model for peer to peer networks, e.g. [3].

In a very recent work [11], a slightly different algorithm was introduced, where each vertex may choose 4 *different* neighbors for push and pull transmission. Surprisingly, this minor change in the ability of the vertices leads to an exponential decrease in the number of trans-

missions which reduces to $\Theta(n \log \log n)$.

A similar, but continuous-time model is the so-called Richardson's growth-model (often also termed as first-passage percolation [14]) serving as a simple model for the spread of disease. In this model, runtime analysis has been often focused on hypercubes [1, 14] which is closely related to the analysis of random subgraphs of hypercubes. Besides the time after all nodes are infected, also the time until the opposite node becomes infected has been studied.

Quasirandomness means that we try to imitate a particular property of a random process deterministically. This concept occurs in several areas of mathematics and computer science. A prominent example are low-discrepancy point sets and Quasi-Monte Carlo Methods (see e.g. Niederreiter [19]), which proved to be superior over random sample points in numerical integration.

An example closer related to our work is a quasirandom analogue of random walks introduced by Priezzhev, Dhar, Dhar, and Krishnamurthy [21] and later popularized by Jim Propp. Here the vertices are equipped with a rotor pointing to a neighbor and 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, e.g., Cooper and Spencer [4] show 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 a vertex at some time deviates from the expected number had the population done a random walk instead, by only a constant c_d . This constant is independent of the number of particles and their initial position. For the case $d = 1$, that is, the graph being the infinite path, the constant c_1 is approximately 2.29 [5]. For the two-dimensional grid the constant is $c_2 \approx 7.87$ [8]. It is also known that for the graph being an infinite k -ary tree ($k \geq 3$), the deviation between both models can be unbounded [6].

The quasirandomness in our broadcasting model lies in the property that a vertex in the long run contacts each of its neighbors approximately equally often, similar to what would have happened in the random push model. In a sense, and this is typical for quasirandomness, we do better in that the deviations are at most one, whereas in the random push model a vertex v after k contacts would have contacted each neighbor only $k/\text{deg}(v) \pm \Theta(\sqrt{k/\text{deg}(v)})$ times.

Graph class	Broadcasting times
General graphs	$\mathcal{R}(G) = \mathcal{O}(\Delta \cdot (\text{diam}(G) + \log n))$ w. h. p. ¹ [13]
	$\mathcal{Q}(G) \leq \Delta \cdot \text{diam}(G)$ w. p. 1 (Theorem 1)
	$\mathcal{R}(G) \leq 12n \log n$ w. h. p. [13]
	$\mathcal{Q}(G) \leq 2n - 3$ w. p. 1 (Theorem 1)
Complete k -ary trees	$\mathcal{R}(G) = \Theta(k \log n)$ w. h. p. (Section 3)
	$\mathcal{Q}(G) = \Theta(k \log n / \log k)$ w. p. 1 (Section 3)
Hypercubes	$\mathcal{R}(G) = \Theta(\log n)$ w. h. p. [13]
	$\mathcal{Q}(G) = \Theta(\log n)$ w. h. p. (Theorem 2)
Random graphs $\mathcal{G}(n, p)$ with $p = (\log n + f(n))/n$, where $f(n) \rightarrow \infty$ and $f(n) = \mathcal{O}(\log \log n)$	$\mathcal{R}(G) = \Theta(\log^2 n)$ w. p. $1 - o(1)$ [13]
	$\mathcal{Q}(G) = \Theta(\log n)$ w. p. $1 - o(1)$ (Theorem 3)
Random graphs $\mathcal{G}(n, p)$ with $p \geq (1 + \varepsilon) \log(n)/n$, $\varepsilon > 0$ (including complete graph)	$\mathcal{R}(G) = \Theta(\log n)$ w. p. $1 - o(1)$ [13, 20]
	$\mathcal{Q}(G) = \Theta(\log n)$ w. p. $1 - o(1)$ (Theorem 3)

Table 1: Broadcasting times of different graphs G in the random ($\mathcal{R}(G)$) and quasirandom ($\mathcal{Q}(G)$) push model (cf. Definition 1).

2 Precise Model, Notations and Preliminaries

As in the classical push model we aim at spreading a rumor in an undirected graph $G = (V, E)$. By $n := |V|$ we shall always denote the number of vertices of the graph (=number of nodes of the network) considered.

Each vertex $v \in V$ is associated with a cyclic permutation $\pi_v: N(v) \rightarrow N(v)$ of its neighbors (usually simply viewed as list of neighbors). While above we said that a vertex when it first obtains the rumor has chosen a position on the list uniformly at random as starting point for its broadcasting campaign, in the analyses the following equivalent model will be advantageous. We assume that initially each vertex has a position on the list chosen uniformly at random, and that it updates this position each time-step even if it is not informed (“ever rolling lists”). More precisely, at the start of the protocol each vertex chooses an *initially contacted* neighbor i_v uniformly at random from $N(v)$. In each time-step $t = 1, 2, \dots$, the vertex v sends the rumor to vertex $\pi_v^{t-1}(i_v)$, if it is informed, and does nothing otherwise. In the first case, $\pi_v^{t-1}(i_v)$ becomes informed (if it is was not already).

The focus of our investigation is how long it takes until some rumor known only by a single vertex is broadcasted to all other vertices. We adopt a worst-case view in that we aim at bounds that are independent of the starting vertex and of all the lists.

Given a graph $G = (V, E)$, the number of iterations

(or time-steps) of a broadcasting procedure until the rumor reaches all the vertices of G is a random variable that depends on the topology of G .

DEFINITION 1. *Let $\mathcal{R}(G)$ be the number of iterations of the random push model until all vertices in G receive the rumor. Analogously, let $\mathcal{Q}(G)$ be the maximal number of iterations of the quasirandom push model until all vertices in G receive the rumor for all starting vertices and all possible lists.*

Note that both random variables are defined on different probability spaces. Our aim is to bound $\mathcal{Q}(G)$ and to compare it with $\mathcal{R}(G)$ for different graph classes.

In the analysis of the quasirandom push model, it will occasionally be convenient to assume that a vertex after receiving the rumor does not transfer it on for a certain number of time-steps (delayed model). By the assumption of ever rolling lists, it is clear that this will only result in other vertices receiving the rumor later. Consequently, the random variable describing the broadcast time of this model strictly dominates $\mathcal{Q}(G)$. Of course, this also holds if several vertices delay the propagation of the rumor.

LEMMA 1. *The random variable describing the broadcast time of the quasirandom push model with arbitrary delays dominates $\mathcal{Q}(G)$.*

Occasionally, we shall need a notation for the fact that some vertex informs another one would have done so if it had the rumor in time. In this case we say the first vertex contacts the second (c.f. [10]). More precisely, a vertex $u_1 \in V$ *contacts* another vertex $u_m \in$

¹w. p. stands for “with probability”. w. h. p. stands for “with high probability”, which refers to an event which holds with probability at least $1 - \mathcal{O}(n^{-1})$.

V within the time-interval $[a, b]$, if there is a path $(u_1, u_2, \dots, u_{m-1}, u_m)$ in G and $t_1 < t_2 < \dots < t_{m-1} \in [a, b]$ such that for all $j \in [1, m-1]$, $\pi_{u_j}^{t_j-1}(i_{u_j}) = u_{j+1}$.

Graph theoretical notation: Throughout the paper, we use the following notation. For a vertex v of a graph $G = (V, E)$, we denote by $N(v) := \{u \in V \mid \{u, v\} \in E\}$ the set of its neighbors and by $\deg(v) := |N(v)|$ its degree. For any $S \subseteq V$, let $\deg_S(v) := |N(v) \cap S|$. Let $\Delta := \max_{v \in V} \deg(v)$ be the *maximum degree*. The *distance* $\text{dist}(x, y)$ between vertices x and y is the length of the shortest path from x to y . The *diameter* $\text{diam}(G)$ of a connected graph G is the greatest distance between any two vertices in G .

All logarithms $\log n$ are natural logarithms to the base e in the following. As we are only interested in the asymptotic behavior, we will sometimes assume that n is sufficiently large. We should also mention that in order to simplify the presentation we will usually not try to minimize the constant factors in our runtime analysis.

3 General Results

In this section, we give two bounds for the broadcast time in general graphs. The corresponding bounds for the random model are $\mathcal{R}(G) = \mathcal{O}(\Delta \cdot (\text{diam}(G) + \log n))$ and $\mathcal{R}(G) \leq 12n \log n$ w. h. p. [13].

THEOREM 1. *For any graph $G = (V, E)$,*

- (i) $\mathcal{Q}(G) \leq \Delta \cdot \text{diam}(G)$ w. p. 1,
- (ii) $\mathcal{Q}(G) \leq 2n - 3$ w. p. 1.

Proof. (i) Let $\mathcal{P} = (u = u_0, u_1, \dots, v = u_{\text{dist}(u,v)+1})$ be a shortest path from u to v . Clearly, u_1 becomes informed after at most $\deg(u_0) \leq \Delta$ time-steps and inductively the claim follows.

(ii) Let $v \in V$ and let $\mathcal{P} = (u = u_0, u_1, \dots, v = u_\ell)$ be a shortest path from u to v . Then, as observed already in [13], any vertex w not lying on \mathcal{P} has at most three neighbors on \mathcal{P} , and these are contained in $\{u_{i-1}, u_i, u_{i+1}\}$ for some $i \in [1, \ell-1]$. If some w not lying on \mathcal{P} has exactly three neighbors u_{i-1}, u_i, u_{i+1} on \mathcal{P} , we call it a counterfeit of u_i (as u_i and w have, apart from themselves, exactly the same neighbors on \mathcal{P}). Denote by $C(u_i)$ the set of counterfeits of u_i . Without loss of generality, we may choose \mathcal{P} in such a way that for all $i \in [1, \ell-1]$, u_i is informed not later than any of its counterfeits.

Note also that any vertex u_i on the path has only u_{i-1} and u_{i+1} (if existent) as neighbors on the path.

Let t_i denote the time that vertex u_i becomes informed. Then, clearly, $t_0 = 0$. By definition of the Propp machine and choice of \mathcal{P} , we have $t_1 \leq t_0 + |N(u_0) \setminus C(u_1)| = t_0 + |N_{V \setminus \mathcal{P}}(u_0)| + 1 - |C(u_1)|$. For $2 \leq i \leq \ell-1$, similarly, we have $t_i \leq t_{i-1} + |N(u_{i-1}) \setminus$

$C(u_i)| = t_{i-1} + |N_{V \setminus \mathcal{P}}(u_{i-1})| + 2 - |C(u_i)|$. Finally, $t_\ell \leq t_{\ell-1} + |N_{V \setminus \mathcal{P}}(u_{\ell-1})| + 2$. We conclude

$$t_\ell \leq \sum_{i=0}^{\ell-1} |N_{V \setminus \mathcal{P}}(u_i)| - \sum_{i=1}^{\ell-1} |C(u_i)| + 2\ell - 1.$$

Now each vertex w not lying on \mathcal{P} can contribute at most 2 to the above expression (if it has three neighbors on \mathcal{P} , then it is also a counterfeit). Hence $t_\ell \leq 2(n - \ell - 1) + 2\ell - 1 = 2n - 3$. \square

The first bound is asymptotically tight for any constant-degree graphs (including e.g. two- or three-dimensional meshes), as $\Omega(\text{diam}(G) + \log n)$ is an obvious lower bound for every graph. Furthermore, a path of length $n - 1$ fulfills the second bound of Theorem 1 with equality.

Compared to the general upper bound on the random push model of $\mathcal{R}(G) = \mathcal{O}(\Delta(\text{diam}(G) + \log n))$ w. h. p. [13], the quasirandom model may save a small factor on graphs G with $\text{diam}(G) = o(\log n)$. Roughly, the additional factor is caused by the Coupon Collector problem [18]. On complete k -ary trees T the quasirandom model outperforms the random model by a factor of $\log k$ as it is easy to see that $\mathcal{R}(T) = \Theta(k \log n)$ w. h. p. and $\mathcal{Q}(T) = \Theta(k \log n / \log k)$ w. p. 1.

4 Quasirandom Broadcast on Hypercubes

We now focus on the hypercube as this is an important network for parallel computation. Let $H = (V, E)$ denote the d -dimensional hypercube, where $V = \{0, 1\}^d$, $d = \log_2 n$ and $E = \{\{u, v\} \mid \|u - v\|_1 = 1\}$.

Theorem 1 gives an upper bound of $\mathcal{Q}(H) = \mathcal{O}(\log^2 n)$ w. p. 1. The following result shows that this is the best possible upper bound if we insist on probability 1.

PROPOSITION 1. *For the hypercube H with n vertices, $\mathcal{Q}(H) = \Theta(\log^2 n)$ holds with non-zero probability.*

Proof. We prove that there exist lists and initially contacted neighbors for each vertex such that $\Omega(\log^2 n)$ steps are required to inform all vertices independent of the initially informed vertex. For any vertex $x \in \{0, 1\}^d$ and $i \in [1, d]$ let $x(i)$ be the vertex obtained by flipping the i -th bit of x . Then, for any vertex x we choose the neighbor list to be $(x(1), x(2), \dots, x(d))$ and the initially contacted neighbor to be $x(1)$. Assume that initially the vertex $s = (s_1, \dots, s_d)$ owns the rumor. Due to the construction, an arbitrary vertex v requires k steps to send the information to neighbor $v(k)$ and by simple induction we require $\sum_{k=1}^d k = \Omega(d^2) = \Omega(\log^2 n)$ steps to inform $\bar{s} = (1 - s_1, \dots, 1 - s_d)$. \square

For the random push model it is known that it informs w. h. p. each vertex in $\mathcal{R}(H) = \Theta(\log n)$ steps [13]. The following theorem proves that the quasirandom model also has a runtime of $\mathcal{Q}(H) = \Theta(\log n)$ w. h. p.

THEOREM 2. *For the hypercube H with n vertices, $\mathcal{Q}(H) = \Theta(\log n)$ w. h. p.*

Proof. By symmetry we may assume that $u = 0^d$ knows a rumor at the beginning. Our proof consists of three stages. In the first stage we show that after $\mathcal{O}(d)$ steps a large set of informed vertices I' which does not contain vertices which are too close to each other exists. Similarly, a large set of uninformed vertices U must exist in order to keep a fixed vertex w at step $\mathcal{O}(d)$ uninformed. In particular, every vertex of I' will be close to some proper vertex $u \in U$. Finally, we show that one of the informed vertices in I' informs one close vertex of U with high probability implying that w is also informed after $\mathcal{O}(d)$ steps. A graphical illustration of our proof can be found in Figure 1. Before going into the details we should remark that for the sake of readability we will frequently use real quantities where integers are required.

Forward Approximation: (from step 0 till step $4d$) We first show that after $4d$ steps a large set of informed vertices exists. Let L_i be the set of vertices with $\|x\|_1 = i$. Note that after $2d$ steps, $L_0 \cup L_1$ has been informed completely.

Consider some time-step t and denote by I_t the set of informed vertices. We may assume that all initially contacted neighbors of $I_t \cap L_i$ are still to be chosen u. a. r. Notice that the set of edges between $I_t \cap L_i$ and L_{i+1} satisfy $|E(I_t \cap L_i, L_{i+1})| = \sum_{v \in L_{i+1}} \deg_{I_t \cap L_i}(v) = |I_t \cap L_i| \cdot (d - i)$. Our goal is to show that a large set of vertices in L_{i+1} will be also informed after $\mathcal{O}(1)$ additional steps. The probability that a vertex $v \in L_{i+1}$ becomes not informed after 10 steps is

$$\begin{aligned} \Pr[v \notin I_{t+10}] &\leq \prod_{u \in N(v) \cap I_t \cap L_i} \left(1 - \frac{10}{d}\right) \\ &= \left(1 - \frac{10}{d}\right)^{\deg_{I_t \cap L_i}(v)}. \end{aligned}$$

By linearity of expectations we get

$$\begin{aligned} \mathbf{E}[|I_{t+10} \cap L_{i+1}|] &= \sum_{v \in L_{i+1}} \Pr[v \in I_{t+10}] \\ &\geq \sum_{v \in L_{i+1}} 1 - \left(1 - \frac{10}{d}\right)^{\deg_{I_t \cap L_i}(v)} \\ &\geq \sum_{v \in L_{i+1}} 1 - e^{-\frac{10 \deg_{I_t \cap L_i}(v)}{d}}. \end{aligned}$$

Let us assume in the following that $i \leq \frac{d}{7}$. Then due to $\deg_{I_t \cap L_i}(v) \leq i+1$ and $1 + \frac{x}{2} \geq \exp(x)$ for $-1.5 < x < 0$ we get

$$\begin{aligned} \mathbf{E}[|I_{t+10} \cap L_{i+1}|] &\geq \sum_{v \in L_{i+1}} \frac{10 \deg_{I_t \cap L_i}(v)}{2d} \\ &= |I_t \cap L_i| \frac{(d-i)10}{2d} \geq 4.25 |I_t \cap L_i|. \end{aligned}$$

Let $f: N(u_1) \times N(u_2) \times \dots \times N(u_{|I_t \cap L_i|}) \rightarrow \mathbb{N}$ describe the random variable of $|I_{t+10} \cap L_{i+1}|$ depending on the choices of the initially contacted neighbors, which are u. a. r. and independent from each other. Since some fixed vertex can only inform at most 10 vertices within 10 steps, f satisfies the Lipschitz condition and the method of independent bounded differences [17] gives

$$\Pr[f \leq \mathbf{E}[f] - t] \leq \exp\left(-\frac{t^2}{2|I_t \cap L_i|10^2}\right)$$

and by setting $t = \frac{1}{4}|I_t \cap L_i|$ we conclude by using $|I_t \cap L_i| \geq \frac{d(d-1)}{2}$ that

$$\Pr[f \leq 4|I_t \cap L_i|] \leq \exp\left(-\frac{\frac{1}{4}(|I_t \cap L_i|)^2}{200|I_t \cap L_i|}\right) \leq 2^{-3d},$$

whenever d is large enough.

Iterating over all levels $0 \leq i \leq d/7$ we require at all $2d + (d/7 - 2) \cdot 10 \leq 4d$ time-steps to get w. p. $1 - d2^{-3d}$ that

$$|I_{4d} \cap L_{d/7}| \geq \frac{d(d-1)}{2} 4^{d/7-2} \geq 4^{d/7}.$$

Consider now the following iterative procedure (cf. [13]) of transforming the set I_{4d} into another set I'_{4d} which is empty at the beginning. As long as I_{4d} is nonempty, we first pick an arbitrary vertex, say v , of I_{4d} and put v into I'_{4d} . Then we remove v of I_{4d} together with all vertices $u \in I_{4d}$ such that $\text{dist}(u, v) \leq d/64$ and go to the next iteration. Note that the vertices of I'_{4d} have a pairwise distance of at least $d/64$. We may also conclude that I'_{4d} is of size

$$\frac{4^{d/7}}{\sum_{k=0}^{d/64} \binom{d}{k}} \geq \frac{4^{d/7}}{(64e)^{d/64}} \geq \left(\frac{11}{10}\right)^d,$$

where we have used the inequality $\sum_{i=0}^m \binom{n}{i} \leq \left(\frac{en}{m}\right)^m$.

Backward Approximation: (from step $6d$ till step $1542d$) The key idea is now to analyse the propagation of the rumor in the reverse order. Roughly, we will show that to keep a vertex w uninformed at some time-step $t' = \Theta(d)$, also a lot of other vertices have to be uninformed at time-step $t' - \mathcal{O}(d)$. More formally, these

vertices contact the vertex w within the time-interval $[t' - \mathcal{O}(d), t']$.

Again due to the symmetry of H , we may restrict our attention to the vertex $w = 1^n$. Similar to the definition of I_t , denote by U_t the set of vertices which contact the vertex w within the time-interval $[t, 1542d]$, where $t \leq 1542d$. We will now show that U_{6d} contains some vertex v such that $\text{dist}(0^d, v) \leq d/256$. Let us assume that for some time-step t , $U_t \cap L_i \neq \emptyset$ for some $d/256 \leq i \leq d$. We will bound the time $t' \leq t$ after $U_{t'} \cap L_{i-1} \neq \emptyset$. Observe that there are i many vertices in L_{i-1} which have an edge to some fixed vertex $v' \in L_i$. Due to the “ever rolling lists”, we may assume that for a vertex in L_{i-1} the number of steps till it contacts vertex v' is uniformly distributed in $[1, d]$. Note that this uniform distribution is stochastically smaller than the exponential distribution with mean d plus 2. Therefore, we may assume that this number of steps is exponentially distributed with parameter $1/d$ (expected value d). Furthermore, recall that the minimum of i exponentially distributed random variables with expectation d is $d/i \leq 256$ [18]. Thus, we may upper bound the random variable X until a vertex in $L_{d/256}$ contacts v' by a sum of $\frac{255}{256}d$ independent exponentially distributed random variables X_i with mean 256. As the moment-generating function for X_i equals $\mathbf{E}[e^{tX_i}] = \frac{1/256}{1/256-t}$, where $-1/256 \leq t \leq 1/256$, we may use a Chernoff bound [18] for $X = \sum_{i=1}^{255/256d} X_i$ to obtain

$$\begin{aligned} \Pr[X \geq 1536 \cdot d] &\leq \frac{\mathbf{E}[e^{tX}]}{e^{t \cdot 1536d}} = \frac{\prod_{k=1}^{255/256d} \frac{1/256}{1/256-t}}{e^{t \cdot 1536d}} \\ &\stackrel{t=1/512}{\leq} \frac{2^d}{e^{1/512 \cdot 1536d}} \leq 2^{-5/2d}, \end{aligned}$$

where the first equality is due to the independence of the X_i .

Hence, a vertex with distance $d/256$ from 0^n is in U_{6d} w. p. $1 - 2^{-5/2d}$. Notice that we may replace 0^n by any other vertex and just ignore any bits which are already ones. With this we can conclude that from each vertex there exists another vertex of distance at most $d/256$ which lies in the set U_{6d} w. p. $1 - 2^{-3/2d}$. Recall again that due to the symmetry of H the vertex w could have been replaced by any other vertex of the graph.

Coupling: (from step $4d$ till step $6d$) From the first part of this analysis we know that at time-step $4d$ there exists a certain, large enough set of informed vertices I'_{4d} such that all vertices have a pairwise distance of at least $d/64$. We just have seen that for all vertices $v \in V$ there is at least one vertex $u(v) \in U_{6d}$ such that $\text{dist}(v, u(v)) \leq d/256$. Therefore, there exists a bijection $\Phi: I'_{4d} \rightarrow U_{6d}$ such that for all $v \in I'_{4d}$ it holds $\text{dist}(v, \Phi(v)) \leq d/256$. It remains to show that at least




Figure 1: A Sketch of the proof of Theorem 2. The circles represent $I'(4d)$, the rectangles the corresponding v'' and the triangles represent $\Phi(I'(4d))$.

one vertex $v \in I'_{4d}$ will contact $\Phi(v)$ within the time interval $[4d + 1, 6d]$ since this implies that w will be also informed after $1542d$ time-steps.

We now derive the probability that a specific $v \in I'_{4d}$ informs its respective $\Phi(v) \in U_{6d}$. Let $v' \in I_{4d}$ with $\text{dist}(v, v') \leq d/256$ be such that $\text{dist}(v', \Phi(v))$ is minimal. W. l. o. g. $u(v) \notin \Phi(v')$. Let v'' be some proper neighbor of v' which gets informed at time $4d, \dots, 4d + d - 1$ and is closer to $\Phi(v)$. The first neighbor to which v'' sends the rumor is u. a. r. and decreases the distance to $u(v)$ with probability $\text{dist}(v'', \Phi(v))/d$. Iterating this process and using the fact that $n! \geq (n/3)^n$ for every integer n gives a probability of

$$\prod_{k=1}^{d/256} \frac{k}{d} \geq \frac{d^{d/256}}{(768d)^{d/256}} \geq 768^{-d/256} \geq \left(\frac{11}{12}\right)^d$$

for reaching $\Phi(v)$ before time-step $6d$.

By construction and the fact that we have only considered shortest paths between $v''(v)$ and $\Phi(v)$, the corresponding events w. r. t. each $v''(v)$ and its respective $\Phi(v)$ are independent from each other. Hence, the probability conditioned on the success of the forward and backward approximation that no path succeeds is at most

$$\left(1 - \left(\frac{11}{12}\right)^d\right)^{(11/10)^d} \leq e^{-\left(\frac{121}{120}\right)^d} \leq 2^{-3/2d}.$$

Therefore, with probability at most $3 \cdot 2^{-3/2d} \leq 2^{-2d}$ the vertex w will remain uninformed at time-step $1542d$. Hence, the probability that at this time all vertices are informed is at least $1 - 2^d 2^{-2d} = 1 - 2^{-d}$. Since $\Omega(\log n)$ is obviously a lower bound, as the number of informed vertices can at most double in each step, the claim of the theorem follows. \square

5 Quasirandom Broadcast on Random Graphs

In this section we show that the quasirandom push model has a logarithmic runtime as well on many random graphs. Let $\mathcal{G}(n, p)$ be the probability space of graphs with n vertices and each edge present independently with probability p . We only consider the case that $p \geq c_n \log(n)/n$, where $(c_n - 1) \log(n) \rightarrow \infty$. By [12], such graphs are connected w. p. $1 - o(1)$.

Feige et al. [13] showed for sparse random graphs $G \in \mathcal{G}(n, p)$ where $c = 1 + \mathcal{O}(\log \log n / \log n)$ that the random push model needs w. p. $1 - o(1)$ $\mathcal{R}(G) = \Theta(\log^2 n)$ iterations to spread a rumor to all vertices. They also showed that for denser random graphs G with $p \geq (1 + \varepsilon) \log(n)/n$, for some fixed $\varepsilon > 0$, that w. p. $1 - o(1)$ $\mathcal{R}(G) = \Theta(\log n)$.

Theorem 3 below shows for a much larger class of random graphs that the quasirandom model only requires $\mathcal{Q}(G) = \Theta(\log n)$ w. p. $1 - o(1)$. Note that this also includes the complete graph ($p = 1$) and the space $\mathcal{G}(n, 1/2)$ which chooses uniformly at random from all graphs on n vertices.

THEOREM 3. *For random graphs $G \in \mathcal{G}(n, p)$ with $p \geq c_n \log(n)/n$ where $(c_n - 1) \log(n) \rightarrow \infty$, the quasirandom push model informs all vertices within $\mathcal{Q}(G) = \Theta(\log n)$ steps w. p. $1 - o(1)$.*

Proof. As the theorem covers a large class of random graphs, we sometimes will have to distinguish between “sparse” random graphs with $p = \Theta(\log(n)/n)$ and “dense” random graphs with $p = \omega(\log(n)/n)$.

For dense random graphs, there exist constants $\alpha < 1$ and $\beta > 1$ such that w. p. $1 - o(1)$ the minimum degree is at least $\alpha D(n)$ and the maximum degree is at most $\beta D(n)$ where $D(n) := p(n - 1)$ is the expected degree of a vertex.

For sparse random graphs, Cooper and Frieze [2] showed that w. p. $1 - o(1)$ there are at most $n^{1/3}$ *small vertices*, i. e., of degree $\leq \log n / 20$, and that no two small vertices are within a distance of $\log n / (\log \log n)^2$ or less. A *large vertex* denotes a vertex which is not small. By definition, the degree of all large vertices is at least $\alpha D(n)$ with $\alpha := 1/(20c_n)$ and $D(n)$ defined as above. Analogously to dense random graphs, there is also a constant $\beta > 1$ such that w. p. $1 - o(1)$ the maximum degree of sparse random graphs is at most $\beta D(n)$. In dense random graphs there are no small vertices w. p. $1 - o(1)$.

Similar to the proof of Theorem 2, the analysis consists of a forward approximation, a backward approximation and a coupling stage. First, we will prove that after $t_1 = \mathcal{O}(\log n)$ time-steps there is a linear number of informed vertices which have not informed any other vertex yet. Then, we will (roughly) show that an un-

informed vertex at some time t_2 implies a logarithmic number of uninformed vertices $\mathcal{O}(\log n)$ steps before t_2 . In the final coupling stage we will prove that it is very likely that within a single step one of the linear many informed vertices contacts one of the logarithmic many uninformed vertices.

Forward Approximation: Let u know some rumor at time-step 0. We consider phases of several steps. Let t denote the current phase (and not the time-step). Let I_t be the set of large vertices informed after the t -th phase. Let $N_t \subseteq I_t \setminus I_{t-1}$ be the set of newly informed large vertices which have got the rumor from vertices in N_{t-1} in the $(t - 1)$ -th phase.

First, we show that after at most two phases there are at least $100 \log n$ large vertices informed w. h. p. If u is small, within a single step we inform a large vertex. Hence, assume w. l. o. g. $\deg(u) \geq \alpha D(n)$. For dense random graphs and sufficiently large n we have $\deg(u) > 100 \log n$ and just do $100 \log n + 1$ steps in which $100 \log n + 1$ neighbors of u get informed. This implies $|N_1| = 100 \log n$ w. p. $1 - o(1)$.

For sparse random graphs we have $D(n) = \mathcal{O}(\log n)$. There, a first phase of $\beta D(n)$ steps informs all neighbors of u . In a second phase of $\beta D(n)$ steps, all vertices within distance ≤ 2 from u get informed. It remains to show that at least $100 \log n$ different large vertices got informed w. h. p. The probability that while $\geq (\beta D(n) - 1)^2 = \Theta(\log^2 n)$ times vertices received the rumor from a large vertex, only $100 \log n$ different large vertices were informed is

$$\begin{aligned} & \binom{n}{100 \log n} \left(\frac{100 \log n}{n - 1} \right)^{\Omega(\log^2 n)} \\ & \leq n^{100 \log n} \left(\frac{100 \log n}{n - 1} \right)^{\Omega(\log^2 n)} \leq n^{-2}. \end{aligned}$$

Therefore, $|N_2| > 100 \log n$ w. h. p.

We now assume to have informed $|N_2| > 100 \log n$ large vertices. The following phases last 5 steps each. In each phase every vertex of N_t contacts at least 4 large vertices of N_{t+1} . The probability that a vertex in N_{t+1} is hit by more than one vertex of N_t within one phase is at most $5|N_t|/(n - 1)$. This implies for $|N_t| \leq n/30$ and $|I_t| \leq n/15$,

$$\mathbf{E}[|N_{t+1}|] \geq 4|N_t| \left(1 - \frac{5|N_t|}{n} \right) \left(1 - \frac{1}{15} \right) > 3|N_t|.$$

The random variable $|N_{t+1}|$ depends on the independent choices of the initially contacted neighbors of N_t and on the states of the 5 different chosen neighbors (informed or not informed). We number the vertices of N_t by $1, \dots, |N_t|$ and denote by Y_i the random variable exposing the 5 transmissions of vertex $i \in [1, |N_t|]$.

Then, $Z_i := \mathbf{E}[|N_{t+1}(Y_1, \dots, Y_{|N_t|})| \mid Y_1, \dots, Y_{i-1}]$ is a Doob martingale [18] which exposes the random variable $|N_{t+1}|$ step by step. In particular we have $\mathbf{E}[Z_0] = \mathbf{E}[|N_{t+1}|]$. Furthermore observe that $|Z_i - Z_{i-1}| \leq 5$ (conditioned on Y_1, \dots, Y_{i-2}) for every $i \in [1, |N_t|]$, since every vertex of N_{t+1} sends the rumor to 5 neighbors and the choices of the initially contacted neighbors are independent from each other. Inserting our findings into Azuma-Hoeffding's Inequality [18] yields

$$\Pr[|N_{t+1}| < 2|N_t|] \leq \exp\left(-\frac{|N_t|^2}{50|N_t|}\right) \leq n^{-2},$$

since $|N_t| \geq 100 \log n$.

This implies that w.h.p. there is a time $t_1 = \Theta(\log n)$ with $|N_{t_1}| \geq n/30$.

Backward Approximation: We now analyse the propagation of the rumor in the reverse order. We will show that in order to keep a vertex w uninformed till some time $t_2 = \Theta(\log n)$, w.h.p. there must be at least $40 \log n$ uninformed vertices at time-step $t_2 - \mathcal{O}(\log n)$.

First, we examine dense random graphs. We show that there are w.h.p. at least $40 \log n$ vertices which contact w within the time-interval $(t_2 - 80 \frac{\beta}{\alpha} \log n, t_2]$. The time until a fixed neighbor $x \in N(w)$ contacts the vertex w can be assumed to be uniformly distributed in $[1, \deg(x)]$ due to the ‘‘ever rolling lists’’. We bound the probability that a particular neighbor $x \in N(w)$ contacts w in the time-interval $(t_2 - 80 \frac{\beta}{\alpha} \log n, t_2]$. As we have $80 \frac{\beta}{\alpha} \log n < \alpha D(n) \leq \deg(x)$, this probability is

$$\frac{80 \frac{\beta}{\alpha} \log n}{\deg(x)} \geq \frac{80 \frac{\beta}{\alpha} \log n}{\beta D(n)}.$$

Since there are at least $\alpha D(n)$ such neighbors, the expected number of successful neighbors is at least

$$\frac{80 \frac{\beta}{\alpha} \log n}{\beta D(n)} \alpha D(n) \geq 80 \log n.$$

Let X be random variable describing the number of successful neighbors. We get by a Chernoff bound [18]

$$\Pr[X \leq 40 \log n] \leq e^{-80 \log n/8} \leq n^{-2}.$$

It remains to examine sparse random graphs. If w is a small vertex, it must be contacted by a large vertex within $\beta D(n)$ time-steps. Hence we assume w.l.o.g. that w is large. Within $\beta D(n)$ steps, all $x \in N(w)$ must have contacted w at least once. At least $\alpha D(n) - 1$ of them are large. Within additional $\beta D(n)$ steps, all $y \in N(N(w))$ must have contacted a $x \in N(w)$. Therefore, at time $t_2 - 2\beta D(n)$ all vertices with distance 2 from w must be uninformed in order

to keep w uninformed at time-step t_2 . As this is a set of size $\Omega(\log^2 n)$ w.p. $1 - o(1)$, there are at least such $40 \log n$ uninformed vertices at time $t_2 - 2\beta D(n)$.

Coupling: From the forward approximation we know that at some time t_1 there are $\geq n/30$ newly informed vertices w.h.p. On the other hand, the backward approximation showed that if some vertex v remains uninformed till time-step $t_2 = \Theta(\log n)$, there must be $\geq 40 \log n$ uninformed vertices at some time $t_2 - \mathcal{O}(\log n)$. The probability that none of the $\geq n/30$ newly informed vertices informs any of the $\geq 40 \log n$ uninformed vertices around v within a single step after time t_1 is

$$< \left(1 - \frac{40 \log n}{n-1}\right)^{n/30} < \exp\left(-\frac{40}{30} \log n\right) = n^{-4/3}.$$

Therefore, by a union bound all vertices are informed after $\mathcal{O}(\log n)$ steps w.p. $1 - o(1)$. \square

6 Conclusion

In this paper, we proposed and investigated a quasirandom analogue of the classical push model for spreading a rumor to all vertices of a network.

We showed that for the network topologies of complete graphs, hypercubes and random graphs $\mathcal{G}(n, p)$, where p only needs to be slightly larger than the connectivity threshold, after $\Theta(\log n)$ iterations all vertices are informed with probability $1 - o(1)$. Hence the quasirandom model achieves asymptotically the same bounds as the random one, or even better ones (for random graphs with p close to $\log(n)/n$).

From the methodological point of view, this work is also interesting. Our proofs show, in particular, that the difficulties usually invoked by highly dependent random experiments can be overcome.

From the general perspective of using randomized methods in computer science, our results, as a number of other recent results, can be interpreted in the way that choosing the right dose of randomness might be a fruitful topic for further research.

7 Acknowledgements

We are thankful to the reviewers for helpful comments improving the quality of this paper. Moreover we wish to thank Robert Elsässer and Martin Gairing for various discussions.

References

- [1] B. Bollobás and Y. Kohayakawa. On Richardson's model on the hypercube. In B. Bollobás and A. Thomason, editors, *Combinatorics, Geometry*

- and Probability, pages 129–137. Cambridge University Press, 1997.
- [2] C. Cooper and A. Frieze. The cover time of sparse random graphs. In *14th ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 140–147, 2003.
- [3] C. Cooper, M. Dyer, and C. Greenhill. Sampling regular graphs and peer-to-peer networks. In *16th ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 980–988, 2005.
- [4] J. Cooper and J. Spencer. Simulating a random walk with constant error. *Comb. Probab. Comput.*, 15:815–822, 2006.
- [5] J. Cooper, B. Doerr, J. Spencer, and G. Tardos. Deterministic random walks on the integers. *European Journal of Combinatorics*. To appear, preliminary version available from arXiv:math/0602300.
- [6] J. Cooper, B. Doerr, T. Friedrich, and J. Spencer. Deterministic random walks on regular trees. In *19th ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 766–772, 2008.
- [7] 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 *6th ACM Symposium on Principles of Distributed Computing (PODC)*, pages 1–12, 1987.
- [8] B. Doerr and T. Friedrich. Deterministic random walks on the two-dimensional grid. Available from arXiv:math/0703453.
- [9] R. Elsässer. On the Communication Complexity of Randomized Broadcasting in Random-like Graphs. In *18th ACM Symposium on Parallelism in Algorithms and Architectures (SPAA)*, pages 148–157, 2006.
- [10] R. Elsässer and T. Sauerwald. Broadcasting vs. Mixing and Information Dissemination on Cayley Graphs. In *24th International Symposium on Theoretical Aspects of Computer Science (STACS)*, pages 163–174, 2007.
- [11] R. Elsässer and T. Sauerwald. On the Power of Memory in Randomized Broadcasting. In *19th ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 218–227, 2008.
- [12] P. Erdős and A. Rényi. On random graphs. *Publ. Math. Debrecen*, 6:290–297, 1959.
- [13] U. Feige, D. Peleg, P. Raghavan, and E. Upfal. Randomized broadcast in networks. *Random Structures and Algorithms*, 1(4):447–460, 1990.
- [14] J. Fill and R. Pemantle. Percolation, first-passage percolation and covering times for richardson’s model on the n -cube. *Annals of Applied Probability*, 3:593–629, 1993.
- [15] A. Frieze and G. Grimmett. The shortest-path problem for graphs with random arc-lengths. *Discrete Applied Mathematics*, 10:57–77, 1985.
- [16] R. Karp, C. Schindelhauer, S. Shenker, and B. Vöcking. Randomized Rumor Spreading. In *41st IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 565–574, 2000.
- [17] C. McDiarmid. On the method of bounded differences. In *Surveys in combinatorics, 1989 (Norwich, 1989)*, volume 141 of *London Math. Soc. Lecture Note Ser.*, pages 148–188. Cambridge Univ. Press, Cambridge, 1989.
- [18] M. Mitzenmacher and E. Upfal. *Probability and Computing: Randomized Algorithms and Probabilistic Analysis*. Cambridge University Press, 2005.
- [19] H. Niederreiter. *Random number generation and quasi-Monte Carlo methods*, volume 63 of *CBMS-NSF Regional Conference Series in Applied Mathematics*. SIAM, Philadelphia, PA, 1992.
- [20] B. Pittel. On spreading a rumor. *SIAM Journal on Applied Mathematics*, 47(1):213–223, 1987.
- [21] V. B. Priezzhev, D. Dhar, A. Dhar, and S. Krishnamurthy. Eulerian walkers as a model of self-organized criticality. *Phys. Rev. Lett.*, 77:5079–5082, 1996.